INTERACTIVE FAQ

Harvest Now, Decrypt Later

Understanding the HNDL Threat, Q-Day 2029, and the PQC+ Solution

Table of Contents on Page 3

A practical reference for mid-enterprise leaders, prepared by your software development partner. This document explains the Harvest Now, Decrypt Later (HNDL) threat, the Q-Day 2029 timeline projected by Google, IBM, and Microsoft, the personal liability now facing executives under 2025 DOJ guidelines, and how PQC+ delivers NIST-certified post-quantum protection in 90 days.

WHAT'S INSIDE

- Executive summary of HNDL, Q-Day, and personal liability
- 30 interactive Q&A entries with one-click navigation
- Sector-by-sector exposure analysis
- Executive personal liability under 2025 DOJ guidelines
- The PQC+ solution: architecture, deployment, and ROI
- 90-day implementation path with no "rip and replace"

PQC+™

Transformativ IP

Regulatory Compliance & HNDL Protection
Federal & all 50 States within 90 Days

Executive Summary

The world is in the middle of a global engineering transition from "theoretical research" to "mandatory implementation" in cryptography. The driver is the maturation of quantum computing, and the most immediate consequence is a threat known as **Harvest Now, Decrypt Later (HNDL)** — adversaries are stealing encrypted data today with the intention of decrypting it once quantum computers become powerful enough.

Q-Day — the day RSA and ECC encryption can be broken — is now projected for 2029. This is not a fringe estimate. It is the published projection of **Google, IBM, and Microsoft**, the three largest investors in quantum computing. When all three industry leaders independently arrive at the same timeline, it stops being speculation.

The U.S. National Institute of Standards and Technology (NIST) has finalized new post-quantum cryptography (PQC) standards — ML-KEM, ML-DSA, and SLH-DSA. The European Union requires critical infrastructure to complete its transition by **December 31, 2030**, and all other systems by 2035. Most importantly for leadership: **2025 DOJ guidelines now treat executive inaction on the quantum threat as a possible "willful violation"** rather than negligence — opening the door to personal criminal prosecution, personal civil penalties, and uninsured liability.

Why this matters now

Large enterprises typically take 12 to 15 years to migrate cryptography. If Q-Day arrives in 2029 — five years from now — most organizations face a guaranteed exposure window. Any data harvested today that needs to remain confidential for longer than 5 to 15 years should be treated as already compromised.

The PQC+ solution

PQC+ is a software-based post-quantum cryptography platform designed for **90-day deployment with no "rip and replace."** It uses NIST-certified algorithms (FIPS 203, 204, 205) and goes further — embedding access control rules directly into the cryptographic key structure, binding decryption to specific hardware, and treating consent as a real-time cryptographic attribute. A stolen PQC+ database becomes mathematically useless: no resale value on the dark web, no future decryption by quantum computers, no way for an attacker to monetize what was stolen.

For executives, PQC+ provides what regulators look for in a "good faith" defense: a documented, dated record of diligent action against a known emerging threat, with NIST certifications and a complete compliance dossier. It is the single strongest action available today to shift breach classification away from willful violation.

> **Why this matters now**
>
> Any data encrypted with today's classical algorithms and intercepted today will eventually be decrypted. If your data needs to remain confidential for longer than 5 to 15 years, it is already at risk. The decision is not whether to migrate — it is whether to migrate proactively in 90 days, or reactively after a breach.

Table of Contents

Click any question to jump directly to its answer.

The PQC+ Solution

PQC+ Architecture

PQC+ in Healthcare

Comparison & Strategy

Frequently Asked Questions

Detailed answers covering the HNDL threat, Q-Day 2029, executive liability, and the PQC+ solution.

THE HNDL THREAT — FUNDAMENTALS

Q01. How does the Harvest Now, Decrypt Later threat change today's data security risks?

In short

HNDL turns tomorrow's quantum computers into today's liability — data stolen now will be readable later.

Harvest Now, Decrypt Later (HNDL) changes the rules of data security by separating the moment data is stolen from the moment it is actually read. Rather than treating encryption as a permanent shield, HNDL forces organizations to treat confidentiality as a perishable asset — one whose protection erodes as quantum computing matures.

A new way to think about risk

Traditionally, data was considered safe as long as today's encryption could not be broken by today's technology. HNDL replaces that with a simple, time-based test: if the data needs to stay secret for longer than it will take an adversary to acquire a quantum computer, then the data is already compromised — even if no one can read it yet.

"Zombie" data: information that retains value for decades

HNDL specifically targets data that holds its value over long periods. Examples include:

- **Genomic and healthcare data:** A person's DNA never changes, so its sensitivity never expires.
- **National security and diplomacy:** Diplomatic cables, intelligence identities, and state secrets often require 30 to 50+ years of protection.
- **Financial and blockchain records:** Even if a blockchain updates its signing algorithms, the entire historical ledger remains exposed and can be deanonymized later.

The threat is irreversible

In a traditional breach, you can rotate keys, patch software, or reset credentials. With HNDL, once encrypted data is intercepted and stored, it cannot be "un-stolen." The attacker simply has to be patient. Any organization that waits for a quantum computer to appear before upgrading its cryptography will have already lost control of its historical data.

From active attacks to silent collection

HNDL also shifts the nature of the attack itself. Adversaries don't need to break encryption in real time — they just need to record and store encrypted traffic, which is largely undetectable. Real-world signals, such as internet traffic being rerouted through adversarial jurisdictions, suggest large-scale passive harvesting is already underway.

Plain-language analogy

> Imagine thieves stealing a reinforced safe they cannot yet open. They don't crack it on site — they take the whole safe home and wait for a tool to be invented that can cut through the door. The contents inside are still worth a fortune ten years from now.

The bottom line: HNDL means the quantum threat is deferred, not distant. Any data encrypted with today's classical algorithms is vulnerable if its value extends past the arrival of fault-tolerant quantum computing.

↑ Back to Table of Contents

SECTOR EXPOSURE

Q02. Which sectors face "Critical Exposure" because of indefinite data retention?

> **In short**
>
> Sectors whose data must remain confidential for 30+ years — or forever — face the highest HNDL risk.

"Critical Exposure" applies to sectors that handle data requiring confidentiality for more than 30 years, or indefinitely. In these areas, the lifespan of the data far exceeds the projected arrival of cryptographically relevant quantum computers, which makes retrospective compromise effectively certain if migration is delayed.

Sectors in the Critical Exposure tier

- **State intelligence and national security:** State secrets, intelligence identities, and classified archives often require protection for 30 to over 100 years.
- **Public blockchains:** Distributed ledgers like Bitcoin retain data indefinitely. Even if future transactions use post-quantum cryptography, the entire historical ledger can be harvested today and deanonymized later.
- **Genomic data and child welfare records:** Genomic data is biologically immutable and remains sensitive for the lifetime of the individual and their descendants (50–100 years). Child welfare records carry similarly long protection mandates.

Across these sectors, the HNDL threat ensures that anything encrypted with today's classical algorithms and intercepted now will eventually be compromised. The required secrecy window of 50+ years is simply longer than the 5–15 years experts project before quantum decryption becomes feasible.

↑ Back to Table of Contents

THE HNDL THREAT — FUNDAMENTALS

Q03. How does the Store Now, Decrypt Later (SNDL) threat change today's risk calculations?

In short

SNDL (the same concept as HNDL) reframes confidentiality from a static state into a time-sensitive equation.

Store Now, Decrypt Later (SNDL) — which is another name for Harvest Now, Decrypt Later — changes risk calculations by turning data security from a static condition into a time-dependent one. Instead of focusing solely on preventing intrusion today, organizations must manage how long their data needs to remain confidential against a future capability that experts agree is coming.

1. A time-based test for failure

Traditionally, data was deemed safe as long as standards like RSA-2048 could not be broken by current technology. SNDL introduces a new test: confidentiality fails the moment the required lifespan of the data exceeds the time remaining until a quantum computer arrives.

The practical implication: Security teams must now weigh how long data needs to stay secret (for example, 25 years for state secrets) against the projected arrival of fault-tolerant quantum computers, currently estimated between 2029 and 2035.

2. Irreversible liabilities

SNDL creates breaches that cannot be undone. Once data is harvested and stored, it cannot be retrieved. Upgrading to post-quantum cryptography in the future will not protect data that was stolen before the upgrade. This is especially damaging for data with a long tail of value: genomic information, intelligence, and immutable blockchain history.

3. From active intrusion to passive collection

The threat model shifts from detecting active network intrusions to defending against passive interception. Attackers can record encrypted traffic via fiber-optic taps, satellite interception, or BGP rerouting — all of which are extremely difficult to detect. This creates an asymmetric advantage: attackers invest cheaply in storage, while defenders must invest heavily in migration.

4. The "risk deficit"

A risk deficit occurs when the time required to migrate to quantum-safe standards is longer than the time remaining before quantum computers arrive. Large enterprises may need 12 to 15 years to fully migrate. If quantum computers arrive in 10 years, the organization faces a guaranteed window of exposure for everything transmitted during that gap. In effect, any long-lived data encrypted with classical algorithms today should be treated as potentially public information in the long term.

Plain-language analogy

Traditional security is like guarding a bank vault — if the lock holds, the money stays safe. SNDL is like robbers stealing the entire vault and taking it home. They don't have to open it tonight. They

can sit on it for a decade until a laser is invented that cuts it open. The bank lost the assets the moment the vault left the building.

↑ Back to Table of Contents

ADVERSARY MODEL

Q04. What are the three evolving resources of an HNDL adversary?

In short

Collection, future decryption power, and patience — the three resources that define an HNDL attacker.

The formal model of the HNDL adversary identifies three evolving resources that set them apart from conventional attackers:

- **Collection capability:** The ability to intercept, index, and store ciphertexts at scale, using cheap and durable storage spread across cloud and terrestrial infrastructure.
- **Decryption capability:** The latent computational power that will come from quantum algorithms (such as Shor's and Grover's), specialized accelerators, and post-Moore architectures — capabilities that will eventually render classical encryption obsolete.
- **Temporal horizon:** Strategic patience — the willingness to defer exploitation for years or decades, bridging the gap between today's cryptographic strength and tomorrow's decryption capability.

These three resources evolve asynchronously. As a result, an adversary's effective power grows over time, even if the target organization's cryptography never changes.

↑ Back to Table of Contents

RISK MODELING

Q05. What mathematical condition defines a failure of data confidentiality?

In short

Confidentiality fails when the secrecy lifetime of the data is longer than the time until decryption becomes possible.

Confidentiality fails when the required secrecy lifetime of the data exceeds the adversary's decryption horizon. In simple terms:

The HNDL failure condition

If the data must remain secret for LONGER than it takes an adversary to acquire decryption capability, the data is already considered compromised — even if it cannot yet be read.

There are two variables in this relationship:

- **Required secrecy lifetime:** How long the data must remain confidential to retain its value or to satisfy regulatory mandates.
- **Adversary's decryption horizon:** The time remaining until the adversary can break the encryption — typically tied to the arrival of a cryptographically relevant quantum computer.

When the secrecy lifetime exceeds the decryption horizon, the data is mathematically considered compromised the moment it is harvested, because the encryption will fail before the data loses its sensitivity. More advanced models also express this as a probability function, estimating the likelihood that adversary capability will outpace the data's required lifespan.

↑ Back to Table of Contents

COMPARISON & CONTEXT

Q06. How do the system impacts and timelines of Y2K and Q-Day compare?

In short

Both require enterprise-wide coordination, but Q-Day is bigger, longer, and lacks a fixed deadline.

The transition to post-quantum cryptography is often compared to the Y2K challenge because both require enterprise-wide coordination. But they differ fundamentally in scope, duration, and the nature of the threat.

1. Timeline and predictability

Y2K: The deadline was fixed and well known — January 1, 2000. Most preparations could be completed in the 6 to 12 months beforehand, and the risk effectively ended within months of the deadline.

Q-Day: The date is unknown and probabilistic. Estimates generally target 2030, with fault-tolerant quantum computer projections ranging from 2028 to 2035. Large enterprises may take 12 to 15 years to fully migrate — a "risk deficit" where the time to fix the problem exceeds the time until it arrives. The threat is also ongoing: HNDL attacks are happening now, and risk will continue past Q-Day.

2. System impact and scope

Y2K was mainly a data formatting issue affecting specific applications. **Q-Day** is a fundamental failure of security infrastructure itself — it affects all interconnected systems, networks, applications, and data infrastructure. Every device, application, and module contains multiple layers of cryptography (protocols, certificates, firmware, signatures) that must be identified and upgraded. Simply swapping algorithms is not enough; enterprises must rethink how cryptography is integrated and managed, moving toward crypto-agility and automation.

3. Resources and remediation

Y2K required hundreds of billions of dollars and millions of labor hours, but the scope was quantifiable. **Q-Day** remediation effort is currently described as "completely unknown" — but the cost of inaction is projected to be exponentially higher than the cost of early implementation.

Y2K vs. Q-Day at a glance

Feature	Y2K	Q-Day
Target Date	Fixed: Jan 1, 2000	Unknown: estimated ~2030
Threat Duration	Finite (ended within months)	Ongoing (retroactive & future)
System Impact	Limited (data/app updates)	Broad (all infrastructure)
Lead Time	6–12 months	Years (potentially decades)
Remedy Effort	Billions of dollars	Completely unknown

Bottom line: Y2K is a useful analogy for the leadership and code-review effort required, but Q-Day is significantly higher risk because it turns historical data into a present liability through HNDL — and has no definitive finish line.

↑ Back to Table of Contents

REGULATORY DEADLINES

Q07. When does the EU roadmap require critical infrastructure transition to be complete?

> **In short**
>
> Critical infrastructure must complete PQC migration by the end of 2030; all other systems by 2035.

The European Union's Coordinated Implementation Roadmap requires critical infrastructure and other high-risk systems to complete their transition to post-quantum cryptography by **the end of 2030**.

This deadline applies to sectors specifically identified as vital — utilities, telecommunications, finance, and government — which must be secured against quantum threats as soon as possible, and no later than December 31, 2030. This 2030 target is distinct from the broader EU goal: all other feasible systems must complete the transition by **2035**.

↑ Back to Table of Contents

REGULATORY DEADLINES

Q08. What is the projected SSL/TLS certificate lifespan by 2029?

> **In short**
>
> TLS certificates will drop from 398 days to just 47 days by March 2029 — forcing monthly renewals.

By March 15, 2029, the maximum allowed lifespan for public SSL/TLS certificates is projected to drop to **47 days**. This is a major shift from today's 398-day standard and is part of a tiered timeline mandated by the CA/Browser Forum:

- **March 15, 2026:** Lifespan drops to 200 days.
- **March 15, 2027:** Lifespan drops to 100 days.
- **March 15, 2029:** Lifespan drops to 47 days.

The 47-day term forces a monthly renewal cadence — effectively creating 12 times more renewal work each year for IT teams. This is widely viewed as a critical "onramp" that pushes organizations to build the automation and crypto-agility they will need for the broader post-quantum transition.

↑ Back to Table of Contents

THE HNDL THREAT — FUNDAMENTALS

Q09. How does the Harvest Now, Decrypt Later threat model actually work?

> **In short**
>
> HNDL operates in three phases: harvest encrypted data today, store it, and decrypt it once quantum is viable.

HNDL (also called Store Now, Decrypt Later) is a temporal attack strategy that separates data theft from data compromise. It works by exploiting the gap between how long sensitive data must remain secret and how long it will take for a cryptographically relevant quantum computer (CRQC) to arrive. The attack unfolds in three phases:

Phase 1: Harvest (collection)

Attackers passively intercept and collect encrypted data. They can't read it yet — but they don't need to.

- **Methods:** Wiretaps on fiber-optic cables, satellite interception, compromised routers and servers. BGP hijacking — rerouting internet traffic through adversarial jurisdictions — is a known real-world indicator.
- **Targets:** Key exchange traffic (such as TLS handshakes) and encrypted payloads. Because interception is passive, it's hard to detect, letting adversaries build massive stockpiles silently.

Phase 2: Store (archival)

Harvested data is archived in long-term storage facilities — data centers or government repositories. This phase relies on cheap storage and strategic patience — the willingness to wait years or decades. The strategy targets data with a long tail of value: national security secrets, genomic data, intellectual property, and financial records. A one-time password isn't worth stealing this way, but data that must remain secret for 10 to 50 years is.

Phase 3: Decrypt (exploitation)

This phase begins once a fault-tolerant quantum computer capable of running Shor's algorithm exists. Shor's algorithm efficiently solves the math problems that underpin RSA and ECC, rendering those keys useless. At that point, the adversary processes the stockpiled ciphertext, recovers the session keys, and reads the historical data retrospectively.

Why this matters: irreversibility

Unlike a typical breach — where you can change a password or apply a patch — harvested data cannot be "un-harvested." The victim has permanently lost control of the asset; the only remaining variable is how long until the adversary has the computing power to unlock it.

> **Plain-language analogy**
>
> HNDL is like thieves stealing an uncrackable safe today, putting it in a warehouse, and waiting ten years for a laser tool that can slice through the door. The bank loses the contents the moment the safe is taken — not when the laser is finally used.

↑ Back to Table of Contents

SECTOR EXPOSURE

Q10. Why is the HNDL threat especially dangerous for healthcare data?

> **In short**
>
> Healthcare data is immutable, multi-generational, and tied to legacy devices — making it uniquely exposed.

HNDL is especially dangerous for healthcare because the lifespan of the data's sensitivity often exceeds the lifespan of the encryption protecting it. While a credit card number can be reissued or a password rotated, healthcare information has unique properties — immutability, multi-generational impact, and dependence on legacy devices — that make retrospective decryption catastrophic.

1. Extreme longevity of value

Personal health records, mental health history, and medical research routinely require confidentiality for the lifetime of the patient — typically 10 to 30 years, and often longer. Because healthcare data must remain protected for 25 to 75+ years, anything harvested today is effectively already compromised. Quantum decryption is projected to be feasible well within that window.

2. The immutability of genomic data

Genomic information is the most critical subset of healthcare data because it is biologically immutable.

- **Permanent vulnerability:** A DNA sequence doesn't change. Once decrypted, the privacy loss is irreversible — there is no "credential reset."
- **Multi-generational impact:** Compromise extends to biological relatives and descendants, exposing them for 50 to 100 years.
- **High-consequence threats:** Decrypted genomic data could enable genetic discrimination by insurers or employers. In the worst case, the convergence of quantum computing and synthetic biology could theoretically allow targeted bio-threats based on individual genetic markers.

3. Infrastructure "security debt" in IoMT

Hospitals run thousands of connected devices — pacemakers, insulin pumps, imaging machines — that operate for 15+ years and rely on embedded, legacy encryption that often can't be patched. This creates ongoing collection vectors: devices deployed today will still be transmitting harvestable data when quantum threats mature.

4. Long-tail liability and erosion of trust

Healthcare breaches are already the most expensive type of data compromise — averaging $9.77 million per incident. HNDL adds long-tail liabilities, where organizations may face lawsuits and regulatory penalties for breaches that effectively occur years after the initial theft. Even more damaging: if patients fear their health secrets will eventually be exposed, they may withhold sensitive information from clinicians, compromising both care quality and public health outcomes.

↑ Back to Table of Contents

SECTOR EXPOSURE

Q11. Which specific sectors are classified as facing critical HNDL exposure?

> **In short**
>
> Sectors with 30+ year or indefinite data lifespans: state intelligence, public blockchains, genomic data, and child welfare records.

Critical Exposure applies to sectors that hold data with confidentiality requirements exceeding 30 years, or that must be retained indefinitely. In these cases, the required lifespan of the data far outlasts the projected arrival of fault-tolerant quantum computers — making retrospective compromise effectively certain if migration is delayed.

Sectors classified as Critical Exposure

- **State intelligence and national security:** State secrets, intelligence identities, and classified archives often carry confidentiality mandates of 30 to 100+ years. Diplomatic cables and sovereign treaties are statistically guaranteed to be exposed if harvested today and stored until a quantum computer arrives.
- **Public blockchains:** Distributed ledgers have indefinite data lifespans because the ledger is immutable. Bitcoin's entire transaction history is preserved forever. An adversary who harvests the ledger today can deanonymize all historical transactions once a quantum computer is available.

- **Genomic data and child welfare records:** General medical records sit in the High Exposure tier (10–30 years), but genomic data is Critical because it is biologically immutable and impacts biological relatives, creating a 50- to 100-year sensitivity window. Child welfare records carry similar protection mandates spanning nearly a century.

How Critical differs from High Exposure

Sectors such as satellite communications, legal and government records, and general financial contracts typically fall into the High Exposure tier — data longevity of 10 to 30 years. They still face imminent risk and require urgent migration, but they don't carry the multi-generational or indefinite vulnerability that defines the Critical tier.

↑ Back to Table of Contents

ADVERSARY MODEL

Q12. What are the four primary resources an HNDL adversary possesses?

> **In short**
>
> Technically there are three — but the literature includes other "groups of four" that are easy to confuse with them.

The formal HNDL adversary model identifies **three** resources, not four. The peer-reviewed source — the Telecom journal article "Harvest-Now, Decrypt-Later: A Temporal Cybersecurity Risk in the Quantum Transition" — defines the adversary by:

- **Collection capability:** Intercept, index, and store ciphertexts at scale using durable, inexpensive storage.
- **Decryption capability:** Latent computational power from quantum algorithms (Shor's, Grover's), specialized accelerators, and post-Moore architectures.
- **Temporal horizon:** Strategic patience — the willingness to defer exploitation for years or decades.

Where the "four" confusion comes from

Other frameworks in the same literature do list four items, which is easy to mix up with the adversary model:

- **Four mitigation approaches:** Post-quantum cryptography (PQC), hybrid key exchange, forward-secure lifecycles, and governance/policy.
- **Four primary precepts:** The U.S. federal PQC migration strategy is built on inventory, early start, prioritization, and identification of unsupported systems.
- **Four critical sectors:** Financial institutions, government agencies, defense contractors, and healthcare providers.

↑ Back to Table of Contents

RISK MODELING

Q13. What mathematical condition determines when HNDL confidentiality failure occurs?

> **In short**
>
> When required secrecy lifetime exceeds the adversary's decryption horizon, confidentiality has already failed.

Confidentiality fails when the required secrecy lifetime of the data exceeds the adversary's decryption horizon. In plainer terms, the data fails the moment it is harvested if the encryption will be broken before the data loses its sensitivity.

The two variables are:

- **Required secrecy lifetime:** How long the data must remain confidential to keep its value or meet regulatory requirements.
- **Adversary's decryption horizon:** How long until the adversary acquires the capability — typically a cryptographically relevant quantum computer — to break the encryption.

What this means in practice

When the secrecy lifetime exceeds the decryption horizon, the encryption has effectively "expired" before the data loses its sensitivity. Data harvested today is already compromised, because the adversary will be able to decrypt it while it is still secret.

Residual exposure: the migration window

Risk models also calculate a residual exposure window — the difference between the data's secrecy lifetime and the adversary's horizon at the point of migration. If that window is positive, the sector faces inevitable retrospective compromise for that duration.

> **Example: the healthcare sector**
>
> If healthcare data has a required lifetime of 25 years, and the adversary's decryption horizon is projected at 19 years, then the residual exposure window is 6 years. Healthcare data harvested today will be exposed for 6 years before encryption catches up.

Because the exact arrival of quantum capability is uncertain, advanced models also express the condition as a probability function — estimating the likelihood that the adversary's capability will mature within the data's required timeframe.

↑ *Back to Table of Contents*

SECTOR EXPOSURE

Q14. Which sectors are classified as having critical HNDL exposure? (detailed view)

In short

State intelligence, public blockchains, genomic data, and child welfare records — anywhere data needs 30+ years of secrecy.

Sectors are classified as Critical Exposure when their data's required secrecy lifetime exceeds 30 years, or is indefinite. In these cases, the data's longevity far outlasts the time remaining until a cryptographically relevant quantum computer (CRQC) is expected to arrive — which makes retrospective compromise a statistical certainty if the data is harvested today.

Critical Exposure sectors

- **State intelligence and national security:** State secrets, diplomatic cables, and classified archives, with confidentiality mandates spanning 30 to 100+ years.
- **Public blockchains:** Indefinite data lifespan due to ledger immutability. An adversary harvesting Bitcoin's ledger today could later deanonymize the entire transaction history.
- **Genomic data:** Biologically immutable, affecting biological relatives — sensitivity window of 50 to 100 years.
- **Child welfare records:** Protection periods that span nearly a century, placing them in the Critical tier alongside genomic data.

High Exposure (a different, but still urgent, tier)

Satellite communications, legal and government records, scientific archives, and general financial contracts typically sit in the High Exposure tier — data longevity of 10 to 30 years. These sectors face imminent risk and require urgent migration, but they don't carry the multi-generational or indefinite vulnerability profile of the Critical tier.

↑ Back to Table of Contents

SECTOR EXPOSURE

Q15. Why is genomic data uniquely vulnerable to retrospective decryption attacks?

In short

DNA can't be changed, affects relatives, and remains exploitable for a century — making it uniquely exposed to HNDL.

Genomic data is uniquely vulnerable to HNDL because of three properties: biological immutability, multi-generational impact, and extreme longevity of value. Unlike other sensitive data types, the eventual compromise of genomic data is permanently irreversible and potentially catastrophic.

1. Biological immutability

The most critical vulnerability is that genomic data cannot be changed. A password can be rotated, a credit card can be reissued, and even a Social Security number can theoretically be reassigned — but a DNA sequence is constant from birth to death.

- **Permanent privacy loss:** Once decrypted, the loss is absolute. There is no remediation strategy.
- **Non-decaying value:** Genomic data does not lose relevance over time. It remains accurate and exploitable for the entirety of the subject's life.

2. Multi-generational and familial risk

Genomic data carries a hereditary risk factor unique among data types. Compromising one person's genetic code mathematically exposes their biological relatives — parents, siblings, and children — who may have never consented to data collection.

- **Descendant exposure:** The sensitivity window is 50 to 100 years because it affects future descendants.
- **National security implications:** Decrypted genomic data can identify biological relatives of government officials, military personnel, or intelligence agents — creating coercion and forensic risks that span generations.

3. Specific high-consequence threats

- **Targeted bio-threats:** The convergence of quantum computing and synthetic biology could theoretically enable bioweapons or treatments targeted at specific individuals or populations based on genetic markers.
- **Genetic discrimination:** Decrypted data could be used by insurers or employers to deny coverage or employment based on genetic predispositions — decades after the original theft.

Why it lands in the Critical tier

Because of these factors, genomic data is classified as Critical Exposure. Its required secrecy lifetime — often a century — far exceeds the adversary's decryption horizon, which is projected around 2030 to 2035. That gap effectively guarantees that genomic data encrypted with today's standards and harvested now will be compromised while it is still highly sensitive.

↑ Back to Table of Contents

CRYPTOGRAPHY & STANDARDS

Q16. How does Shor's algorithm threaten classical asymmetric encryption standards?

> **In short**
>
> Shor's algorithm lets a quantum computer break RSA and ECC exponentially faster than any classical method.

Shor's algorithm threatens classical asymmetric encryption by giving quantum computers a way to solve the math problems behind public-key cryptography — integer factorization and discrete logarithms — exponentially faster than any known classical algorithm.

The underlying mathematical vulnerability

Current asymmetric encryption standards rely on problems that are too hard for classical computers to solve in any reasonable time:

- **RSA:** Relies on the difficulty of factoring large composite integers into their prime factors.
- **ECC (Elliptic Curve Cryptography):** Relies on the difficulty of solving the discrete logarithm problem.

Classical computers would need millions or trillions of years to solve these problems for sufficiently large key sizes. Shor's algorithm uses quantum mechanical properties — superposition and entanglement — to convert these intractable problems into tasks that can be solved in polynomial time.

Exponential speedup and key derivation

The threat is existential because Shor's algorithm allows an attacker to derive a private key directly from a public key.

- **RSA impact:** A classical computer might need 2^50 steps to break a scheme; a quantum computer running Shor's might need only 50. While a classical supercomputer would take trillions of years to break RSA-2048, a sufficiently powerful quantum computer could theoretically do it in hours — or even seconds.
- **ECC impact:** Shor's algorithm also computes discrete logarithms, leaving ECC vulnerable to the same kind of rapid decryption.

Why making keys bigger doesn't help

Symmetric cryptography (like AES) is threatened by a different quantum algorithm — Grover's — and can be defended by doubling the key size (AES-128 → AES-256). Asymmetric encryption cannot be saved this way:

- Because Shor's offers exponential speedup, doubling an RSA modulus (e.g., 2048 → 4096 bits) only makes the problem roughly 8 times harder — a polynomial factor, not exponential.
- As a result, RSA and ECC are considered mathematically fragile. They must be completely replaced by post-quantum cryptography (PQC) algorithms — not simply strengthened.

↑ Back to Table of Contents

CRYPTOGRAPHY & STANDARDS

Q17. What are the primary performance trade-offs when implementing PQC algorithms?

In short

Larger keys, more computation, more memory, and hybrid overhead — PQC isn't a drop-in replacement.

The transition to post-quantum cryptography introduces significant performance trade-offs, primarily around bandwidth, computation, and memory. Unlike past upgrades (DES to AES), PQC requires fundamental architectural changes rather than simple key lengthening.

1. The "bandwidth tax": key and signature sizes

The most immediate trade-off is a dramatic increase in the size of cryptographic artifacts, which impacts network latency and protocol stability.

- **Massive size increases:** An ML-KEM-768 (Kyber) public key is 1,184 bytes versus 32 bytes for classical ECDH — a 30–40x increase. ML-DSA-65 (Dilithium) signatures are roughly 3,309 bytes versus 64 bytes for ECDSA — a 50x increase.
- **Protocol impact:** TLS handshakes quadruple in size — from about 4 KB to over 15 KB when using hybrid schemes.
- **Fragmentation and packet loss:** Larger packets cause fragmentation, which can degrade performance on high-latency links (satellite, cellular). Legacy middleboxes — firewalls, load balancers — may treat the larger packets as anomalies and drop connections, requiring network re-architecture.

2. Computational asymmetry and latency

Performance varies significantly depending on the operation and the hardware environment.

- **Encryption vs. verification:** Lattice-based schemes like ML-KEM are efficient on modern hardware and often outperform classical algorithms in encryption speed. But ML-DSA signature verification can be heavier — a real concern in Zero Trust architectures where every API request needs verification.
- **Hardware dependency:** Modern CPUs with vector instructions (AVX2) handle PQC well; embedded systems do not. ML-KEM on embedded devices without hardware acceleration can run 10 to 100 times slower.
- **Trusted execution environments (TEE):** Inside secure enclaves like Intel SGX, ML-KEM performance can drop to roughly 30% of standard speed due to memory access constraints.

3. Memory constraints in IoT and embedded systems

Memory requirements may render many legacy devices obsolete.

- **RAM exhaustion:** Low-power microcontrollers typically have 32 to 64 KB of RAM. ML-KEM-768 alone needs roughly 15 KB of working memory — leaving little room for application logic and potentially forcing a hardware replacement rather than a software update.
- **Smart card limitations:** EMV chip cards would need 4 to 6 times more RAM to support PQC algorithms — an economic and engineering challenge for mass deployment.

4. The hybrid overhead

To manage risk during the transition, most organizations are adopting hybrid cryptography — running both a classical algorithm (like ECC) and a post-quantum algorithm (like ML-KEM) at the same time.

- **Double the work:** This provides defense in depth but effectively doubles the computational overhead for key generation, signing, and verification.
- **Operational cost:** Maintaining hybrid environments creates an estimated 20–40% overhead for cryptographic operations staff due to the complexity of patching and managing two distinct stacks.

5. Algorithm-specific trade-offs

- **Lattice-based (ML-KEM, ML-DSA):** Best balance of speed and size, but relatively new and complex to implement securely against side-channel attacks.

- **Hash-based (SLH-DSA, SPHINCS+):** Extremely high security confidence (conservative math), but with massive signature sizes (8 KB to 40 KB) — impractical for bandwidth-constrained applications like mobile apps or real-time communications.

↑ Back to Table of Contents

Q-DAY & TIMELINE

Q18. What is Q-Day, and who is predicting it will happen in 2029?

In short

Q-Day is the day quantum computers break RSA and ECC. Google, IBM, and Microsoft project it for 2029–2030.

Q-Day is the day when a quantum computer becomes powerful enough to break the public-key cryptography that secures most of the world's digital infrastructure — specifically RSA and Elliptic Curve Cryptography (ECC). It is not a hypothetical event in some distant future. It is a date that the three largest quantum computing investors — **Google, IBM, and Microsoft** — now project will arrive in 2029.

Why 2029?

The 2029 projection comes from published quantum hardware roadmaps. IBM, Google, and Microsoft have each publicly committed to delivering fault-tolerant quantum computing in the late 2020s. Their hardware milestones — qubit counts, error correction breakthroughs, and logical qubit demonstrations — are now converging on the same window. When all three industry leaders independently arrive at the same timeline, it stops being speculation.

What "breaks" on Q-Day

The asymmetric encryption protecting your TLS sessions, your VPNs, your digital signatures, your code-signing certificates, your banking APIs, and almost every secure connection on the internet today. RSA-2048 — currently estimated at trillions of years to break — could be cracked in hours or seconds.

Why this matters today, not in 2029

Q-Day is a future event, but its consequences are happening now. Adversaries are already harvesting encrypted data through Harvest Now, Decrypt Later (HNDL) attacks — storing today's traffic for tomorrow's decryption. By the time Q-Day arrives, the damage to historical data is already done. The question for every organization is no longer "if" Q-Day will happen — it's whether your data will already have been harvested before you migrate.

↑ Back to Table of Contents

Q-DAY & TIMELINE

Q19. If Q-Day is 5+ years away, why do we have to act now?

In short

Because adversaries are harvesting your encrypted data today and migration takes years — not days.

There are three reasons to act now, even though Q-Day itself is several years out. Each one alone justifies immediate action. Together they make delay indefensible.

1. Your data is being stolen right now

Adversaries don't need a quantum computer to start the attack. They only need one for the final step. Encrypted traffic is being intercepted and stockpiled today — through fiber-optic taps, satellite interception, and BGP rerouting through adversarial jurisdictions. By the time you migrate to post-quantum cryptography, the historical data that was already harvested will eventually be decrypted.

2. Migration takes longer than the time you have

Large enterprises typically need 12 to 15 years to fully migrate cryptography. If Q-Day arrives in 2029 — five years from now — most organizations face a guaranteed exposure window of 7 to 10 years. This is what risk modelers call the "risk deficit": the time required to fix the problem is longer than the time before the problem becomes critical.

3. The legal and regulatory windows have already closed

U.S. Department of Justice 2025 guidelines now treat inaction on emerging cryptographic threats as a possible "willful violation" rather than negligence. The EU requires critical infrastructure to complete post-quantum migration by **December 31, 2030**, with all other systems by 2035. SSL/TLS certificate lifespans drop to 47 days by March 2029 — forcing the automation that PQC migration also requires. Regulators have already accepted that PQC migration is non-optional.

The bottom line

Waiting for Q-Day to arrive before migrating is mathematically equivalent to accepting that your historical data will be exposed. The only question is whether you migrate before or after that exposure happens.

↑ Back to Table of Contents

EXECUTIVE LIABILITY

Q20. Can executives now be held personally liable for not addressing the quantum threat?

In short

Yes. Under 2025 DOJ guidelines, inaction can escalate from negligence to "willful violation" — with criminal exposure.

Personal executive liability for data security failures has become a real and present risk. Under **2025 Department of Justice guidelines**, inaction on known emerging threats — including the quantum threat — can shift the legal status of a breach from negligence to "willful violation." The distinction matters enormously: willful violation opens the door to personal criminal charges, personal civil penalties, and uninsured liability.

What changed in 2025

DOJ guidance now treats failure to address publicly known emerging threats — when the threat is documented, the standards exist, and viable mitigations are available — as a potential willful violation by the responsible executives. The quantum threat checks every box: it is publicly documented (NIST has finalized PQC standards), the mitigations exist (ML-KEM, ML-DSA), and the timeline is established (Q-Day projected for 2029).

What executives are personally exposed to

- **Criminal prosecution:** Federal and state statutes attach significant prison sentences for willful violations involving healthcare data, AI use, and consumer privacy.
- **Personal civil penalties:** Fines that can be assessed against individual officers and directors — not just the company.
- **Uninsured liability:** Most Directors & Officers (D&O) policies exclude penalties and willful misconduct. Personal assets become exposed.
- **Clawbacks:** Compensation may be recoverable from executives whose decisions contributed to the breach.
- **State-level exposure:** A growing number of states impose criminal penalties for failures in healthcare, AI, and privacy compliance — and jurisdiction follows the patient, not the company headquarters.

Why "we were evaluating it" is no longer a defense

With NIST-certified post-quantum cryptography solutions now deployable in 90 days, prolonged "evaluation" looks less like prudence and more like willful inaction. Regulators and prosecutors are aware that viable, certified solutions exist.

↑ *Back to Table of Contents*

EXECUTIVE LIABILITY

Q21. How does deploying PQC+ protect executives personally?

In short

PQC+ creates a documented "good faith" defense that shifts breach classification away from willful violation.

Deploying NIST-certified post-quantum cryptography is one of the strongest "good faith" defenses available to executives today. It directly addresses every category of personal liability that has emerged under 2025 DOJ guidelines and state-level enforcement. Here is what it does, specifically:

Mitigates "willful violation" designations

Deploying NIST-certified PQC is a concrete, documentable action against a known emerging threat. It is exactly the kind of good-faith compliance measure regulators look for when deciding whether to classify a breach as negligent (lower exposure) or willful (criminal exposure).

Defends against criminal prosecution

Federal and state statutes that attach prison sentences to willful violations require proof of willfulness. A documented PQC+ deployment — with NIST FIPS 203/204/205 certifications, audit trails, and a compliance dossier — establishes the factual record that defeats willfulness claims.

Reduces personal civil penalty exposure

Many of the personal civil penalties now appearing in state law require evidence of inaction. A deployed, certified solution creates a paper trail that documents the opposite: timely, diligent action.

Covers uninsured costs

D&O policies typically exclude penalties and willful misconduct. PQC+ helps keep breaches in the "negligence" category — where insurance still applies — and provides documentation against clawbacks.

Meets multi-state regulatory thresholds

PQC+ is designed to meet the strictest jurisdictional requirements across the states that now impose criminal penalties for healthcare, AI, and privacy violations. Because the system enforces consent and access rules automatically at the data layer, compliance evidence is generated continuously — not reconstructed after a breach.

Closes the "deployment gap"

With a 90-day deployment path available, the historical defense of "we needed time to evaluate" no longer holds. PQC+ makes the diligent path the practical path.

> **What protection looks like in practice**
>
> A documented record showing NIST-certified PQC deployment, dated audit trails of executive decision-making, compliance dossiers ready for regulators, and four-layer cryptographic protection that makes harvested data mathematically useless. That record is what shifts a breach from willful violation to a defensible incident.

↑ Back to Table of Contents

THE PQC+ SOLUTION

Q22. What is PQC+ and how is it different from traditional encryption?

In short

PQC+ embeds access control directly into the cryptographic key — making stolen data mathematically useless.

PQC+ (Post-Quantum Cryptography Plus) is a data-centric security platform that combines NIST-certified post-quantum encryption with access control rules embedded directly inside the cryptographic key structure. It is built on three pillars: post-quantum algorithms (FIPS 203, 204, 205), hardware-bound keys, and dynamic attribute-based access control (ABAC). Together they shift security from "protecting the network" to "protecting the data itself."

Where traditional encryption falls short

Traditional security treats encryption and access control as two separate systems. Data is encrypted, and then a separate layer of policy decides who can decrypt it. Once an attacker is inside the network — or once they steal a database — the access control layer is bypassed, and a single master key opens everything.

What PQC+ does differently

PQC+ fuses encryption and access control together. The access policy is not a separate check — it is mathematically embedded in the key itself. A stolen database is not a collection of files protected by a master key. It is a collection of millions of individually quantum-locked blobs, each tied to a specific identity, specific hardware, and specific authorization context.

- **Indecipherable blobs:** Stolen files are random-looking values without the specific decapsulation keys.
- **Hardware lock:** Even with a stolen password, data cannot be read without the specific hardware-defined keys linked to an authorized device.
- **No resale value:** Because data is tied to specific identities and environments, it cannot be sold on the dark web — buyers also lack the keys and MFA access.

Why this is called the "Holy Grail"

The fundamental flaw in modern cybersecurity is that once an attacker is inside, they have free rein. PQC+ changes the game by tying security to the data itself, not the perimeter. The data becomes a portable vault — secure regardless of where it travels or who steals it.

↑ *Back to Table of Contents*

THE PQC+ SOLUTION

Q23. How does PQC+ neutralize the Harvest Now, Decrypt Later threat?

In short

> By using NIST FIPS 203/204 algorithms designed to resist quantum attacks — and binding decryption to hardware and identity.

PQC+ neutralizes HNDL on two fronts. First, it uses cryptographic algorithms that quantum computers cannot break. Second, even if those algorithms were broken, the data still could not be read without the specific hardware and identity context bound to each record.

Layer one: quantum-resistant algorithms

PQC+ is built on the NIST post-quantum cryptography standards finalized in 2024:

- **FIPS 203 (ML-KEM):** Module-Lattice-Based Key-Encapsulation Mechanism — used for key exchange. Replaces RSA and ECDH.
- **FIPS 204 (ML-DSA):** Module-Lattice-Based Digital Signature Algorithm — replaces RSA and ECDSA signatures.
- **FIPS 205 (SLH-DSA):** Stateless Hash-Based Digital Signature — provides additional signature protection with extremely high security confidence.

These algorithms are based on mathematical problems (lattice problems and hash-based constructions) that are hard for both classical and quantum computers — they are not vulnerable to Shor's algorithm the way RSA and ECC are.

Layer two: hardware-bound, identity-bound decryption

Even if quantum computers somehow advanced beyond what's anticipated, PQC+ data still cannot be read without:

- The **specific hardware-defined keys** of an authorized device.
- The **authorized identity** with active MFA.
- A **dynamic ABAC context match** — including location, time, IP reputation, and hardware integrity.

Tiered key strength for long-lived data

PQC+ supports tiered key sizes matched to how long data needs to remain secret:

Tier	Key Size	Intended Protection
Commercial	1024 bits	Standard business operations
Military	2048 bits	Sensitive defense and mission-critical data
Intelligence	5120 bits	Highly classified intelligence
Max Strength	**10240 bits**	Designed to protect sensitive data for decades

For industries where data must remain secret for 50+ years — oil and gas geological data, national intelligence, healthcare records — the 10240-bit tier provides entropy designed to outlive the advancement of decryption technology itself.

↑ Back to Table of Contents

THE PQC+ SOLUTION

Q24. Can PQC+ really be deployed in 90 days?

In short

Yes — because PQC+ is software-based and sits alongside existing systems, with no hardware changes or "rip and replace."

Yes. PQC+ is designed for 90-day deployment because it does not require new hardware, does not replace existing systems, and does not require organizations to migrate their data. It deploys as a software layer that sits alongside the systems already in place.

Why traditional cryptography migrations take years

Most cryptographic migrations require swapping out hardware, reissuing certificates across thousands of endpoints, coordinating across every application that uses encryption, and rewriting integration code. That work routinely takes large enterprises 12 to 15 years.

Why PQC+ is different

- **Software-only:** No new hardware to procure, install, or commission. No firmware updates to legacy devices.
- **Cloud-based platform:** Scaling and updates are handled centrally. Standard HTTPS outbound traffic is typically all the network change required.
- **Coexistence, not replacement:** PQC+ does not replace your EHR, your lab system, your imaging PACS, your billing platform, or your existing interface engine. It sits in the middle as a connecting layer.
- **Format-agnostic:** Works with HL7 v2, C-CDA, FHIR, X12, and DICOM out of the box. Existing interfaces continue to operate.

What a 90-day deployment looks like

- **Weeks 1–2 — Contracting and planning:** Scope definition, Business Associate Agreement execution, project kickoff.
- **Weeks 3–6 — Foundation (Phase 1):** Core platform live, EHR connected, basic consent operational, QHIN connectivity established.
- **Weeks 7–10 — Clinical enhancement (Phase 2):** Clinical modules deployed, SMART on FHIR apps embedded in EHR workflows.
- **Weeks 11–13 — Revenue and compliance (Phase 3):** Revenue cycle modules active, partnership integrations generating measurable value, compliance dossier complete.

Why this matters for executive liability

A 90-day deployment path eliminates the "we needed time to evaluate" defense. It also creates the dated record of diligent action that protects executives personally if a breach occurs later.

↑ Back to Table of Contents

PQC+ ARCHITECTURE

Q25. What is the "four-layer cryptographic envelope" and what does it actually protect?

In short

Four reinforcing layers — PQC encryption, embedded ACLs, hardware-bound keys, and dynamic context — that make stolen data useless.

The four-layer cryptographic envelope is the architectural backbone of PQC+. Each layer addresses a different attack vector. Together they ensure that a stolen database cannot be decrypted, even with a future quantum computer, even with stolen credentials, even on a compromised device.

Layer 1: Post-quantum encryption

Every data record is encrypted using NIST-certified post-quantum algorithms — FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), and FIPS 205 (SLH-DSA). These algorithms resist both classical and quantum attacks, neutralizing the HNDL threat at the cryptographic level.

Layer 2: Embedded access control list (ACL)

Access rules are not enforced by a separate system — they are mathematically baked into the key structure itself. The data and its authorization rules are inseparable. There is no way to bypass the access check by going around it, because there is nothing to go around.

Layer 3: Hardware-bound keys

Decryption requires the specific hardware-defined key of an authorized device. A stolen password is not enough. A stolen device is not enough without the authorized identity. This breaks the economic model of credential theft.

Layer 4: Dynamic ABAC context

Before decryption is permitted, the Dynamic Attribute-Based Access Control engine evaluates three context pillars in real time:

- **Subject attributes:** PQC-verified identity, clearance level, current MFA status.
- **Resource attributes:** Sensitivity tier of the data (Commercial, Military, Intelligence, Max Strength).
- **Environment attributes:** GPS location, time of day, IP reputation, hardware integrity. A drone that flies into a restricted zone, or a device that has been tampered with, simply stops getting valid ACLs generated.

What the layers achieve together

A stolen database becomes millions of individual, quantum-locked files. Without the unique combination of identity, MFA, hardware ID, and active context match, the data is mathematically useless. There is no resale value on the dark web. There is no future decryption by quantum computers. There is no way to monetize what was stolen.

↑ Back to Table of Contents

PQC+ ARCHITECTURE

Q26. How does PQC+ handle consent — and why is consent treated as a cryptographic attribute?

In short

Consent is encoded as a real-time attribute inside the encrypted packet. Revoke consent and access dies everywhere, instantly.

In most systems, consent is a yes/no checkbox stored in a database. Someone checks it before allowing access. The problem: that check happens far from the data, after the data has already been retrieved and decrypted. If the check is bypassed, the data is exposed.

PQC+ treats consent as a **dynamic cryptographic attribute** — part of the same ABAC evaluation that controls all other access. Consent is not a label attached to data. It is a condition that must be satisfied before the data can be decrypted at all.

How it works in practice

- **Consent becomes an attribute:** Each consent decision is an attribute the ABAC engine evaluates in real time alongside identity, role, and hardware integrity.
- **Granular masking at the data layer:** If a patient consents to share heart rate data but not GPS location, the GPS fields are masked at the cryptographic layer — before the data is sent. Not after retrieval. Before.
- **Privacy-by-design enforcement:** Because the consent attribute is part of the encrypted packet, no remote system needs to call home to check consent. The data packet itself knows whether it is allowed to open for that requester.

The revocation kill-switch

When a patient or data subject withdraws consent, the Privacy Management system updates the consent attribute immediately. The next time the data is accessed — anywhere in the world, on any device — the ABAC engine sees "Consent = False" and refuses to generate a valid ACL. Access dies everywhere, instantly. This is the closest equivalent to a true kill-switch for data that has already been distributed.

Why this collapses the dark web data market

The dark web data market depends on bulk data utility — buy a database, sell or exploit the contents. PQC+ makes that model worthless. A stolen PQC+ database is not data — it is a heap of locked envelopes that will never open without the correct identity, hardware, MFA, and active consent attribute. There is no payload to monetize.

> **Why this fulfills GDPR and HIPAA automatically**
>
> These regulations require that data processing cannot occur without valid consent. With PQC+, that requirement is mathematical — not procedural. The data physically cannot be processed without the valid consent attribute being met.

↑ Back to Table of Contents

PQC+ IN HEALTHCARE

Q27. How does PQC+ integrate with existing healthcare systems like Epic, Oracle Health, and MEDITECH?

> **In short**
>
> PQC+ sits as a connecting layer alongside existing EHRs — no rip-and-replace, no data migration.

PQC+ was designed around the reality that hospitals have invested millions in systems like Epic, Oracle Health (formerly Cerner), MEDITECH, and Allscripts. Replacing those systems is not on the table. Instead, PQC+ sits as a secure, compliant fabric in the middle — connecting everything, securing everything, and adding capabilities the existing systems lack.

The architecture

Your existing systems remain at the bottom of the stack: EHR, lab system, PACS, billing, practice management. PQC+ sits in the middle, providing SMARTInteroperability, SMARTEntityResolution, SMARTCompliance, SMARTDataLake, and other modules. At the top are external connections: QHINs, insurance payers, labs, imaging centers, and HIEs.

Epic environments

- **FHIR R4 APIs:** Bidirectional flow of demographics, encounters, diagnoses, medications, labs, and clinical notes. PQC+ modules surface within Epic via App Orchard and SMART on FHIR.
- **Care Everywhere integration:** Extends Epic's external data exchange to non-Epic sources, adding consent management that Epic does not provide natively.
- **ADT/HL7 feeds:** Real-time operational data via standard HL7 v2 ADT messages. The existing interface engine keeps working.

Oracle Health (Cerner) environments

- **FHIR R4 APIs:** Clinical data exchange through Oracle's FHIR R4 implementation.

- **Millennium data integration:** Deeper access to clinical documentation, orders, results, scheduling, and revenue cycle data.
- **CommonWell connectivity:** Complements Oracle's CommonWell membership — adds data from networks CommonWell does not reach directly.

MEDITECH, Allscripts, and other systems

PQC+'s interoperability engine is format-agnostic. For systems with less mature FHIR implementations, it ingests HL7 v2 (ADT, ORM, ORU, MDM), parses C-CDA documents, handles X12 transactions (eligibility, claims, remittance), and supports Direct secure messaging. Existing interfaces don't need to change.

> **Your data doesn't move**
>
> You are not migrating 15 years of patient data out of your EHR. PQC+ creates a unified index that knows where data lives and retrieves it on demand. The data stays where it is. PQC+ provides the intelligence layer that makes it accessible and compliant.

↑ Back to Table of Contents

PQC+ IN HEALTHCARE

Q28. How does PQC+ govern data flowing into AI systems?

> **In short**
>
> Every AI data request passes through an architecturally enforced gateway using Model Context Protocol (MCP).

AI is already in hospitals — clinical decision support, ambient documentation, predictive analytics, imaging interpretation. More is coming. The challenge is that most existing consent frameworks were designed before AI existed: patients consented to "treatment, payment, and healthcare operations" without contemplating that their data might train a language model or feed an automated diagnostic algorithm. Regulators are catching up fast.

How the AI Compliance Gateway works

PQC+'s AI Compliance Gateway uses Model Context Protocol (MCP) — a standardized way for AI models to request and receive data from external systems. Every AI data request must pass through this checkpoint. It is not optional. It is architecturally enforced.

- **Step 1:** The AI application sends a data request via MCP.
- **Step 2:** The gateway checks the patient's consent settings for that specific AI category.
- **Step 3:** If consent is granted, data is returned with any configured masking applied. If not, the request is denied and the attempt is logged.
- **Step 4:** A full audit trail is recorded — what data was shared, to which AI vendor, under which consent basis.

Separate consent for each AI category

PQC+ supports granular consent that recognizes patients have different feelings about different AI uses:

- **Clinical decision support:** AI tools that help clinicians diagnose or treat. Most patients accept this — it directly benefits their care.
- **Automated decision-making:** AI that makes decisions without human review. Higher stakes — patients may want to opt out.
- **LLM training data:** Using patient data to train language models. Data leaves the institution; patients often object.
- **Third-party AI applications:** External vendors' AI tools. Patients may not trust unknown third parties.
- **Research AI:** AI used in clinical research. Different regulatory framework (IRB oversight).

Automatic data masking before AI ever sees the data

When AI is authorized, the gateway still applies protections: PII masking strips names and SSNs; date shifting prevents re-identification while preserving clinical meaning; geographic generalization replaces specific addresses with regions; sensitive category redaction removes substance use, reproductive health, or behavioral health data based on consent settings.

> **Ready for the regulatory question**
>
> When regulators ask "How is patient data flowing into AI systems at your hospital?" — you have a complete, auditable answer. Every request, every approval, every denial, every masking decision is logged with timestamps.

↑ Back to Table of Contents

PQC+ IN HEALTHCARE

Q29. What measurable financial return does PQC+ generate for a hospital?

> **In short**
>
> Medication pricing transparency, value-based care optimization, and consolidated point solutions generate measurable ROI from day one.

PQC+ generates measurable financial return in three categories: new revenue, cost reduction, and risk reduction. Each one is concrete and quantifiable.

1. Medication pricing transparency

Through a strategic pharmacy network partnership covering 22,000+ pharmacies, PQC+ delivers real-time medication pricing at the point of care. Pharmacy Benefit Manager (PBM) markups — typically 8–15% on transactions — drop to just 2%. The hospital participates in transaction-based revenue sharing on facilitated

pharmacy transactions. SMARTRxPricing also includes automatic HCPCS and CPT code mapping for the clinician's order.

2. Value-based care optimization

SMARTVBR calculates Risk Adjustment Factor (RAF) scores accurately, capturing reimbursement that is routinely lost to undercoding. SMARTRCM uses AI to analyze eligibility across state, federal, and private programs, and produces "what-if" projections for different payment models.

3. Consolidation of point solutions

Hospitals typically pay for many separate tools that PQC+ replaces with consolidated modules:

- **Consent management software** → SMARTCompliance
- **Data masking and anonymization tools** → AI Compliance Gateway
- **FHIR server and interoperability engine** → SMARTInteroperability
- **Identity resolution / MPI software** → SMARTEntityResolution
- **Audit logging and compliance documentation** → SMARTCompliance audit trails
- **AI governance tools** → AI Compliance Gateway + MCP controls

4. Liability and breach cost reduction

The average HIPAA breach now costs $9.77 million per incident. Time-based audit trails document every data access with verifiable records — reducing exposure if a breach does occur. More importantly, PQC+'s embedded access control makes stolen data mathematically useless, which changes the breach calculus entirely.

> **The combined effect**
>
> QHIN connectivity gives you comprehensive patient data. AI compliance controls let you use that data with AI tools while maintaining trust and compliance. Financial partnerships create revenue that funds the platform investment. Patient consent sits underneath all of it.

↑ Back to Table of Contents

COMPARISON & STRATEGY

Q30. Why is PQC+ called the "Holy Grail" of cybersecurity?

> **In short**
>
> Because it solves the fundamental flaw of perimeter security — once inside, attackers have free rein. PQC+ secures the data itself.

Modern cybersecurity has a single, unfixable flaw: once an attacker is inside the perimeter, they typically have free rein. Firewalls, intrusion detection, network segmentation — all of these protect the "pipes." None of them protect

the "water" flowing through the pipes. The Holy Grail of cybersecurity has always been a system that secures the data itself, regardless of where the data travels or who manages to steal it. That is what PQC+ delivers.

From securing pipes to securing water

Traditional security: protect the network, hope the data stays inside. PQC+ security: make the data its own portable vault. Even if the network is breached, even if a database is exfiltrated, even if an insider walks out with a hard drive, the data cannot be opened without the correct identity, hardware, and active context.

What this delivers across industries

- **Healthcare:** Individual-centric security. Even if a hospital server is breached, individual medical records remain encrypted blobs that can only be opened by the specific clinician and individual's matched hardware keys.
- **Open Finance:** Secure data exchange where consent is physically baked into the data packet. A bank cannot share what the ACL does not permit — even by accident.
- **Industrial, drones, supply chain:** Hardware-defined keys prevent spoofing. An attacker cannot send a fake command to a drone or a factory robot because they do not possess the physical hardware key embedded in the instruction's ACL.
- **Oil and gas:** Protecting proprietary geological data for decades requires the 10240-bit key tier to withstand quantum computers that do not yet exist.

Why this is historically significant

The shift from perimeter security to data-centric security is the same scale of change as the shift from castles to mobile warfare. Castles were eventually irrelevant. Perimeters are becoming irrelevant for the same reason: the attackers can always get inside. PQC+ accepts this reality and protects what actually matters — the data — directly.

> **The bottom line**
>
> By tying security to the data itself, PQC+ solves the fundamental flaw of modern breaches. The treasure becomes impossible to touch — not because the walls are higher, but because the treasure itself is locked at the cryptographic, identity, hardware, and context layers simultaneously.

↑ Back to Table of Contents

PQC+™

Transformativ IP

Regulatory Compliance & HNDL Protection
Federal & all 50 States within 90 Days