February 2026

ICIT

Quantum-resilient convergence:

The shared defense of AI, space, and critical infrastructure

David Mussington, Ph.D., CISSP, DDN QTE
ICIT Fellow, Co-Chair, ICIT FCEB Resilience Center

www.icitech.org

Table of Contents

About ICIT

The Institute for Critical Infrastructure Technology (ICIT) is a nonprofit, nonpartisan, 501(c)3 think tank with the mission of modernizing, securing, and making resilient critical infrastructure that provides for people's foundational needs. ICIT takes no institutional positions on policy matters. Rather than advocate, ICIT is dedicated to being a resource for the organizations and communities that share our mission. The views and opinions expressed in this essay are solely those of the author(s) and do not necessarily reflect the official policy or position of ICIT. Any assumptions made within the analysis are not reflective of the position of any entity other than the author(s). To learn more, please visit www.icitech.org

CyberRisk ALLIANCE | Thank you to our Strategic Partner **CyberRisk Alliance** | cyberriskalliance.com

Executive summary

1. The strategic imperative

AI data centers and Low Earth Orbit (LEO) ground segments have evolved into the coordination layer for national critical infrastructure, sharing control planes with power grids, financial networks, and logistics systems. Consequently, the migration of these assets to Post-Quantum Cryptography (PQC) is no longer a niche IT upgrade but a **Tier-1 resilience requirement**.

This document argues that AI infrastructure and PQC migration are inextricably linked. They share the same capital expenditure cycles, the same 2030–2035 planning horizon, and the same operational necessity for deep visibility and control.

2. The risk: Integrated attack surface

We face a "**Harvest-Now, Decrypt-Later**" threat environment. Adversaries are actively archiving encrypted traffic today — including proprietary model weights, long-lived telemetry, and firmware signing artifacts — to decrypt them once quantum capabilities mature.

- **Concentration risk**: AI clusters centralize high-value intellectual property and critical control paths, creating a single point of failure for long-term confidentiality.
- **The LEO bridge**: Vulnerabilities in AI control planes provide indirect access to satellite constellations and OT environments. We must assume adversaries target the entire data path — from ground terminals and RF links to the satellite bus itself — treating these "high-speed passthroughs" as active attack surfaces.

3. The solution: AI as the migration engine

Contrary to viewing AI solely as a risk, this analysis posits that **AI infrastructure is a strong substrate capable of executing PQC migration at scale**. The observability, automated policy enforcement, and telemetry pipelines native to modern AI/Cloud platforms are tools that can analyze and correlate information provided by Automated Cryptographic Discovery and Inventory (ACDI) toolsets, significantly supplement the discovery of cryptographic dependencies (CBOM), enforce crypto-agility, and detect "downgrade" attacks in real-time.

4. Key recommendations & levers

- **Align procurement with GSA standards:**
  Security is now an acquisition problem. Program offices must integrate requirements from the GSA PQC Buyer's Guide (2025) directly into RFI/RFPs for AI and LEO refreshes, mandating Cryptographic Bill of Materials (CBOM) delivery and defined migration milestones.
- **Manage the hardware "Valley of Death" (2025–2028):**
  We face an immediate supply chain gap where current AI accelerators and SmartNICs lack native PQC support. The strategy must explicitly require "crypto-agile wrappers" or software-hybrid modes for hardware procured in this window to prevent locking in a generation of quantum-vulnerable silicon.
- **Navigate geopolitical fragmentation:**
  Operators must design for a splintered regulatory landscape. Systems spanning US, EU, and PRC jurisdictions will require parallel PQC stacks to satisfy diverging sovereignty and lawful access mandates.

5. Conclusion

The window to secure the 2030s is open now. If PQC is treated as a "bolt-on" after AI infrastructure is built, the cost and complexity of remediation will be prohibitive. Automated Cryptographic Discovery and Inventory of encryption vulnerabilities must begin now. This is delivered by ACDI toolsets, enhanced by AI, whose outputs can then be further consumed by AI for analysis, correlation, and agentic use-cases. By embedding cryptographic requirements into the current wave of AI and LEO build-outs, we can transform a potential vulnerability into a defensible, quantum-resilient posture.

1. Introduction

AI data centers and LEO/ground-segment sites are turning into the coordination layer for everything else. They host the models that shape decisions, the data sets that organizations cannot easily rebuild, and the control systems that reach into power grids, payment systems, logistics networks, and satellite constellations.[1][2]

When those environments talk to each other, they rely on the same classical public-key cryptography that has been in place for decades — TLS for APIs, VPNs for management and backhaul, code-signing for firmware and software, PKI for identity. Those are exactly the mechanisms that a future, capable quantum adversary can retrospectively break if they have captured the traffic and artifacts today.[1][3]

Defining the AI cluster

For the purposes of this analysis, an **AI cluster** is defined as a high-performance computing environment characterized by massive parallelism, specialized accelerators (GPUs/TPUs), and high-bandwidth interconnects. These environments host the **AI/ML models that drive autonomous critical decisions** and process sensitive training data. Unlike standard enterprise networks, they utilize unique telemetry and control planes that require distinct security architectures."

The basic argument in this piece is that AI infrastructure and post-quantum migration are now joined at the hip. They share the same time window — roughly the 2030–2035 period that NIST and the NCCoE are using as the planning horizon for getting high-value systems off purely classical crypto.[1][4][5]

They also share the same operational assumptions: visibility into what is actually running, centralized configuration and policy, the ability to roll out changes gradually and measure the impact, and enough telemetry to know when something has gone wrong.[1][6] Those are not properties of a typical legacy enterprise network; they are characteristics of mature AI/cloud platforms.

That leads to two simple but uncomfortable conclusions. First, the AI and LEO infrastructure that is being built and refreshed over the next decade will either be designed to make PQC practical — meaning cryptographic choices are visible, controllable, and monitored — or it will bake in patterns that are fundamentally quantum-vulnerable and very expensive to reverse later (for example, long-lived signing keys that cannot be changed, or protocol stacks that only "kind of" support PQC and silently fall back to classical).[1][7]

Second, attackers will not wait for standards to fully settle; they will actively look for these gaps and delays and incorporate them into how they discover, prioritize, and hold access in and around AI clusters.[8][9]

The rest of the paper unpacks that claim in a structured way:

Section 2 looks at why AI expansion and PQC migration are on the same clock: the transition timelines NIST is using, the refresh cycles for data centers and ground segments, and the way US, EU, and PRC choices around standards, cloud regulation, and lawful access are already pulling the technical options in different directions.[10][11]

Section 3 walks down the stack and asks, concretely, where PQC actually shows up: in protocols like TLS and IKE, in shared crypto libraries, in HSMs and accelerators, in PKI and service meshes, in application-level formats and firmware. It is explicit about where "hybrid" deployments and partial migrations create real exposure rather than reducing it.[1][5]

Section 4 focuses on the risk surface and the adversary. It looks at what is actually being concentrated in AI clusters—long-lived data, control planes, update channels—and how harvest-now, decrypt-later thinking naturally extends into long-term integrity attacks once classical signatures and key-establishment can be broken. It also treats AI, LEO, and OT as one connected surface, not three separate domains.[1][2]

The last three sections shift from description to levers:

Section 5 treats AI assets as the natural engine for PQC: the place where ACDI-derived cryptographic inventory, staged rollout, and anomaly detection can actually run at the needed scale.[12][13]

Section 6 looks at how different oversight roles—critical-infrastructure security, intelligence, law enforcement and CALEA, regulators, and procurement/IT governance—shape what is realistic to deploy, and where there are built-in tensions that must be surfaced rather than ignored.[10][6][14]

Section 7 then zooms in on supply-chain and acquisition, on the assumption that the real inflection points are buried in RFP language, vendor roadmaps, CBOM and telemetry expectations, and how firmly PQC requirements are tied to AI and LEO refresh cycles.[15][16]

The through-line is that by the mid-2030s, the systems that are meaningfully quantum-resilient will not be the ones with the "best" algorithm on a slide; they will be the ones where cryptography was treated as a design and procurement property of AI and space/OT infrastructure from the start.

2. Shared drivers, timelines, and geopolitics

AI data center build-outs and PQC transition are effectively the same program of record on different letterheads. Both are long-horizon, capex-intensive transformations that land in the 2030–2035 window NIST, OMB, and CISA/NSA/NIST are now implicitly treating as the deadline for quantum-relevant systems.[1][17][9]

Timelines and refresh windows

- NIST IR 8547 assumes 10–20-year crypto transitions and explicitly anchors PQC changes to major infrastructure refreshes — protocol stacks, HSMs, PKI, directory services, and core applications.[1][7] That is precisely where hyperscale AI expansions and LEO/OT digitalization are already funded.
- The CISA/NSA/NIST quantum-readiness work and NIST's CSWP 48 mappings push cryptographic inventory, risk assessment, and vendor engagement into the same planning cycles that govern data center and ground-segment modernization, making PQC planning a first-class component of AI/cloud governance rather than a parallel track.[9][6][14]

Compliance and risk-framework integration

- NIST CSWP 48 maps NCCoE's PQC migration capabilities directly into NIST CSF 2.0 and SP 800-53 controls — crypto inventory, crypto-agility, telemetry, key management — turning PQC from an R&D topic into a set of expectations auditors and regulators can assess.[6][14]
- For AI clusters and LEO/ground-segment operators, that means PQC readiness will show up as part of "normal" risk and compliance reporting: inventory completeness, migration coverage on high-value flows, crypto-telemetry depth, and vendor roadmap quality.

Geopolitical alignment and divergence

- US-aligned ecosystems are converging on the NIST portfolio — ML-KEM (FIPS 203) and ML-DSA/related signatures (FIPS 204/205) — with NCCoE patterns for PQ-TLS/SSH/IKE and PQ-ready HSMs as the reference implementation.[1][4]
- Other blocs are not standing still: PRC-aligned infrastructures are moving toward domestic PQC stacks for clouds and intelligent computing centers under their own regulatory regimes; EU regulators are increasingly explicit about digital sovereignty expectations around key residency, HSM auditability, and cloud/space PKI.[11][10]
- For AI data centers and LEO ground nodes operating across jurisdictions, the implication is straightforward: expect to run multiple PQC and key-governance profiles in parallel, and expect those choices to be read as geopolitical alignment signals, not just technical selections.[10][11]

3. Technical-stack overlap and operational pitfalls

"NIST IR 8547 slices PQC transition into five layers---protocols, software crypto libraries, cryptographic hardware, PKI/infrastructure, and applications/services. AI clusters span all five, plus a 'shadow layer' of firmware and embedded control that will ultimately bound how much **quantum risk can be systematically mitigated**."

3.1 Protocols: Where PQ KEMs Land First

Almost every control and data path in an AI cluster terminates on a small set of protocols that are now the primary PQC battleground.[1][4]

TLS, QUIC, SSH, IPsec/IKE

- Public APIs, operator consoles, and many service-mesh edges ride on TLS 1.3 and QUIC; admin and automation flows run over SSH; regional backhaul and some ground-segment paths use IPsec/IKE.[4][5]
- IETF work is actively defining PQC recommendations and profiles: PQ and hybrid KEMs for TLS-based protocols, hybrid KEMs for IKEv2, and experimental PQ key-exchange for SSH.[5][11]

For US-aligned stacks, FIPS 203 (ML-KEM) is the anchor KEM; other families appear in niche or jurisdiction-specific roles.[1] KEMs dominate handshake size, latency, and CPU cost — exactly the dimensions that matter for API gateways, dense service meshes, and satellite-affected links.[5]

The immediate pitfall is dual-stack crypto: these protocols will support both classical and PQ (or hybrid) variants for an extended period. Without hard policy and end-to-end telemetry, "supporting PQC" can coexist with quiet fall-backs to classical on critical paths.[6][14]

3.2 Libraries, frameworks, and orchestration

Below the protocols sit the crypto libraries and frameworks that actually implement ML-KEM, ML-DSA, and alternatives.[1][7]

Shared stacks across AI services

- AI frameworks, data platforms, control-plane components, and observability stacks typically depend on a small set of crypto libraries (OS crypto APIs, OpenSSL descendants, language runtimes). Once those stacks expose PQ KEMs and signatures, the estate can often adopt PQC via configuration and policy rather than bespoke rewrites.[1][4]

- NCCoE's TLS/SSH work and CSWP 48's capability mappings assume this model: push PQC into common primitives, then surface it through policy, crypto-agility mechanisms, and central configuration, not hand-coded crypto in every service.[6][12]

The failure modes are familiar: pinned old library versions, local cipher-suite overrides, and hard-coded algorithm IDs and key sizes that prevent services from taking advantage of new PQC capabilities even when the platform supports them.[7]

3.3 Cryptographic hardware: HSMs, TPMs, and offload engines

Cryptographic hardware is both a trust anchor and a schedule constraint.[1]

HSM/TPM/KMS realities

- AI data centers already lean on HSMs, TPMs, and KMS for key storage, certificate issuance, KDFs, and code-signing. PQC migration requires those devices to implement ML-KEM/ML-DSA (or approved equivalents), expose PQ-capable key slots, and handle larger keys and signatures without collapsing throughput or SLA margins.[1][4]
- This implies firmware updates and new validations (for example, FIPS 140-3 coverage for PQ algorithms), and changes to shard/backup/recovery patterns for keys in multi-tenant environments.[4][5]

Some hardware will be upgradable via firmware; some will need replacement. A non-trivial subset of accelerators and SmartNICs embed fixed-function crypto offload that will never speak PQC. Those components either move out of the critical path or become permanent quantum-vulnerable enclaves in otherwise modernized systems.[4]

3.4 PKI, identity, and mesh trust

PKI and identity infrastructure encode who is allowed to do what where, and how that trust is adjudicated.[1][6]

Certificate chains and service identity

- PQC migration here means issuing certificates with PQ signatures (often hybrid/composite during transition), updating clients and middleboxes to understand new algorithm identifiers and key formats, and ensuring revocation/enrollment mechanisms handle PQ CAs correctly.[6][14]
- Service meshes and workload-identity systems must carry PQ keys and identities, rotate them on mesh tempos, and enforce policies that do not quietly accept classical-only chains for high-value paths.

Partial modernization is the obvious trap: CAs that issue PQC leaves but retain classical-signed roots/intermediates; meshes that present PQ identities but chain them back to classical trust anchors. On the surface the estate is "PQC-enabled"; in reality, the root of trust remains quantum-vulnerable.[1][7]

3.5 Applications, artifacts, and the firmware layer

At the top, NIST groups "applications and services," but in AI clusters you have to include firmware and embedded control if you want an honest view of residual risk.[1][7]

Application-level cryptography and formats

- AI-adjacent applications implement client-side encryption, key-wrapping, and signed artifacts (model packages, pipeline manifests, data exports). Unless they move to PQC-aware libraries and formats, they quietly preserve classical exposure above modernized transports and PKI.[4]
- Any data format that embeds keys or signatures must be extended to carry PQ material and, where hybrid is used, to encode it unambiguously enough for enforcement and analytics.[6]

Firmware and embedded control

- Accelerators, SmartNICs, BMCs, satellite modems, and OT gateways are governed by cryptographic boot and update chains that often rely on long-lived classical signatures with minimal crypto-agility.[1]
- If those chains cannot be re-keyed or redesigned with PQC (or at least robust hybrid) within operational lifetimes, they become hard structural limits on how quantum-resilient an AI cluster or ground segment can ever be. In a 2030–2035 horizon, expect a disproportionate share of residual risk to sit here.[1][7]

3.6 Hybrid modes and partial migration: Where things break

Hybrid modes are the transitional reality, but they introduce their own complexity budget.[1][5]

Hybrid done right vs. hybrid-in-name-only

- Proper composite constructions require both classical and PQ components to fail, providing real defense-in-depth during the transition. Many proposed "hybrids," however, effectively bolt PQ onto an unchanged classical scheme, leaving a single classical failure point under a "quantum-safe" label.[5][11]
- Certificate size and path-length constraints — especially in IoT, LEO, and middlebox-constrained environments — limit how aggressively composite chains can be deployed, forcing trade-offs about where hybrids are actually viable.[7]

Performance, MTU, and fragmentation

- Larger keys and signatures increase handshake sizes and can trigger MTU/fragmentation issues for TLS, QUIC, and IKE, particularly across middleboxes and satellite links; this is an operational issue, not an academic footnote, in AI+LEO topologies.[4][5]
- The pain is not uniform: dense control planes with many short-lived connections (microservice RPC, API gateways) are far more sensitive to KEM overhead than long-lived bulk channels.

Partial migration as an attack surface, not a "phase"

- In AI clusters, uneven adoption across layers is the baseline. Protocols may support PQC while hardware does not; PKI may issue PQC leaves while roots remain classical; management VPNs and firmware signing chains may not move at all.[1][6]
- In that world, "we support PQC" is not a useful claim. The meaningful questions are: which algorithms, on which flows, anchored to which trust chains, with what downgrade behavior and what telemetry?

4. Quantum risk concentration and adversary tradecraft

AI clusters are where quantum risk, state tradecraft, and infrastructure fragility converge. They concentrate the secrets and control surfaces that NIST and CISA/NSA/NIST explicitly flag for "harvest-now, decrypt-later," and they sit inside cloud, OT, and space topologies that nation-states already treat as pre-positioning terrain.[1][9]

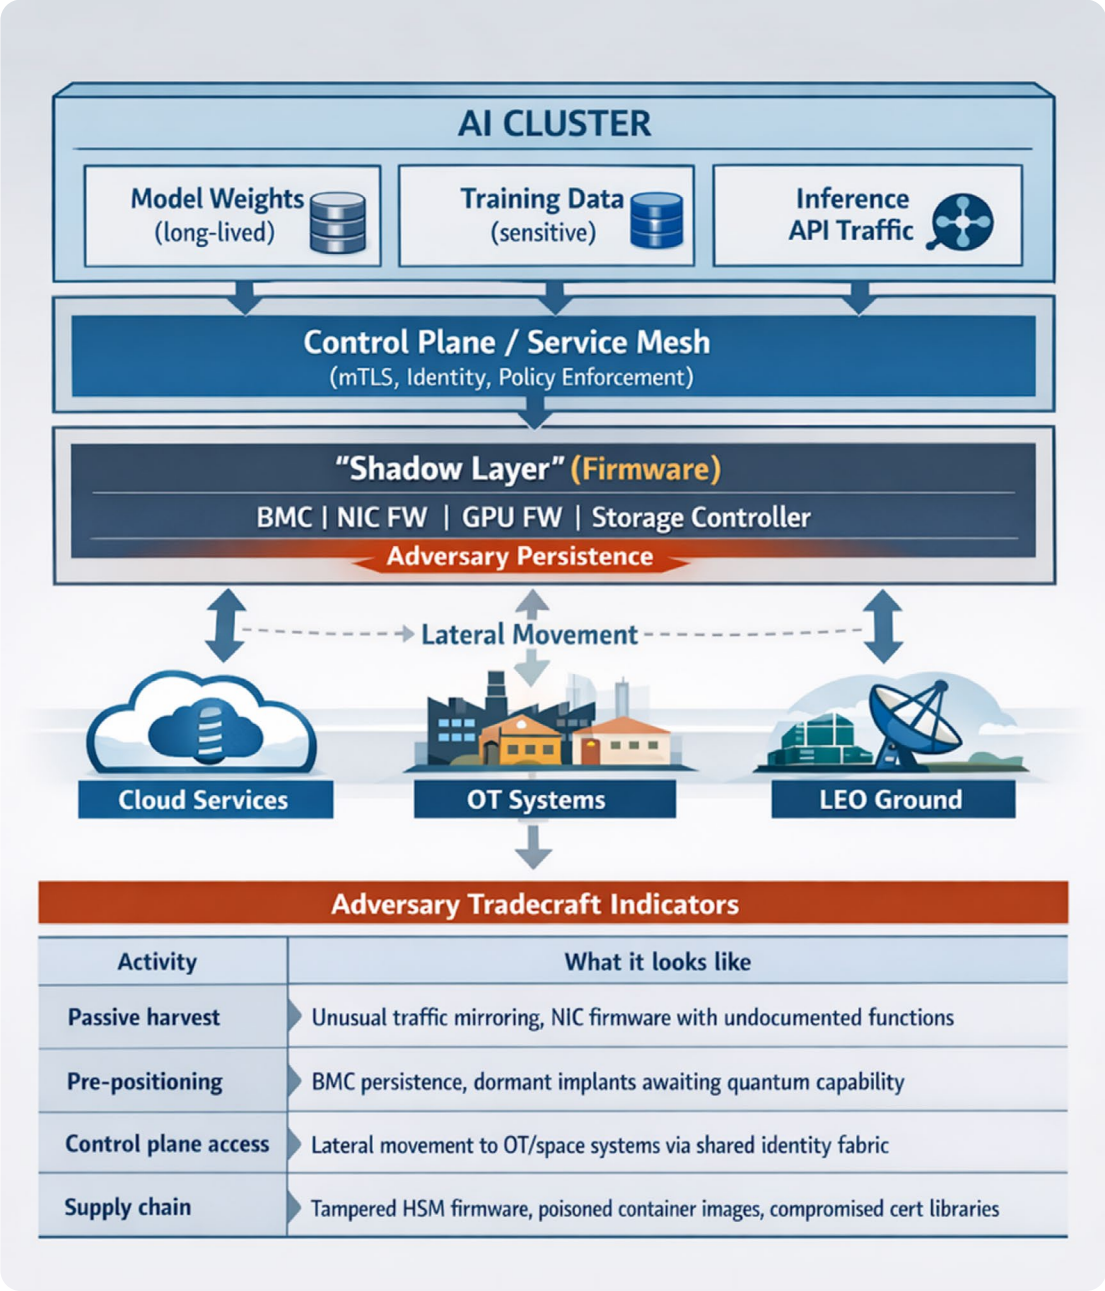

Adversary Tradecraft Indicators

Activity	What it looks like
Passive harvest	Unusual traffic mirroring, NIC firmware with undocumented functions
Pre-positioning	BMC persistence, dormant implants awaiting quantum capability
Control plane access	Lateral movement to OT/space systems via shared identity fabric
Supply chain	Tampered HSM firmware, poisoned container images, compromised cert libraries

Figure 1. AI Clusters

4.1 What is actually concentrated in AI clusters

From a quantum-risk perspective, AI data centers are not generic compute — they are dense aggregations of long-lived data and high-leverage control paths.

Long-lived data at rest and in motion

- Model weights, proprietary training corpora, telemetry lakes, and cross-domain log archives persist for years and are reused across model generations and services, giving adversaries a confidentiality horizon that aligns with plausible quantum timelines.[1][7]
- These assets ride on exactly the key-establishment and bulk-encryption mechanisms NIST prioritizes for early PQC transition: TLS/IKE/VPN for API ingress and east-west traffic, storage encryption for data lakes, and inter-region replication paths that may lack robust forward secrecy.[1][3]

Control planes and update/signing channels

- Cluster control APIs, orchestration backbones, and cross-domain interfaces (including LEO ground links and IT-OT bridges) route through AI-adjacent networks; compromise here yields durable administrative control over both digital workloads and linked physical systems.[1][2]
- Firmware and software update flows for accelerators, SmartNICs, BMCs, satellite modems, and OT gateways often depend on a small number of long-lived signing keys and CA hierarchies — the "hard to upgrade" crypto surfaces NIST calls structurally risky in a PQC transition.[1][7]

The combined picture is straightforward: AI clusters host durable confidentiality targets and privileged integrity targets in the same footprint.

4.2 How state and proxy actors will exploit PQC lag

Given the direction of cloud-centric APT activity and the rapid normalization of AI-assisted tooling, it is reasonable to assume PQC lag and downgrade paths will be first-class campaign inputs, not incidental findings.[8][9]

PQC lag and downgrade mapping as a product

- Automated (and increasingly LLM-augmented) tooling can fingerprint crypto posture across exposed and internal surfaces: classical-only TLS at API edges, legacy VPNs on management planes, non-PQC IKE on inter-region or ground-segment links, signing flows pinned to non-agile HSMs.[8][4]
- Over time, that yields a PQC lag map by operator, region, and function: which clusters remain harvested behind classical key-establishment, which control-plane paths accept silent downgrade from PQ-capable to classical, and where long-lived signing keys are likely to survive beyond 2035.[1][7]

Agentic orchestration of pre-positioning and persistence

- Agentic offensive frameworks can treat cryptographic posture as a scoring feature: prioritize footholds where control-plane and east-west traffic remain classical, and where firmware and boot chains are signed with non-agile or weakly governed schemes.[8]
- Campaign logic can be tuned to lock in access behind OT/LEO gateways whose VPNs and management channels will lag PQC, and to prioritize exfiltration of captures and key material from AI assets least likely to complete migration on time.

Supply-chain and firmware-level quantum exposure

- Front-ending services with PQ-TLS and PQ-VPN does not neutralize legacy signing infrastructures and classical-only supply-chain components. If an adversary archives signed images and control-plane traffic now, and those artifacts rely on classical signatures that become breakable, they can mint "legitimate" updates or credentials later against systems that otherwise kept up at the protocol layer.[1][7]
- The result is a long-horizon attack path: harvest firmware, control-plane captures, and signing artifacts today; when classical cryptography is practically breakable, repurpose that material to seize AI clusters and ground nodes via apparently valid updates and identities.[1]

4.3 AI, LEO, and OT as a shared quantum-risk surface

AI clusters are increasingly co-located with, or logically coupled to, LEO ground segments and OT environments via shared facilities, fabrics, and bridge points.

- PQC gaps at AI nodes — classical inter-region links, non-PQC management channels, non-agile signing — provide indirect access paths into satellite control, GNSS analytics, and OT control services, even if those downstream systems are partially hardened.[2][11]
- Classical choke points in LEO and OT paths (non-PQC VPNs on ground-segment backhaul, legacy crypto on telemetry links) can grant durable visibility into, or leverage over, AI workloads that depend on those links for command, sensing, or data ingestion.[2][4]

In practice, PQC lag or misconfiguration around AI clusters is not a localized defect; it is a structural vulnerability across an integrated AI–LEO–OT surface that state and proxy actors can systematically map and exploit with AI-enabled tooling.[11]

5. AI as PQC migration engine and sensing layer

AI infrastructure is not just something to be "wrapped" in PQC; it is a powerful operational substrate that accelerates PQC migration at scale. NCCoE's migration work effectively assumes the kind of telemetry, topology, and control planes that hyperscale AI assets already operate.[6][12][13]

5.1 Discovery and CBOM at AI scale

NCCoE's starting point is unglamorous but decisive: without automated cryptographic discovery, everything else is theater.[6][13]

Cryptographic inventory as a graph problem

- Extend existing asset and dependency graphs (services, meshes, subnets, LEO/ground nodes, OT gateways) with cryptographic attributes — algorithms, key sizes, protocol versions, libraries, HSM bindings, CA chains, rotation policies.[6][13]
- Treat this as a cryptographic bill of materials (CBOM) problem: every service, firmware image, library bundle, and control-plane path carries a cryptographic profile that can be queried, scored, and mapped to risk frameworks.[13][14]

AI-assisted inventory and clustering

- Apply ML to telemetry and ACDI scan outputs (PCAPs, TLS fingerprints, SSH banners, binary/code scans, HSM logs) to automatically cluster usage patterns — classical-only TLS edges, firmware signing tied to legacy HSMs, inter-region tunnels with no PQ KEMs, and so on.[13][14]
- Map these clusters into CSWP 48–style capability mappings so cryptographic posture is explicitly tied into CSF 2.0 and SP 800-53 controls, rather than trapped in spreadsheets.[6][14]

AI observability and data-engineering pipelines become an ACDI and ACDI output augmentation for PQC: cryptography stops being invisible plumbing and becomes part of the system graph.

5.2 Using AI ops patterns to stage PQC

Once the cryptographic map exists, PQC migration looks a lot like any other large-scale change in a mature AI estate: staged rollouts, policy-driven routing, and continuous measurement.[12][13]

Treat PQC as a feature, not a flag

- Encode PQC policies into the same declarative configuration systems used for service meshes and CI/CD — for example: "require ML-KEM on inter-region replication of model weights," "enforce PQ signatures for firmware signing hierarchies," "disallow downgrade on defined control-plane SNI/paths."[6][12]
- Use canaries and dark launches: bring up PQ-TLS/SSH/VPN alongside classical, mirror production load, and observe latency, handshake behavior, error rates, and downgrade patterns before flipping over traffic.[12]

Figure 2. AI Ops and Risk Tracking

Exploit AI hardware and schedulers for measurement

- Use spare capacity and flexible schedulers to run synthetic PQC load across realistic topologies, including LEO/ground-segment links, instrumented with NCCoE's dimensions (CPU, memory, bandwidth, handshake rates).[6][13]
- Feed empirical data back into policy: where overhead is acceptable, enforce hard cut-overs; where it is not, explicitly document residual classical exposure and timelines in CBOM and risk registers, rather than letting legacy paths persist by inertia.[6][14]

The operational loop is: discover → stage → measure → ratchet coverage — using the AI estate's own control and telemetry fabric as the engine.

5.3 Cryptographic anomaly detection and PQC assurance

PQC posture is not static; it will be subject to misconfiguration, regression, and deliberate downgrade. NIST's mappings and CBOM discussions place cryptographic telemetry — what algorithms, keys, and chains are used — into the set of first-class signals that must be monitored.[6][14]

Behavioral baselines for cryptography

- Train models on "known-good" handshake and certificate behavior across AI ingress/egress, internal meshes, LEO ground links, and OT bridges: which cipher suites appear where, which CAs/issuers are legitimate for which roles, what normal key-usage patterns look like in HSMs.[6][12]
- Flag deviations such as unexpected re-introduction of classical-only suites on high-value paths, issuance from unusual or non-approved CAs, or **signals indicative of side-channel attacks and fault injection attempts aimed at forcing a downgrade** (for example, anomalous spikes in signing key usage).

Connect crypto anomalies to operational change and supply chain

- Correlate cryptographic anomalies with actual changes — firmware updates, library releases, new dependencies, tenant onboarding — so PQC regressions appear as concrete change-management failures rather than abstract "crypto concerns."[6]
- Treat quantum-relevant anomalies (for example, satellite backhaul falling back from PQ-capable IKE to classical IKE under load) as high-priority incidents with runbooks that pull in the governance and oversight roles described in Section 6.[13]

Hardware-rooted baselines

- To harden these baselines against impersonation and key exfiltration, apply FPGA security patterns such as Physically Unclonable Function (PUF)-backed device identity and key-binding. Enroll 'known-good' endpoints with hardware-bound keys so that anomalous cipher-suite selection can be attributed to a specific physical device and correlated with firmware integrity signals."

In this model, AI infrastructure is the nervous system for PQC: it knows where cryptography lives, can actuate changes at scale, and can detect when that posture is drifting or being attacked.

6. Governance, sovereignty, and functional oversight

PQC for AI data centers and LEO/ground-segment infrastructure is not just a cryptographic selection problem; it is a governance and oversight problem with multiple constituencies pulling in different directions. Resilience, intelligence gain/loss, lawful access, competition, and sovereignty all press on the same cryptographic substrate.[9][10]

6.1 Governance and supply-chain framing

At the governance level, PQC is framed as a cross-cutting risk-management and supply-chain program.

From crypto hygiene to risk register entry

- Cryptographic inventories and CBOMs are tied explicitly to asset and data inventories and rolled into enterprise risk registers and board-level reporting, rather than living as specialist artifacts.[6][14]
- PQC requirements — algorithms (for example, ML-KEM/ML-DSA or local equivalents), crypto-agility, PQ-capable HSM/PKI, telemetry, migration milestones—are encoded in acquisition language, configuration baselines, and external service contracts for AI clusters, network fabrics, LEO ground equipment, and OT gateways.[6][15]

Digital sovereignty and lawful access as design inputs

- Divergent expectations across US/EU/PRC and others about key location, HSM residency, regulator access to PKI/telemetry, and lawful-access hooks drive different PQC stack and key-governance choices.[10][11]
- For cross-border AI and LEO operators, that implies running multiple PQC profiles and governance models side-by-side, with crypto-engineering choices read as geopolitical commitments as much as technical decisions.[10]

6.2 Functional oversight roles

Thinking in terms of functional roles clarifies how oversight pressures translate into concrete crypto-engineering constraints.

Critical infrastructure security and resilience functions

- Define sector baselines for quantum readiness across AI, cloud, LEO, and OT dependencies, elevating cryptographic posture (key-establishment, signatures, firmware signing) to a core resilience metric.[9][11]
- Drive continuous crypto-agility: required inventories, PQC roadmaps, exercises, and minimum cryptographic telemetry/logging in AI and ground-segment environments.[6]

Intelligence and national security functions

- Provide deep technical guidance on algorithm families, hybrid constructions, and key-management patterns suitable for high-value AI and LEO/OT systems, including what residual classical exposure is tolerable for which missions.[1][12]
- Map foreign PQC adoption and harvest-now/decrypt-later campaigns around AI, cloud, and space systems, feeding that back into domestic prioritization (where to accelerate PQC, where classical can persist, and what tradecraft to expect—especially downgrade hunting and supply-chain abuse).[11]

Law enforcement and public safety functions (CALEA-constrained)

- Ensure PQC adoption in AI platforms, communications, and ground-segment systems remains compatible with targeted, court-authorized access under CALEA-style regimes, without demanding algorithmic backdoors or permanent classical "side doors."[9]
- Push architectures where lawful access is implemented at endpoints, mediation services, or tightly governed intercept functions, while transport and key-establishment layers remain quantum-resistant; cryptographic logs and CBOMs are part of the evidentiary chain, not a vector for weakening crypto.[9]

Regulatory and oversight functions (communications, financial, space, privacy)

- Integrate PQC posture into licensing, prudential supervision, incident-reporting, and disclosure obligations for AI-intensive and satellite-dependent services, including expectations around PQC incident detection, downgrade monitoring, and reporting.[14][11]
- Make digital sovereignty operational: require auditability of HSM/PKI and key-management for extraterritorial cloud and ground-segment operations — who holds which keys, under which law, with what regulator visibility and recourse.[10][15]

Standards, procurement, and civilian IT oversight functions

- Map NIST's PQC migration capabilities into CSF 2.0, SP 800-53, and reference architectures explicitly tailored to AI clusters and LEO ground segments, including expectations for PQ-TLS/QUIC/SSH/VPN, PQ-ready HSMs, crypto-agile firmware signing, and cryptographic telemetry.[6][14]
- Use federal and quasi-federal buying power to normalize PQC-ready components and telemetry: decline to procure accelerators, SmartNICs, satellite modems, or AI control-plane software that cannot publish CBOMs, support crypto-agility, and meet PQC support timelines aligned with system lifetimes.[15][16]

The upshot is that crypto-engineering decisions in AI and LEO environments increasingly sit under these functional pressures; "what curve do we use?" is the least interesting part of the problem.

7. Supply-chain and procurement levers

Sections 3–6 described the technical and governance landscape. Section 7 is about the policy and shaping levers: how AI and LEO-adjacent programs use acquisition and supply-chain governance to force PQC into real hardware, software, and services on meaningful timelines.[6][15][18]

7.1 Make PQC a first-class procurement requirement

AI data center and ground-segment programs are where the money is, which means they are where meaningful requirements stick.

Push PQC into the front of the acquisition process

- RFI/RFPs should explicitly call out cryptographic capabilities: supported KEM/signature families (for example, FIPS 203/204/205 or accepted equivalents), PQ profiles for TLS/QUIC/IKE/SSH, crypto-agility features, and support for the inventory and telemetry functions outlined in Sections 3 and 5.[6][13]
- "The **GSA PQC Buyer's Guide (2025)** provides the authoritative federal baseline for these requirements. It is no longer just a recommendation but the standard for defensible procurement. AI and LEO program offices should copy-paste its requirements for CBOM export, PQ-capable HSM/PKI, and structured cryptographic telemetry directly into acquisition documents. Vendor roadmaps must align with the milestones in the **GSA PQC Roadmap**, specifically the 2030–2035 transition window."[15][18]

Use AI and LEO RFPs as market-shaping instruments

- For chipset vendors, NIC/switch OEMs, space integrators, and platforms, these contracts are strategic. If the RFP says "no PQC roadmap, no bid," PQC is effectively a market entry condition.[15]
- The same principle applies to OT and LEO integration wrapped around AI workloads: no new control stack, gateway, or ground-station program without a credible PQC plan for its control, telemetry, and signing paths and integration of its crypto telemetry into the AI-scale sensing plane.[18]

7.2 Turn roadmaps and CBOM into hard edges

Visibility and time-phased commitments are where procurement can move beyond checkbox security.

Roadmaps as contractual artifacts

- Serious PQC roadmaps largely agree on structure: inventory/pilots (mid-2020s), broad pilots and early cut-overs (late-2020s), critical-system transition (early-2030s), cleanup by around 2035.[1][5]
- AI and LEO/OT contracts should force this down to named components: which firmware families, protocol stacks, management planes, and HSMs will support which PQC profiles, by which dates, with what assurance (for example, validated modules). Slippage then becomes a contract and performance problem, not an amorphous "we're evaluating."[6][15]

CBOM and cryptographic SLAs

- CBOM is the right abstraction for ongoing visibility: a machine-readable description of algorithms, key sizes, protocol uses, PKI dependencies, and HSM usage for each product or service.[13][15]
- For AI clusters and ground segments, acquisition should require:
  - Initial CBOM at purchase and updated CBOM on major releases.
  - Alignment with CSWP 48 / CSF-mapped capabilities (coverage of key-establishment, data-at-rest, code-signing, update channels).[6][14]
  - Cryptographic SLAs: specific classical algorithms retired by defined dates, PQ KEM required on listed flows, and mandatory emission of structured crypto telemetry to support inventory and anomaly detection.[15]

7.3 Attach PQC to existing capex cycles

PQC is tractable when it rides the refresh waves organizations are already committed to.

Tag lifecycle events as PQC hooks

- NIST and NCCoE explicitly recommend aligning PQC with refresh windows for hardware and software — protocol stacks, HSMs, PKI/IDM, core applications.[1][6] In AI/LEO/OT contexts, this includes:
  - GPU/accelerator and SmartNICs generations that allow retirement or isolation of non-upgradable crypto offload.
  - Ground-station and OT control upgrades that replace non-agile VPNs, modems, and management planes.
  - Major AI platform/service-mesh/observability revisions that can absorb PQC and crypto-telemetry changes.
  - Portfolio and procurement functions should mark these as PQC decision points by default; skipping PQC at those points should require explicit justification.[6][5]

Current AI accelerators and SmartNICs often rely on fixed-function crypto offload engines that cannot support FIPS 203/204 algorithms. For systems deployed in this window, acquisition must mandate 'crypto-agile wrappers' or software-based hybrid modes that allow PQC to run on general-purpose cores, bypassing the legacy hardware offload. This prevents a generation of 'zombie hardware' that is permanently locked to classical crypto."

Contain what cannot realistically move on time

- Some components — certain LEO payloads, deeply embedded OT devices, third-party managed systems — will not be replaceable or re-keyable by 2035. They should be treated as constrained classical enclaves with tight interfaces, explicit cryptographic compensating controls, and clearly articulated residual risk.[7][18]
- For AI workloads dependent on such enclaves for data or control, that residual risk must surface into model-governance and mission-risk decisions, not vanish under "infrastructure assumptions."[1][6]

7.4 Coordinate within sectors, absorb friction across jurisdictions

Supply-chain levers are most effective when there is sector-level coherence but must also survive a fragmented regulatory environment.

Sector-level Alignment on Minimum PQC Posture

- Sector guidance (finance, energy, telecom, health, space) can define a baseline PQC posture for AI, cloud, and LEO/OT dependencies: algorithm deprecation dates, PQ KEM requirements for inter-operator links, minimum CBOM and crypto-telemetry expectations.[9][11]
- Shared infrastructures — payment rails, clearing houses, satellite operators, common AI platforms — can publish PQC timelines so connected parties can align dependencies rather than each negotiating separately.[11]

Design for PQC fragmentation across jurisdictions

- US, EU, and PRC-aligned regimes will diverge on acceptable PQ algorithms, key locations, HSM residency, regulator access, and lawful-access hooks. Cross-border AI and LEO systems will end up running multiple PQC and key-governance stacks in parallel.[10][11]
- The architecture and procurement question becomes: where can heterogeneity be tolerated (for example, per-region HSMs and PKIs), and where must a single PQC posture be enforced for systemic reasons (for example, cross-region AI control planes, shared backbones)? Those decisions should show up explicitly in RFPs, contracts, and integration patterns, not be left to vendor defaults.[15][18]

8. Conclusion: Treating cryptography as infrastructure in an AI–LEO world

Stepping back from the layers and timelines, the basic picture is not complicated. AI clusters and LEO/ground-segment nodes are turning into the coordination fabric for critical infrastructure. They concentrate long-lived data, high-leverage control planes, and hard-to-change update paths — the exact combination that makes "harvest-now, decrypt-later" a practical concern rather than a thought experiment.[1][9] At the same time, they are the only places in most architectures where there is enough observability, control, and operational discipline to run the kind of PQC migration NIST and NCCoE are describing.[6][13] In all environments, that migration starts with cryptographic visibility — discovery, inventory, and risk assessment — that must begin immediately using ACDI tools even in legacy environments, scaled up and accelerated by AI, and output of which is then analyzed by an AI layer.

The implication is that PQC cannot be treated as a later add-on to AI programs or space/OT modernization. It has to be wired in as a property of the infrastructure itself: how protocols are negotiated, how HSMs and accelerators are selected, how PKI is structured, how firmware is signed, how telemetry is shaped, and how RFPs are written.[1][6][15] If those decisions are made without a quantum lens now, the result will be a generation of AI and LEO systems that are operationally elegant but cryptographically brittle—difficult to defend today and cheap to break once credible quantum capabilities exist.[1][5]

The good news is that most of the ingredients are already on the table. ACDI tools exist, the PQC standards are stabilizing; NIST has provided a transition stack and mappings into CSF and SP 800-53; NCCoE has published a workable model for crypto-agility and migration; procurement guidance exists in the form of buyer's guides and reference language.[1][6][13][15] The missing piece is alignment: treating AI and LEO build-outs, PQC migration, and sector-specific governance as a single, coupled program rather than three separate conversations.

If there is a practical takeaway, it is this: between now and the early 2030s, the most important PQC decisions will not be made in cryptography working groups. They will be made by AI platform owners, space and OT program offices, cloud and telecom providers, and the regulators and acquisition teams that set their constraints. Wherever those actors treat cryptography as a first-order design and lifecycle parameter, there is a path to a sane quantum posture. Wherever they do not, the window for fixing it later is much narrower than it looks on paper.[5][18]

Bibliography

References

1. National Institute of Standards and Technology. (2024). *Transition to Post-Quantum Cryptography Standards* (NIST IR 8547). U.S. Department of Commerce. https://csrc.nist.gov/pubs/ir/8547/ipd
2. White House. (2024). *National Security Memorandum on Critical Infrastructure Security and Resilience (NSM-22)*. https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2024/04/30/national-security-memorandum-on-critical-infrastructure-security-and-resilience/
3. CISA, NSA, & NIST. (2023). *Quantum-Readiness: Migration to Post-Quantum Cryptography*. Cybersecurity and Infrastructure Security Agency. https://www.cisa.gov/resources-tools/resources/quantum-readiness-migration-post-quantum-cryptography
4. National Institute of Standards and Technology. (2024). *FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard*. U.S. Department of Commerce. https://csrc.nist.gov/pubs/fips/203/final
5. Kampanakis, P., et al. (2025). *Post-Quantum Hybrid Key Exchange in TLS 1.3* (IETF Internet-Draft). Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-ietf-tls-hybrid-design/
6. National Institute of Standards and Technology. (2025). *Mappings of Migration to PQC Project Capabilities to Risk Management Frameworks* (NIST CSWP 48). https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.48.ipd.pdf
7. National Cybersecurity Center of Excellence. (2024). *Migration to Post-Quantum Cryptography* (NIST SP 1800-38). https://www.nccoe.nist.gov/crypto-agility-considerations-migrating-post-quantum-cryptographic-algorithms
8. Mandiant/Google Cloud. (2025). *APT Trends: The Shift to Network Edge and Virtualization Platforms*. https://www.mandiant.com/resources/blog/shattered-network-security-edge-devices
9. Industrial Cyber. (2023). *CISA, NSA, NIST factsheet addresses migration to post-quantum cryptography*. https://industrialcyber.co/cisa/cisa-nsa-nist-factsheet-addresses-migration-to-post-quantum-cryptography-ahead-of-standards-rollout/
10. Encryption Consulting. (2025). *Understanding NIST CSWP 48 for PQC Migration and Security*. https://www.encryptionconsulting.com/understanding-nist-cswp-48-for-pqc/
11. Decent Cybersecurity. (2025). *NIST Unveils Comprehensive Roadmap for Post-Quantum Cryptography Transition*. https://decentcybersecurity.eu/nist-unveils-comprehensive-roadmap-for-post-quantum-cryptography-transition/
12. NCCoE. (2024). *Crypto-Agility Considerations: Migrating to Post-Quantum Cryptographic Algorithms*. https://www.nccoe.nist.gov/crypto-agility-considerations-migrating-post-quantum-cryptographic-algorithms
13. NIST. (2025). *Frequently Asked Questions about Post-Quantum Cryptography*. https://pages.nist.gov/nccoe-migration-post-quantum-cryptography/
14. NIST CSRC. (2025). *PQC Migration Mappings to Risk Framework Documents*. https://csrc.nist.gov/News/2025/pqc-migration-mappings-to-risk-framework-documents
15. GSA. (2025). *Post Quantum Cryptography Buyer's Guide*. General Services Administration. https://buy.gsa.gov/docviewer?id=60022

16. GSA. (2025). *Post Quantum Cryptography Buyer Overview*. https://buy.gsa.gov/docviewer?id=60022&docTitle=Post+Quantum+Cryptography+Buyer

17. Quantum.gov. (2024). *NIST Draft Report on PQC Transition*. https://www.quantum.gov/nist-draft-report-on-pqc-transition/

18. GSA. (2025). *Post Quantum Cryptography Roadmap*. https://buy.gsa.gov/docviewer?id=60023

Appendix A: Glossary of Terms

Core concepts & threats

- **Harvest-now, decrypt-later (HNDL)**: A strategic threat model where adversaries intercept and archive encrypted traffic today (while it is still secure) with the intent of decrypting it in the future once a cryptographically relevant quantum computer becomes available.
- **Post-quantum cryptography (PQC)**: Cryptographic algorithms (usually running on classical hardware) that are thought to be secure against an attack by a quantum computer. Unlike Quantum Key Distribution (QKD), PQC relies on complex mathematical problems (e.g., lattice-based) rather than quantum physics.
- **Cryptographically relevant quantum computer (CRQC)**: A quantum computer large and stable enough to run Shor's algorithm effectively, thereby breaking current public-key cryptography (RSA, ECC).
- **Crypto-agility**: The capacity of a security system to switch between cryptographic algorithms or parameters (e.g., moving from RSA to ML-KEM) via configuration or policy updates without requiring significant infrastructure re-engineering or downtime.
- **Hardware "Valley of Death"**: The transitional period (approx. 2025–2028) where PQC standards are finalized, but commercial off-the-shelf hardware (accelerators, SmartNICs) has not yet integrated native silicon support for them, forcing reliance on slower software implementations.

Infrastructure & hardware

- **AI cluster**: A high-performance computing environment optimized for artificial intelligence workloads (training and inference), characterized by massive parallelism, high-bandwidth interconnects, and heavy reliance on GPU/TPU accelerators.
- **LEO ground segment**: The terrestrial infrastructure required to control Low Earth Orbit satellite constellations. This includes tracking stations, antennas, and the backhaul networks that connect satellites to the terrestrial internet.
- **OT (operational technology)**: Hardware and software that detects or causes a change through the direct monitoring and/or control of physical devices, processes, and events (e.g., SCADA systems, industrial control systems).
- **HSM (hardware security module)**: A physical computing device that safeguards and manages digital keys, performs encryption and decryption functions for digital signatures, and provides strong authentication. PQC migration requires HSMs to support lattice-based keys.
- **SmartNIC (smart network interface card)**: A programmable network adapter that offloads processing tasks (encryption, firewalls, routing) from the host CPU. Legacy SmartNICs with fixed-function crypto blocks are a major PQC bottleneck.
- **BMC (baseboard management controller)**: A specialized service processor that monitors the physical state of a computer, network server, or other hardware device. BMCs often have long-lived, hard-to-patch firmware signing keys.

Cryptographic protocols & standards

- **ML-KEM (module-lattice-based key-encapsulation mechanism)**: The NIST standard (FIPS 203) for post-quantum key establishment. It replaces mechanisms like Diffie-Hellman and RSA key transport.
- **ML-DSA (module-lattice-based digital signature algorithm)**: The NIST standard (FIPS 204) for post-quantum digital signatures. It replaces RSA and ECDSA signatures.
- **Hybrid (composite) mode**: A transition strategy where two algorithms (one classical, one post-quantum) are used simultaneously. The data is only compromised if both algorithms are broken. This provides defense-in-depth during the transition years.
- **KEM (key encapsulation mechanism)**: A class of encryption techniques used to secure a symmetric key for transmission. In PQC, KEMs are the primary replacement for classical key exchange.
- **TLS 1.3 (Transport Layer Security)**: The current standard for securing communications over a computer network. PQC migration involves updating TLS handshakes to support hybrid or PQC-only key exchanges.
- **IKEv2 (Internet Key Exchange version 2)**: The protocol used to set up a security association (SA) in the IPsec protocol suite. Critical for securing VPNs and ground-segment backhaul.

Governance & Operations

- **CBOM (cryptographic bill of materials)**: A structured inventory (often machine-readable) that lists all cryptographic assets, algorithms, libraries, and dependencies within a piece of software or hardware. Essential for identifying quantum vulnerability.
- **NIST IR 8547**: A NIST internal report outlining the timeline and technical benchmarks for the migration to PQC standards.
- **CSWP 48 (Cybersecurity White Paper 48)**: A NIST document mapping PQC migration capabilities to existing risk management frameworks (like the NIST Cybersecurity Framework 2.0).
- **FIPS (Federal Information Processing Standards)**: Publicly announced standards developed by NIST for use in computer systems by non-military American government agencies and contractors. FIPS 203, 204, and 205 are the PQC standards.
- **CALEA (Communications Assistance for Law Enforcement Act)**: A US act requiring telecommunications carriers to ensure their equipment, facilities, and services are able to enable the government to intercept communications pursuant to a lawful authorization. PQC implementations must balance security with these lawful access requirements.

Dr. David Mussington

Fellow, Institute for Critical Infrastructure Technology (ICIT),
Co-Chair, ICIT's Center for FCEB Resilience,
Professor of the Practice at the University of Maryland's School of Public Policy

Dr. David Mussington is a Fellow of the Institute for Critical Infrastructure Technology (ICIT) and Co-Chair of ICIT's Center for FCEB Resilience. Additionally, he is a Professor of the Practice at the University of Maryland's School of Public Policy.

Prior to rejoining UMD in January of 2025, David served as the Executive Assistant Director for Infrastructure at the Cybersecurity and Infrastructure Agency, U.S. Department of Homeland Security. At CISA, David was one of three presidentially appointed officials charged with implementing the nation's critical infrastructure security and resilience strategies and plans across 16 critical infrastructures.

He also led interagency efforts on counter- and anti- terrorism efforts, playing a leading role in reducing the risks of domestic targeted violence, school safety, and physical infrastructure security standards. He was also a founding member of CISA's Cyber Safety Review Board.

David has extensive public and private sector experience in cyber and infrastructure security, selected for the Senior Executive Service and assigned to the Office of the Secretary of Defense in the role of Senior Advisor for Cyber Policy, later joining the NSC staff as Director for Surface Transportation Security Policy.

As a researcher at RAND Corporation and later at the Institute for Defense Analyses, David directed cybersecurity studies for the Department of Homeland Security (DHS), the Office of the Director of National Intelligence (ODNI), the Federal Communications Commission, the Bank of Canada, and NATO.

David has a Ph.D. in Political Science from Canada's Carleton University, and M.A. and B.A. degrees from the University of Toronto. He undertook postdoctoral study at Harvard's Belfer Center and at the UK's International Institute for Strategic Studies. In 2021 David was elected a life member of the Council on Foreign Relations.

In 2023 David was awarded Homeland Security Today's Mission Award, for contributions to the U.S. Critical Infrastructure Security and Resilience mission. In 2024 he received the 2024 Impact Award from the Institute for Critical Infrastructure Technology (ICIT) for leadership in critical infrastructure policy and strategy.

ICIT

www.icitech.org