March 2026

ICIT

Entangled Migrations

PQC, QKD, and US–PRC Risk Postures for Critical Infrastructure

David Mussington, Ph.D., CISSP, DDN QTE
ICIT Fellow, Co-Chair, ICIT FCEB Resilience Center
Professor of the Practice, University of Maryland School of Public Policy

www.icitech.org

Table of Contents

Bottom Line Up Front 03

I. Introduction: From Convergent Attack Surfaces to Divergent Risk Postures 05

II. QKD as a Concurrent Challenge: Global Maturation on the PQC Migration Timeline 07

III. Structural Coupling: How PQC and QKD Fail Together 09

IV. What Is Actually Deployed: US and PRC QKD Programs in Empirical Context 15

V. Divergent Risk Postures: PRC Defense-in-Depth vs. US Single-Assumption Concentration 19

VI. Geopolitical Fragmentation and Alliance Implications 25

VII. Implications for US CI Risk Frameworks and Deterrence Policy 29

VIII. Findings and Conclusion: The Question Defense-in-Depth Requires 34

Bibliography 37

About 39

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

BLUF (Bottom Line Up Front)

The ICIT Quantum-Resilient Convergence paper [1] established that Post-Quantum Cryptography (PQC) migration and AI/Low Earth Orbit (LEO) infrastructure modernization are a single coupled program, and that the window to embed cryptographic resilience into that infrastructure closes in the early 2030s. This paper extends that analysis into a dimension the first paper identified but did not develop: the parallel emergence of Quantum Key Distribution (QKD) as a live infrastructure investment on the same 2026–2035 timeline, and the structurally divergent risk choices the United States and the People's Republic of China (PRC) are making across both PQC and QKD for their most consequential critical infrastructure (CI) links.

The United States applies defense-in-depth as a foundational principle in nuclear deterrence, missile defense, cyber architecture, and critical infrastructure protection. Quantum-era cryptographic resilience is the conspicuous exception. Current US policy concentrates all CI protection on a single class of mathematical assumption — post-quantum cryptography — implemented through one standards track, executed on a migration timeline already under institutional strain. The PRC has made a different choice: it is building a sovereign QKD infrastructure layer at continental scale while simultaneously pursuing a domestic PQC stack, purchasing defense-in-depth against the possibility that either approach alone proves insufficient.

This paper does not argue that the US should replicate PRC-scale QKD deployment. It asks whether concentrating all quantum-era cryptographic resilience on PQC — without any physics-based fallback — is an acceptable risk posture for the small number of Tier-1 CI links — defined in this paper as the critical infrastructure communication paths where confidentiality horizons are permanent or multi-decade, where compromise enables physical consequences or systemic financial disruption, and where the cost of cryptographic failure is not recoverable through patching or migration after the fact — specifically nuclear command segments, financial settlement backbones, and bulk power control networks. The answer may be yes. But the question must be asked with the same rigor the United States applies to every other domain where it stakes national security on layered defense, and current policy does not ask it.

Thesis

PQC and QKD are conventionally framed as competing responses to the quantum cryptographic threat — mathematical hardness versus physics-based key exchange. That framing obscures two facts that matter for CI resilience.

First, the technologies are structurally coupled. QKD's operational security depends on PQC at the authentication layer; both share partial-deployment downgrade vulnerabilities; both face hardware maturity constraints on overlapping timelines; and both introduce classical chokepoints that reintroduce the threat surface they were designed to escape. Neither is a standalone answer. A risk framework that evaluates one without the other will systematically misestimate quantum-era exposure.

Second, they encode different bets against different failure modes — and the United States and the PRC have chosen different sides of that bet for their highest-value infrastructure. PQC is a bet on the durability of mathematical hardness assumptions. QKD is a bet on the stability of physical laws governing quantum measurement. The PRC has paid for both bets simultaneously. The US has chosen one. The divergence is not a deployment-scale gap measurable in kilometers of fiber. It is a structural difference in how each state allocates risk across algorithmic failure, infrastructure chokepoints, and governance for the CI links where the cost of being wrong is highest.

This paper's contribution is to extend the ICIT convergence analysis into that risk-posture dimension: modeling PQC and QKD as coupled surfaces with shared structural challenges, mapping the US and PRC choices as distinct risk allocations over those surfaces, and assessing whether the resulting US single-assumption posture warrants the same defense-in-depth scrutiny applied to every other domain of national security.

I. Introduction: From Convergent Attack Surfaces to Divergent Risk Postures

In February 2026, ICIT published *Quantum-Resilient Convergence: The Shared Defense of AI, Space, and Critical Infrastructure* [1], which argued that AI data centers, LEO ground-segment infrastructure, and post-quantum cryptographic migration are a single coupled program sharing capital expenditure cycles, the 2030–2035 planning horizon, and the same operational requirement for deep cryptographic visibility and control. That paper established that PQC migration is not a niche IT upgrade but a Tier-1 resilience requirement — and that the window to embed cryptographic resilience into the current wave of AI and LEO build-outs closes in the early 2030s. If PQC is treated as a bolt-on after infrastructure is built, the cost and complexity of remediation will be prohibitive.

This paper is a direct successor. It assumes the ICIT convergence analysis as baseline — the convergent attack surface, the Harvest-Now-Decrypt-Later (HNDL) threat environment, the hardware Valley of Death, the firmware and management-plane exposure, the governance tensions across jurisdictions [1] — and extends it into a dimension that the first paper identified but did not develop: the parallel emergence of QKD as an operational technology on the same timeline, and the structurally divergent choices the United States and the PRC are making across both PQC and QKD for their most consequential critical infrastructure links.

The policy conversation still treats PQC and QKD as competing paradigms. The National Security Agency (NSA), NIST, and the UK's National Cyber Security Centre occupy the QKD-skeptic position [12][36]: QKD provides no authentication, requires dedicated hardware, does not integrate with software-defined networks, and depends on trusted nodes for distances beyond metropolitan scale — making PQC dramatically more cost-effective and scalable for general-purpose deployment. The PRC and EU occupy the complementary-deployment position: QKD offers eavesdropping detection that PQC cannot provide, creates a cryptographic layer whose security depends on physics rather than algorithmic assumptions, and provides a hedge against PQC hardness assumptions weakening over the 2030–2035 horizon. Both positions are structurally coherent for the use cases they prioritize.

What neither position fully addresses is the space between them — the interdependency cascade between PQC and QKD, and the risk-acceptance asymmetry between a PQC-only strategy and a PQC-plus-QKD strategy on the small number of CI links where failure would be catastrophic. QKD's operational security depends on PQC at the authentication layer. Both PQC and QKD share downgrade vulnerabilities during partial deployment. Both are constrained by hardware maturity gaps on overlapping timelines. Both introduce classical chokepoints that reintroduce the threat surface they were designed to escape. Against that backdrop, the PRC has chosen to pay for a physics-based hedge against PQC failure on its Tier-1 links. The United States has not. The question is whether that divergence reflects examined risk acceptance on both sides — or examined acceptance on one side and institutional inertia on the other.

This paper provides the analytical structure to answer that question. It proceeds in seven sections.

Section II establishes QKD as a concurrent challenge maturing on the PQC migration timeline — driven not by US investment but by PRC and EU programs and commercial adoption that is entering the US CI environment regardless of federal policy.

Section III identifies five structural coupling mechanisms through which PQC and QKD share failure modes, and examines how living-off-the-land adversary tradecraft — already documented in US critical infrastructure — exploits the transition conditions those mechanisms create.

Section IV provides the empirical deployment comparison: what the PRC has built, what the US has built, and what the deployment gap reveals about risk allocation rather than technological capability.

Section V structures the US–PRC comparison across four risk dimensions — algorithmic failure, infrastructure chokepoints, governance and abuse, and exploitation and HNDL — and documents how decentralized US QKD adoption is producing emergent dual-deployment complexity without the corresponding algorithmic hedge.

Section VI maps how the bilateral divergence compounds into a fragmented global landscape affecting multinational CI operators, allied interoperability, and the specific structural constraints that the Communications Assistance for Law Enforcement Act (CALEA) imposes on US telecommunications QKD deployment.

Section VII translates the analysis into implications for US CI risk frameworks — the gaps in current frameworks, the extension requirements, and the components of a Tier-1-specific defense-in-depth evaluation — and reframes the strategic consequences through the lens of persistent cyber competition rather than classical deterrence.

Section VIII presents the findings and returns to the central question: whether the United States has applied to quantum-era cryptographic resilience the same defense-in-depth scrutiny it applies to every other domain where it stakes national security on layered protection.

II. QKD as a Concurrent Challenge: Global Maturation on the PQC Migration Timeline

The ICIT convergence paper anchored its analysis to a specific planning horizon: the 2026–2035 window during which AI infrastructure, LEO ground-segment modernization, and PQC migration share capital expenditure cycles, operational assumptions, and — if the migration is deferred — a compounding quantum vulnerability that becomes prohibitively expensive to remediate after the window closes. That timeline was derived from NIST IR 8547's [2] transition assumptions, the National Cybersecurity Center of Excellence's (NCCoE) migration framework [30], Commercial National Security Algorithm Suite 2.0 (CNSA 2.0) compliance deadlines [4], and the hardware refresh cadence for the AI and space systems that now serve as critical infrastructure's coordination layer.

QKD is maturing on the same timeline — not because the United States has programmed a parallel QKD initiative, but because other states and commercial actors have. The result is that QKD is becoming an operational feature of the global CI environment within the same window that US policy treats as exclusively a PQC migration problem.

The PRC's investment is the most consequential. The China Quantum Communication Network (CN-QCN) now spans more than 10,000 kilometers of backbone, with 145 fiber nodes linking 20 metropolitan networks across 17 provinces and 80 cities. [5][6] Satellite QKD from the Micius program is operational [8], and a geostationary earth orbit (GEO) QKD capability is planned for the late 2020s. These are not research testbeds. QKD-protected links are integrated into Chinese banking, grid control, government communications, and urban transit systems. The PRC is simultaneously pursuing a domestic PQC stack on a timeline comparable to US NIST-track migration. By the early 2030s, PRC Tier-1 critical infrastructure is designed to operate behind both layers.

The European Union is on a parallel trajectory at sovereignty-infrastructure scale. The European Quantum Communication Infrastructure (EuroQCI) initiative [10], backed by all 27 member states and integrated into the IRIS² secure connectivity program, combines terrestrial fiber QKD with a satellite segment. The European Space Agency's (ESA) Eagle-1 satellite is targeted for demonstration in 2026, and EuroQCI aims for initial operational capability by 2027. The EU's January 2026 proposed amendment to NIS2 [11] includes the first explicit PQC requirement in directive text — and the European Commission's Quantum Strategy simultaneously endorses QKD as a complementary layer for protecting government institutions, data centers, hospitals, and energy grids. Europe is not choosing between PQC and QKD. It is deploying both under a digital sovereignty rationale.

The United States occupies a structurally different position. NSA has stated that it does not support QKD for national security systems (NSS). [12] NIST has published skeptical assessments of QKD's operational viability for general-purpose network security. The Department of Defense (DoD) Chief Information Officer's PQC directive [13] established that outdated cryptographic solutions must be replaced with NIST-approved PQC algorithms, and the directive's framing excludes QKD from the approved remediation path. US quantum

communications investment is concentrated in Department of Energy (DOE) national laboratory testbeds and the Defense Advanced Research Projects Agency's (DARPA) QuANET program — research-stage efforts with no CI deployment mandate and no procurement hook connecting them to the PQC migration timeline that Office of Management and Budget (OMB) M-23-02 [14] and Executive Order 14144 [15] are driving across federal agencies. The National Quantum Coordination Office under the Office of Science and Technology Policy coordinates quantum networking research across DOE, NSF, DARPA, and other agencies under the National Quantum Initiative [16], and the five renewed National Quantum Information Science Research Centers include quantum communications in their portfolios. Whether that coordination extends to CI-specific QKD risk assessment or integration with PQC migration planning is not documented in any published policy instrument. If interagency deliberation on QKD's role in CI resilience exists below the level of published policy, it has not produced visible outputs in the form of guidance, mandates, or procurement language — and CI operators making independent QKD decisions are proceeding without a federal framework to reference.

The US position is structurally coherent for general-purpose network security. PQC is software-deployable across existing infrastructure at dramatically lower cost than dedicated quantum fiber. It scales to enterprise, cloud, and operational technology (OT) environments where QKD's hardware requirements and distance constraints are prohibitive. The NSA/NIST skepticism is well-founded for the overwhelming majority of CI links by volume.

But the US position creates a condition that the ICIT convergence paper's framework makes visible: the United States faces a two-technology risk environment — one in which both PQC and QKD are shaping the CI postures of peer competitors, allies, and commercial ecosystems — while resourcing a one-technology response. US CI operators who interact with PRC or EU infrastructure will encounter QKD-protected links and QKD-dependent systems whether or not US policy acknowledges them. US financial institutions are already experimenting with QKD on their own initiative. US energy utilities are hosting DOE-funded QKD demonstrations on operational grid infrastructure. The technology is arriving in the US CI environment through commercial and allied channels regardless of the federal policy posture.

The remainder of this paper examines what that mismatch means for Tier-1 CI risk. The next section establishes that PQC and QKD are not independent choices with separable risk profiles — they are structurally coupled in ways that affect any operator deploying either or both. Section V then maps how the US and PRC allocate risk across that coupled surface differently, and what the resulting asymmetry implies for the CI links where the cost of failure is highest.

III. Structural Coupling: How PQC and QKD Fail Together

The policy debate frames PQC and QKD as alternative approaches — one mathematical, one physical — and treats the choice between them as a technology selection problem. That framing is analytically incomplete. PQC and QKD share failure modes, and those shared failure modes interact in ways that affect CI operators regardless of which technology they deploy. A risk framework that evaluates PQC in isolation will miss the ways QKD's emergence changes the risk landscape even for PQC-only operators. A framework that evaluates QKD as an independent hedge will miss the ways QKD inherits PQC's failure modes at the authentication layer.

This section identifies five structural coupling mechanisms and then examines how existing adversary tradecraft — specifically living-off-the-land (LOTL) techniques already documented in US critical infrastructure — exploits the conditions these mechanisms create during transition.

III.A. Authentication Dependency

QKD's quantum channel can, in principle, deliver information-theoretically secure key material. But every deployed QKD system also operates a classical channel that performs endpoint authentication, basis reconciliation, error correction, and privacy amplification. That classical channel must use conventional or post-quantum cryptography. Once a cryptographically relevant quantum computer (CRQC) exists [2] [4], a QKD deployment whose authentication layer has not been migrated to PQC is compromised at the classical channel — regardless of how secure the quantum channel is. An adversary with CRQC capability can impersonate endpoints and run parallel QKD sessions, performing a person-in-the-middle attack that fully defeats the supposed quantum security.

This creates a false confidence problem. A CI operator that has deployed QKD on a Tier-1 link may classify that circuit as quantum-secure while in practice inheriting the full PQC Valley of Death exposure at its authentication layer. From an HNDL perspective, the quantum channel protects key material, but the classical authentication channel leaks session metadata and authentication material that can be harvested today and exploited once CRQC capability matures.

The implication is structural: QKD deployment timelines are not independent of PQC deployment timelines. They are constrained by them. The PRC's CN-QCN, the EU's EuroQCI, and any future US QKD pilot all depend on successful PQC migration at their authentication layers. Even where QKD is considered as a hedge for Tier-1 links, it cannot be treated as a standalone answer. It is coupled to the very PQC migration that the ICIT convergence paper identified as a Tier-1 resilience requirement.

III.B. Partial-Deployment Downgrade

The ICIT convergence paper documented how partial PQC migration creates a distinct attack surface: dual-stack configurations allow silent fallback from post-quantum to classical key establishment on critical paths, and without end-to-end cryptographic telemetry, an operator cannot distinguish genuine PQC coverage from advertised-but-unenforced coverage. Management planes — the administrative and orchestration interfaces used to configure, monitor, and provision infrastructure — firmware signing chains, and inter-region links can remain entirely classical even after a vendor reports PQC support. [1][3]

QKD has an analogous partial-deployment problem with a compounding property. When a QKD link degrades — through photon loss, detector saturation, environmental interference, or operational misconfiguration — the system falls back to classical key establishment. That fallback is operationally necessary. It also carries a priority inversion that PQC downgrade does not: organizations do not deploy QKD on routine traffic. The existence of QKD on a given link fingerprints the traffic as high-value in the operator's own risk assessment. Adversary collection logic can therefore treat QKD fallback events as highest-priority harvest opportunities. The operator's deployment decision has identified the traffic as worth quantum-protecting; the fallback event signals that the protection has temporarily lapsed.

A CI operator running PQC and QKD in combination faces downgrade opportunities at both layers — the PQC layer, where classical fallback occurs under load or misconfiguration, and the QKD layer, where quantum link degradation forces reversion to classical key establishment. The two downgrade signals are correlatable: an adversary monitoring both layers simultaneously can identify when a high-value link has lost both its algorithmic and its physics-based protections.

III.C. Hardware Maturity Gaps

The ICIT convergence paper identified the PQC Valley of Death — the 2025–2028 period during which PQC standards are finalized but commercial hardware (AI accelerators, SmartNICs, VPN appliances, OT management devices) lacks native silicon support, forcing reliance on slower software implementations or crypto-agile wrappers that impose performance penalties.

QKD faces an analogous but less-discussed hardware maturity constraint. Fiber QKD is limited to approximately 100 kilometers per segment without signal regeneration. Backbone distances require trusted-node relay chains — the PRC's original Beijing-Shanghai backbone used 32 trusted nodes over 2,000 kilometers [8]; the subsequent China Quantum Communication Network expansion added 145 fiber backbone nodes across more than 10,000 additional kilometers, for a combined national backbone exceeding 12,000 kilometers. [5][6] True quantum repeaters — entanglement-based devices that would extend quantum key distribution without trusted intermediaries — remain pre-commercial. Laboratory demonstrations exist, but no field-deployable system operates at the scale, speed, or reliability required for CI-grade links. The 2030–2035 PQC migration planning horizon overlaps with the earliest optimistic estimates for operational quantum repeater availability.

Dense LEO constellations partially relax the scheduling and availability constraints that limit current satellite QKD. A constellation with thousands of satellites in overlapping orbital planes could provide near-continuous line-of-sight between any satellite and any ground station, dramatically increasing key generation availability compared to single-satellite systems like Micius. However, current satellite QKD architecture requires the satellite to generate separate keys with each ground endpoint, making each satellite a trusted node that possesses key material for both sides of the conversation during key exchange. A constellation of thousands of QKD-capable satellites is thousands of trusted nodes in orbit. The path to eliminating this dependency — satellite-based entanglement distribution, where the satellite distributes entangled photon pairs without ever possessing the key material — has been demonstrated in principle (Micius demonstrated entanglement distribution over 1,200 km in 2017 [37]) but does not operate at the rates, reliability, or scale required for CI-grade links. The constellation form factor does shift the trusted-node risk profile: orbital nodes are harder to physically access, harder to persistently compromise through LOTL techniques, and subject to a narrower, more auditable supply chain than terrestrial relay nodes. That shift is meaningful for the risk-posture comparison in Section V but does not eliminate the trusted-node dependency within the planning horizon.

Satellite QKD also faces operational constraints independent of the trusted-node problem. Current systems operate primarily at night to reduce background photon noise. Key generation rates per pass remain low relative to terrestrial fiber QKD. Weather disrupts the optical link. The PRC's planned GEO QKD satellite would eliminate the LEO scheduling constraint but does not resolve the key-rate limitation. Ground station infrastructure shares the LEO ground-segment attack surface that the ICIT convergence paper already identified as quantum-risk-concentrated.

The combined implication is that neither PQC nor QKD offers a complete solution within the planning horizon that both occupy. PQC cannot protect what it cannot reach — firmware layers, legacy OT, non-agile signing chains. QKD cannot scale beyond its hardware constraints — distance, key rate, trusted-node dependency. Risk frameworks that assume one technology compensates for the other's gaps will systematically undercount residual exposure during the period when both are most needed.

III.D. Concentrated-Node Vulnerability

The ICIT convergence paper identified PQC migration lag at identifiable infrastructure chokepoints: legacy Hardware Security Modules (HSMs) that cannot support post-quantum algorithms, classical-signed Public Key Infrastructure (PKI) root certificates, and firmware signing chains pinned to non-agile key hierarchies. These chokepoints share a common property — they are classical systems, at known locations, whose compromise yields disproportionate returns because they serve as trust anchors for large portions of the cryptographic estate.

QKD's trusted-node architecture replicates this pattern at the infrastructure layer. Every relay node in a trusted-node QKD backbone stores and processes key material in a classical system at a fixed physical location. Compromise of a single trusted node breaks key confidentiality for all traffic transiting that node.

CN-QCN's 145 backbone nodes represent 145 instances of this concentrated-vulnerability pattern — classical systems with key material at known, enumerable locations whose compromise is high-yield.

The ICIT convergence paper's analysis of CALEA [44] mediation points as concentrated-vulnerability targets applies directly. [1] QKD trusted nodes are functionally equivalent: classical systems handling concentrated cryptographic material at fixed locations whose compromise yields outsized intelligence or operational value. In both cases, the security of the larger system is bounded by the security of the weakest concentrated node — which is a classical system, subject to the full classical threat surface that quantum-era cryptography was deployed to escape.

This mechanism is analytically distinct from the authentication dependency described in Section III.A. Authentication dependency is a protocol-layer problem: the QKD session's classical control channel requires PQC. Concentrated-node vulnerability is an infrastructure-layer problem: the relay architecture places classical attack surface at enumerable physical locations regardless of what authentication protocol the sessions use.

III.E. Interaction Effects During Simultaneous Migration

The four mechanisms above are not independent. They compound during the transition period in ways that a mechanism-by-mechanism assessment will underestimate.

Authentication dependency means QKD deployment cannot outrun PQC migration — deploying QKD before the authentication layer is PQC-hardened creates false confidence rather than genuine protection. Partial-deployment downgrade means both PQC and QKD create adversary collection opportunities during their respective transitions, and the two downgrade surfaces are correlatable. Hardware maturity gaps in PQC (no native silicon for post-quantum algorithms) and QKD (no quantum repeaters for trusted-node-free operation) overlap on the same 2026–2035 timeline, preventing either technology from compensating for the other's constraints during the period when both are most needed. Concentrated-node vulnerability means that deploying QKD to hedge against PQC algorithmic failure reintroduces classical chokepoints — trusted nodes — that PQC migration was supposed to eliminate.

The net effect is that the transition period is the period of maximum compound exposure. It is also the period during which both the United States and the PRC are making the deployment and investment decisions that will determine their cryptographic postures for the following decade. The risk-posture comparison in Section V is grounded in this compound exposure: neither state's choices eliminate risk during transition; each state is choosing which combination of risks to carry.

III.F. Living-off-the-Land Exploitation of Transition Conditions

The five coupling mechanisms describe structural vulnerabilities that emerge during the PQC/QKD transition. This subsection addresses how those vulnerabilities are most likely to be operationally exploited — not through novel quantum attacks, but through adversary tradecraft already documented in US critical infrastructure.

LOTL techniques — in which adversaries use tools, protocols, and credentials already present in the target environment rather than introducing detectable foreign tooling — are the exploitation modality most advantaged by dual-migration conditions. The reason is structural: the transition period creates exactly the operational environment in which LOTL thrives.

During PQC migration, systems legitimately negotiate both classical and post-quantum cipher suites, and fallback to classical is an expected operational condition. An adversary with pre-positioned access does not need to force a cryptographic downgrade. They persist in an environment where downgrade occurs routinely and collect during the classical windows that operators have accepted as transitional noise. The adversary activity is indistinguishable from the environment's own behavior.

If QKD is simultaneously deployed on high-value links, the same logic applies with the priority-inversion amplifier: QKD fallback to classical key establishment is an expected event during quantum link degradation. The adversary waits for physics — weather, detector limits, fiber degradation — to create the classical window, then collects. No quantum-specific tooling required.

The management plane expansion associated with dual migration compounds the opportunity. PQC migration introduces crypto-agility mechanisms, Cryptographic Bill of Materials (CBOM) tooling, and cryptographic telemetry pipelines. [1][3] QKD deployment adds quantum channel monitoring, trusted-node management, and fallback configuration systems. These management planes run on classical infrastructure — the same infrastructure where LOTL operators already dwell. An adversary with existing management-plane access gains visibility into the transition itself: which links are migrating, which fallback configurations are active, where QKD is deployed and therefore which traffic the operator considers highest-value.

QKD trusted nodes are particularly exposed. They are classical systems — standard servers, HSMs, and network equipment — running standard operating systems with standard management interfaces. They are native LOTL habitat. A trusted node compromised through credential reuse, legitimate remote management protocols, or scheduled task manipulation gives the adversary access to transiting key material without any quantum-specific attack capability. The adversary does not break the quantum channel; they inhabit the classical node the quantum channel depends on.

Finally, the credential and certificate churn inherent in dual migration — PQC re-keying, PKI restructuring, HSM provisioning, QKD key management layer deployment — produces a volume of legitimate cryptographic operations that far exceeds steady-state norms. LOTL operators exploit high-churn environments because each additional certificate request, key rotation, or management session disappears into operational noise. The transition period that the coupling mechanisms identify as maximum compound exposure is simultaneously the period when LOTL exploitation is least detectable.

The compound LOTL surface is proportional to transition complexity. Deployments that manage both PQC and QKD simultaneously present a larger operational exploitation surface during migration than PQC-only deployments — a tradeoff that the PRC has evidently accepted and that the United States avoids by default rather than by explicit risk assessment. This observation does not argue for or against either posture. It identifies a cost of the PRC's defense-in-depth approach that partially offsets its hedge against algorithmic failure, and it identifies a benefit of the US single-technology approach that partially offsets its single-assumption concentration. Neither advantage is free.

IV. What Is Actually Deployed: US and PRC QKD Programs in Empirical Context

The risk-posture comparison in Section V requires an empirical foundation — not a technology survey, but an honest accounting of what each state has built, what it is building, and where QKD has crossed from laboratory demonstration into operational infrastructure or funded CI-adjacent experimentation. This section provides that accounting. It is deliberately factual in register. The analytical weight falls in Section V; this section supplies the evidence base.

IV.A. PRC: Continental-Scale Operational Infrastructure

The PRC's QKD deployment is the most extensive in the world by every quantitative measure, and it is the only national program that has moved definitively from research demonstration to operational CI integration.

CN-QCN spans more than 10,000 kilometers of backbone fiber, incorporating 145 trusted relay nodes and linking 20 metropolitan QKD networks across 17 provinces and 80 cities. [5] Combined with the original Beijing-Shanghai Backbone Network, total fiber mileage exceeds 12,000 kilometers. The backbone architecture uses two ring structures [5] — Beijing-Jinan-Hefei-Wuhan and Shanghai-Hangzhou-Hefei-Nanjing — providing ring-topology protection with an average inter-node distance of approximately 70 kilometers and average link attenuation of 18.6 dB. Of the 145 backbone nodes, 41 are designated as backbone access nodes with metropolitan network access capability; the remaining 104 serve as trusted relay nodes.

The satellite layer is operational and expanding. The Micius LEO satellite demonstrated intercontinental QKD in 2018 [8] and continues to operate. The Jinan-1 quantum microsatellite extends the space segment. In March 2025, a Chinese-led research team demonstrated quantum-secured communication over 12,900 kilometers between Beijing and Stellenbosch, South Africa [7][9] — the longest intercontinental quantum link to date. A GEO QKD satellite is planned for launch around 2027, which would provide continuous coverage of a fixed geographic footprint and eliminate the pass-scheduling constraint that limits LEO QKD utility.

The critical distinction is that CN-QCN is not a testbed. QKD-protected links are integrated into PRC critical infrastructure operations: banking and financial settlement, electrical grid control, government communications, and urban transit systems. China Telecom Quantum Group has deployed quantum-secure systems for commercial applications including fuel station tax monitoring across multiple provinces. Hefei Rail Transit is piloting QKD-secured data transmission for ticketing and passenger systems. The PRC treats QKD as sovereign operational infrastructure, not as a research program awaiting a deployment decision.

The PRC is simultaneously pursuing a domestic PQC stack — analogous to but independent of the US NIST portfolio — under its own regulatory framework. The result is that PRC Tier-1 CI is being designed to operate behind both PQC and QKD layers by the early 2030s. Total PRC government spending on quantum science and technology has reached approximately $15 billion [6][39], with quantum communications

receiving sustained investment through national science and technology megaproject funding and a recently established National Venture Guidance Fund that includes quantum technology as a priority sector.

IV.B. United States: Research Demonstrations and Sector-Led Experimentation

The US presents a structurally different picture: no national QKD deployment program, an explicit federal policy against QKD for national security systems, and a small number of research demonstrations and commercially driven experiments that are advancing without connection to the PQC migration timeline federal agencies are executing.

DOE national laboratory testbeds. The most advanced US QKD-for-CI work is the multi-year program led by Oak Ridge National Laboratory (ORNL) and Los Alamos National Laboratory (LANL) in partnership with EPB, a Chattanooga-based utility and telecommunications company. Beginning in 2019, the ORNL/LANL team placed QKD systems in operational electrical substations [38] connected by EPB's fiber-optic network, demonstrating trusted-node relay of quantum keys across city-scale distances using three distinct vendor systems. The program demonstrated QKD interoperability across vendor-diverse hardware — a practical prerequisite for any eventual grid-scale deployment. ORNL researchers have noted that QKD systems integrated into grid infrastructure could remain secure for decades, matching or exceeding the service life of physical grid components — a longevity argument directly relevant to OT environments where hardware refresh cycles are 20 to 30 years. The program is funded by DOE's Office of Cybersecurity, Energy Security, and Emergency Response, making the CI security framing explicit. In 2025, ORNL began testing laboratory-developed automatic polarization compensation technology on EPB's commercially available quantum network — the first time DOE-developed QKD equipment has operated on a commercial network.

EPB itself has become a significant node in the US quantum ecosystem. In 2022, it launched the first commercially available quantum network in the United States [20], built on its existing fiber infrastructure. In April 2025, EPB and IonQ announced a $22 million partnership to establish the EPB Quantum Center [19] [20], housing a trapped-ion quantum computer alongside the quantum network. An NVIDIA DGX system has been installed at the center through an ORNL agreement, creating a hybrid classical-quantum research resource. The University of Tennessee at Chattanooga hosts the first US university node on a commercial quantum network.

DOE quantum research centers. In November 2025, DOE announced $625 million to renew five National Quantum Information Science Research Centers [16] for up to five years. These include Q-NEXT (headquartered at SLAC, focused on distributed quantum entanglement), the Quantum Science Center (Oak Ridge, national security applications), and three others spanning superconducting systems, quantum materials, and quantum computing architectures. The Chicago Quantum Exchange operates an 80-mile quantum network testbed linking Argonne National Laboratory, Fermi National Accelerator Laboratory, and the University of Chicago. A separate 80-mile testbed operates in New York State through Brookhaven National Laboratory and Stony Brook University. DOE has also awarded $25 million specifically for regional-scale quantum internet testbeds. These programs serve dual open-science and national-security research missions, but none carries a CI deployment mandate or procurement connection.

DARPA. The QuANET program, launched in March 2024 [17][18], is explicitly working on integrating quantum systems into existing US communication infrastructure and network protocols. The program conducted its first phase test event in late 2025, using fielded fiber optics supporting both quantum and classical links. QuANET's framing — interoperability with existing infrastructure rather than parallel quantum networks — is structurally distinct from the PRC's approach and more aligned with how QKD might eventually enter US CI environments if policy permitted.

Financial sector. JPMorgan Chase has deployed a quantum-secured crypto-agile network (Q-CAN) connecting two data centers with QKD-secured VPNs over 100 Gbps fiber. [22][23] The bank's Global Chief Information Officer has stated publicly that JPMorgan is pursuing a dual PQC and QKD remediation strategy. JPMorgan participates in DOE's Q-NEXT center, with its quantum research leadership citing the need for QKD development applicable to financial data protection. HSBC has trialed QKD alongside PQC in a simulated €30 million foreign exchange trade [26] — described as the first quantum-safe protection of a trading terminal. Danske Bank completed a live QKD-protected data center transfer under the EU's OpenQKD initiative. The financial sector is the one US CI vertical where commercial institutions are funding QKD experimentation alongside and independent of national laboratory research.

Federal policy posture. NSA has stated that it does not anticipate certifying or approving any QKD products for national security use. [12][40] The DoD CIO's November 2025 directive [13] requires replacement of outdated cryptographic solutions with NIST-approved PQC algorithms and mandates that all PQC-related technologies be approved by the DoD CIO PQC Directorate before testing, evaluation, or deployment — a framing that effectively excludes QKD from the defense PQC migration pathway. NIST's public posture is skeptical of QKD for general-purpose network security, citing authentication dependency, hardware requirements, distance limitations, and trusted-node constraints — the same structural limitations this paper documents in Section III.

A precise characterization of the federal posture requires distinguishing between what the published policy record establishes and what may exist in channels not visible to open-source analysis. The published record is clear: no US policy instrument establishes a QKD deployment strategy, CI integration mandate, or procurement framework comparable to the PQC migration stack. NSA's position against QKD for national security systems, NIST's skeptical assessments, and the DoD CIO directive's exclusionary framing are documented and quotable. What the published record cannot resolve is whether interagency coordination on QKD's CI implications exists in pre-decisional, classified, or informal channels — through the National Quantum Coordination Office, the NSC cyber directorate, or the interagency process around quantum technology — without having produced publicly visible outputs. The analytical point this paper makes does not depend on the answer to that question. Whether or not interagency deliberation exists, the CI operators documented in this section — JPMorgan, EPB, DOE testbed partners, financial sector experimenters — are making QKD deployment decisions in the absence of a published federal framework. The decentralized adoption pattern described in Section V.A is proceeding in that vacuum, and its compound risk implications are accumulating regardless of what coordination may be occurring at the interagency level.

The asymmetry in federal guidance is stark. CISA's January 2026 product categories list [32] identifies technologies where PQC-capable products are widely available and directs agencies to procure only PQC-capable products in those categories. CISA's Cybersecurity Performance Goals 2.0, released December 2025, aligns with NIST CSF 2.0 and integrates governance functions for cryptographic risk management. NIST CSWP 48 [3] maps PQC migration capabilities into SP 800-53 controls. CISA published OT-specific post-quantum considerations in late 2024. The PQC side of the quantum-era transition has a dense, layered controls and oversight framework — standards, migration mappings, procurement guidance, product category lists, and sector-specific considerations. No equivalent instrument exists for QKD. Not restrictive guidance, not hostile guidance — no guidance at all beyond NSA's blanket exclusion for national security systems. The CI operators deploying QKD on their own initiative are doing so in a space where the federal apparatus has produced extensive controls for the technology they are migrating from and to, and nothing for the technology they are simultaneously experimenting with.

IV.C. The Deployment Gap as Analytical Input

The quantitative gap is stark. The PRC operates 12,000+ kilometers of QKD backbone with 145 nodes across 80 cities, integrated into operational CI. The US has 80-mile testbeds, one commercial municipal network, a handful of financial sector experiments, and an explicit federal policy against QKD for national security applications. Comparing the two as a "quantum race" measured in kilometers or nodes is tempting but analytically incomplete.

What matters for Section V's risk-posture comparison is not the scale difference per se, but what the scale difference reveals about risk allocation. The PRC has made a deliberate, funded, and executed decision to build physics-based cryptographic fallback infrastructure for its Tier-1 CI links, accepting the infrastructure chokepoint risk, governance risk, and LOTL-exploitable complexity that entails. The United States has made a deliberate decision — expressed through NSA policy, DoD directives, and NIST guidance — to concentrate quantum-era resilience on PQC, accepting single-assumption exposure while maintaining a simpler and more defensible operational posture during transition.

The question is not which decision is correct. The question — which Section V takes up — is whether each decision reflects an examined risk acceptance appropriate to the assets at stake, or whether one side's posture is the product of explicit strategic assessment and the other's is the product of institutional inertia.

V. Divergent Risk Postures: PRC Defense-in-Depth vs. US Single-Assumption Concentration

The deployment reality documented in Section IV is often described as a gap — one state has built more than the other. That vocabulary misidentifies the nature of the divergence. What the US and PRC have done is make structurally different bets about which failure modes are tolerable on their highest-value critical infrastructure links. Neither posture eliminates risk. Each accepts exposures the other avoids. The question that matters is not which state has built more, but which failure modes each state has chosen to live with on the links where failure is catastrophic — and whether those choices were made by design or by default.

This section structures the comparison across four risk dimensions that matter for Tier-1 CI — nuclear command segments, financial settlement backbones, and bulk power control networks. These are the links where confidentiality horizons are permanent or multi-decade, where compromise enables physical consequences or systemic financial disruption, and where the cost of algorithmic surprise is not recoverable through patching or migration after the fact.

V.A. Algorithmic Failure Risk

This is the dimension on which the two postures diverge most sharply — and the one where the consequences of the divergence extend furthest into the future.

The United States has concentrated its quantum-era cryptographic resilience on PQC — specifically, on the NIST-standardized algorithm families ML-KEM [27], ML-DSA [28], and SLH-DSA [29], with CNSA 2.0 [4] providing the compliance framework for national security systems. The security of this posture depends on the mathematical hardness assumptions underlying those algorithms — principally lattice-based problems — holding through the planning horizon and beyond. If those assumptions hold, the US approach is dramatically more cost-effective and scalable than any alternative. PQC integrates with established protocol stacks and risk frameworks and does not require purpose-built parallel physical infrastructure — though, as the ICIT convergence paper documented [1], it does require hardware upgrades and replacements where existing devices cannot support post-quantum key sizes and algorithm implementations, a constraint the Valley of Death analysis identified as a binding schedule risk through the late 2020s. [2][33]

If those assumptions weaken, the consequences depend on the character of the weakening. The history of cryptanalysis suggests that partial, family-specific degradation — a reduction in effective security margin, an advance that affects one algorithm family but not others — is more common than catastrophic overnight failure across all deployed primitives. NIST's own decision to select HQC as a backup key-encapsulation mechanism from a different mathematical family [2] and to launch additional signature competitions reflects institutional recognition that single-family dependence is a risk.

A partial weakening is the scenario where the US single-assumption posture is most consequentially exposed. In a partial-weakening scenario, the algorithms are not definitively "broken" — they are weakened enough

to create ambiguity about whether the security margin remains adequate for Tier-1 confidentiality horizons measured in decades. That ambiguity produces a policy paralysis problem: the cost of premature migration (disruption, interoperability loss, resource diversion) must be weighed against the cost of delayed response (continued exposure on links whose protection is degrading). For links with permanent confidentiality requirements — nuclear command data, certain financial settlement records — any period of ambiguous protection is a period during which harvested traffic may be retrospectively exposed.

The PRC has made a different bet. By building QKD infrastructure alongside a domestic PQC stack, the PRC has purchased defense-in-depth against algorithmic failure. If the PRC's PQC algorithms prove vulnerable, Tier-1 links operating behind QKD retain a physics-based key-establishment channel whose security does not depend on any computational hardness assumption. The key material generated over a properly functioning QKD channel is information-theoretically secure — it does not degrade as adversary computational resources increase. A CRQC that weakens lattice-based PQC does not affect the security of a QKD-generated key, because the QKD key was never protected by a lattice problem. It was generated through a physical process whose security is governed by quantum mechanics, not computational complexity.

In a partial-weakening scenario, the PRC's Tier-1 links maintain a physics-based confidentiality floor while the algorithmic layer is assessed, patched, or replaced. The PRC operator loses defense-in-depth but retains one intact layer. The US operator, on the same timeline, faces unmitigated exposure on every Tier-1 link while debating whether the weakening is severe enough to justify disruptive migration. The asymmetry is sharpest not in the dramatic scenario of total PQC failure — where both states would face crisis conditions regardless — but in the ambiguous middle ground where algorithmic confidence erodes gradually and the cost of action must be weighed against the cost of inaction.

This utility of QKD as a PQC-failure hedge is real but conditional. Section III.A established that QKD's classical authentication channel must itself be PQC-hardened; otherwise, the quantum channel is subject to person-in-the-middle attacks that defeat its security regardless of the physics. QKD's fallback value is highest when the PQC weakening is family-specific — affecting, say, lattice-based key encapsulation but not code-based or hash-based signature schemes — so that the authentication layer can be maintained on an unaffected algorithm while the key-establishment layer relies on QKD. It is lowest in the catastrophic scenario where all deployed PQC families fail simultaneously and the authentication layer falls with them. Given that partial, family-specific weakening is the historically more probable pattern, and that a well-engineered QKD deployment can diversify its authentication across multiple PQC families, the hedge has real expected value even after the authentication dependency is discounted.

The decentralized deployment problem. The algorithmic-failure risk dimension is further complicated, in the US context, by the fragmented infrastructure ecosystem through which Tier-1 traffic actually moves.

A financial institution like JPMorgan Chase that deploys QKD between its own data centers is securing a specific point-to-point segment. But Tier-1 settlement traffic does not stay within a single operator's infrastructure. It transits hyperscaler cloud environments — AWS, Azure, Google Cloud — that control the physical backbone, service mesh fabric, HSMs, and protocol stacks through which that traffic moves

between regions. It may cross LEO satellite backhaul segments operated by constellation providers whose cryptographic architecture decisions are driven by their own cost, performance, and engineering constraints. It terminates at clearing houses and counterparties whose migration timelines are independent of the originator's.

The major US hyperscalers are investing heavily in PQC but have not announced QKD integration into their backbone or customer-facing infrastructure. Their architecture — software-defined, globally distributed, multi-tenant, dynamically routed — is structurally mismatched with QKD's requirements for dedicated fiber, point-to-point topology, and single-tenant key management. LEO backhaul operators have not yet deployed either PQC or QKD at scale on their satellite-to-ground links. The result is that a CI operator's QKD-secured segment exists as a cryptographic island embedded in a PQC-only or, during transition, still-classical ocean. The end-to-end confidentiality of the transaction is bounded by the weakest segment in the path, and the weakest segments are controlled by organizations the CI operator does not govern.

This fragmentation produces a condition that shares more with the PRC's compound-deployment complexity than the clean single-assumption model the US federal posture implies. Decentralized QKD adoption by individual CI operators, hyperscaler PQC-only deployment on backbone infrastructure, LEO backhaul segments on independent migration timelines, and commercial quantum network experiments like EPB's Chattanooga deployment are producing an emergent patchwork in which some Tier-1 traffic segments carry QKD protection, others carry PQC, and others remain classical during transition. The seams between these segments — where QKD-protected traffic enters a PQC-only backbone, or where PQC-protected traffic crosses a still-classical backhaul hop — reproduce many of the partial-deployment failure modes Section III documented: downgrade surfaces at organizational boundaries, management plane complexity spanning multiple operators' orchestration environments, priority inversion where QKD presence on some segments fingerprints high-value traffic on adjacent unprotected segments, and LOTL-exploitable classical infrastructure in the transit path that no single operator monitors end to end.

The PRC faces analogous compound conditions — but as the product of a centrally planned architecture whose risks are visible, in principle, from a single governance vantage point. The US compound exposure is emergent. It arises from independent decisions by independent actors and is visible to no single entity in its totality. No one holds the end-to-end view of the cryptographic posture of a Tier-1 transaction that crosses organizational boundaries between a CI operator, a hyperscaler, a backhaul provider, and a counterparty.

The consequence is that the United States is not maintaining the clean single-assumption posture that federal policy describes. It is drifting into a partial version of the transition complexity costs that a deliberate dual PQC+QKD deployment would entail — the same categories of seams, downgrades, and management plane expansion — without the corresponding algorithmic-failure hedge that a deliberate dual deployment would provide, and without the centralized visibility or governance coordination to manage the compound risks it is accumulating. The US position is acquiring the costs of both postures while securing the benefits of neither: single-assumption concentration on the links that matter most, with emergent dual-deployment complexity on the links where market actors have made independent decisions.

This is not an argument that the US should replicate the PRC's centralized model. It is an observation that the US is not, in practice, operating the PQC-only posture its policy describes — and that the gap between policy and operational reality is itself a risk that current frameworks do not assess.

V.B. Infrastructure Chokepoint Risk

On this dimension, the PRC's defense-in-depth advantage reverses.

CN-QCN's 145 trusted backbone nodes are classical systems at fixed, known physical locations, each storing or processing key material for traffic transiting the node. Section III.D established that these nodes replicate the concentrated-vulnerability pattern the ICIT convergence paper documented for legacy HSMs and CALEA mediation points: compromise of a single node yields disproportionate returns. The more nodes in the relay chain, the more enumerable chokepoints exist, and each is a classical system subject to the full classical threat surface.

The United States, by not deploying QKD backbone infrastructure, does not carry this class of risk. US PQC migration introduces its own chokepoints — the legacy HSMs, non-agile PKI roots, and firmware signing chains the ICIT paper identified — but those are chokepoints of transition, not chokepoints of permanent architecture. Once PQC migration is complete, those chokepoints are retired. QKD trusted nodes, by contrast, are permanent features of the operational architecture for as long as quantum repeaters remain pre-commercial — which, as Section III.C established, overlaps with the entire planning horizon.

The PRC mitigates this risk through physical security (substations, government buildings, military facilities), centralized operational authority, and a governance model in which the state controls every node in the chain. Those mitigations are real but not absolute. The LOTL analysis in Section III.F applies: trusted nodes are classical systems with standard management interfaces, and the PRC's own CI networks are not immune to sophisticated intrusion. The PRC has accepted 145 permanent classical chokepoints as the infrastructure cost of its algorithmic hedge. The US carries zero QKD-specific chokepoints but has no algorithmic hedge to show for it.

V.C. Governance and Abuse Risk

The governance dimension is where the PRC's QKD architecture reveals a structural property that is simultaneously a national advantage and a systemic vulnerability — depending on who is assessing it.

QKD trusted-node architectures are structurally compatible with state access to transit traffic. A government that controls every trusted node in a QKD backbone possesses — by architectural necessity — the key material for all traffic transiting those nodes. In a closed national system where the state is the operator, this property is not a bug; it is a feature that aligns QKD infrastructure with domestic surveillance and lawful-intercept requirements. The same architectural property that makes CALEA problematic for US telecom QKD deployment makes trusted-node QKD convenient for PRC governance, where state access to communications is a design requirement rather than a tension to be managed.

This means the PRC's QKD investment serves dual purposes: it provides a physics-based cryptographic hedge for Tier-1 CI, and it provides a state-controlled communications infrastructure whose architecture inherently supports centralized key access. Those purposes are complementary within the PRC governance model. They would be contradictory within the US governance model, where the Fourth Amendment, CALEA's court-authorization requirement, and the structural separation between CI operators and intelligence authorities create tensions that a trusted-node QKD architecture would amplify rather than resolve.

For US CI operators, this governance dimension is relevant in two ways. First, it partially explains why the PRC can invest in QKD at continental scale: the infrastructure serves governance objectives beyond cryptographic resilience, which broadens the political and budgetary constituency for the investment. Second, it means that any US or allied QKD deployment would need to be architected differently — with governance constraints, access controls, and oversight mechanisms that a PRC-style centralized trusted-node model does not require. The US is not simply choosing not to build QKD; it faces a harder architectural problem if it ever chooses to build it.

The US PQC-only posture avoids this governance complexity entirely. PQC does not require centralized key-handling infrastructure. It is implemented in distributed protocol stacks, libraries, and HSMs that operate within existing governance and oversight frameworks. The governance simplicity of PQC-only is a genuine advantage — but it is an advantage purchased at the cost of the single-assumption exposure described in Section V.A.

V.D. Exploitation and HNDL Risk

On the exploitation dimension, the two postures produce different exposures across different time horizons.

In the near term — the transition period through the early 2030s — the PRC's dual deployment creates a larger exploitation surface. Section III.F established that simultaneous PQC and QKD migration produces compound LOTL opportunities: expanded management planes, dual-stack fallback conditions, credential churn, and trusted-node classical infrastructure that is native LOTL habitat. The PRC accepts this transitional exploitation exposure as the cost of building defense-in-depth. The US, by running a single PQC migration, carries a narrower transitional exploitation surface.

In the longer term — post-transition, once deployments mature and stabilize — the exposure profiles reverse. A PRC Tier-1 link operating behind both PQC and QKD layers is less vulnerable to retrospective HNDL exploitation than a US link operating behind PQC alone. If an adversary has harvested encrypted traffic from both states' Tier-1 links during the transition period, the PRC link retains a physics-based layer whose security does not degrade with advances in cryptanalysis or quantum computation. The US link's long-term security depends entirely on whether the PQC algorithms protecting the harvested traffic withstand whatever computational capabilities the adversary develops over the confidentiality horizon — which, for nuclear command and financial settlement data, may be measured in decades.

The HNDL dimension also interacts with the authentication dependency documented in Section III.A. Both states' QKD deployments (where they exist) are vulnerable to HNDL at the classical authentication layer until

PQC migration is complete. But once PQC migration succeeds at the authentication layer, the PRC's QKD links gain a property the US estate does not possess: key material that was exchanged over the quantum channel is information-theoretically secure against future computational advances, including advances that might weaken the PQC algorithms protecting the rest of the session. The US PQC-only estate has no equivalent long-horizon property. Its confidentiality guarantee is bounded by the lifespan of the algorithmic assumptions, not by physics.

V.E. The Asymmetry That Matters

Across the four dimensions, a pattern emerges that is more nuanced than "China is ahead" or "the US approach is more efficient."

The PRC has purchased defense-in-depth against algorithmic failure at the cost of permanent infrastructure chokepoints, governance entanglement, and elevated transitional exploitation exposure. Each of these costs is real, documented, and — the evidence suggests — deliberately accepted within a governance model that can absorb them.

The United States has purchased operational simplicity and a narrower transitional attack surface at the cost of single-assumption concentration on its highest-value CI links. That cost is also real, but it is less clear that it has been deliberately accepted through explicit risk assessment rather than inherited through institutional inertia and policy framing that treats QKD as categorically impractical.

The distinction matters. The PRC's posture has the properties of an examined risk acceptance: the investment is deliberate, the costs are structural, and the tradeoffs are visible in the architecture. The US posture has the properties of an unexamined default: NSA/NIST guidance correctly identifies QKD's limitations for general-purpose deployment, and the policy apparatus has treated that general-purpose assessment as dispositive for all CI links — including the Tier-1 links where the risk calculus is structurally different from the general case.

The question this paper poses is not whether the US should build a CN-QCN equivalent. It is whether the Tier-1 CI links — nuclear command, financial settlement, bulk power control — warrant a separate risk assessment that considers defense-in-depth options the general-purpose policy forecloses. The PRC's investment — its scale, duration, operational integration, and state-directed consolidation — has the characteristics of a deliberate strategic choice, even if the assessment process behind it is not visible in open sources. The United States has not made an equivalent choice, and the absence of a published Tier-1-specific risk assessment suggests it has not conducted the evaluation that would inform one. That gap — between a posture that bears the marks of examined risk acceptance and one that bears the marks of institutional inertia — is the asymmetry that matters.

VI. Geopolitical Fragmentation and Alliance Implications

The US–PRC divergence described in Section V does not exist in a bilateral vacuum. It compounds into a fragmented global landscape in which CI operators, multinational financial institutions, allied military systems, and LEO constellation providers must navigate at least three distinct and partially incompatible quantum-era cryptographic regimes — US/NIST-aligned, PRC-domestic, and EU-sovereign — each with its own algorithm selections, key-governance requirements, QKD infrastructure postures, and lawful-access expectations. The fragmentation is not a side effect of the PQC/QKD divergence. It is a structural consequence of it, and it introduces risks that neither the US nor PRC posture accounts for independently.

VI.A. Three Regimes, Not Two

The US-aligned ecosystem is converging on the NIST PQC portfolio — ML-KEM (FIPS 203) [27], ML-DSA (FIPS 204) [28], SLH-DSA (FIPS 205) [29] — with NCCoE migration patterns [30] and the GSA PQC Buyer's Guide [31] providing the procurement and implementation framework. QKD is excluded from the approved remediation path for national security systems and absent from the federal compliance timeline. The operative assumption is PQC-only, implemented through crypto-agility mechanisms within existing infrastructure.

The PRC is building a parallel stack on both dimensions: domestic PQC algorithms developed under Chinese Academy of Sciences and CCSA standardization processes, and QKD infrastructure at continental scale under state direction. PRC-aligned systems — including QKD equipment marketed through Digital Silk Road channels by QuantumCTek and its state-affiliated successors — are positioned to operate behind algorithm families and key-governance models that are not interoperable with NIST-track implementations. A government that procures QKD infrastructure from PRC-aligned suppliers operates, architecturally, under PRC cryptographic standards and trusted-node governance whether the host country recognizes that dependency or not.

The EU occupies a third position that is allied to the US on PQC algorithm selection but divergent on QKD and sovereignty. EuroQCI is an explicit commitment to sovereign quantum communications infrastructure. The EU's proposed NIS2 amendment requires PQC migration — aligning with the NIST timeline — but the European Commission simultaneously endorses QKD as a complementary layer, and EU digital sovereignty directives impose key-residency, HSM auditability, and data-localization requirements that create operational constraints for US hyperscalers and LEO operators serving European CI customers. The EU is not choosing the US model or the PRC model. It is building a third path that selects PQC interoperability with the US and QKD infrastructure development closer to the PRC's approach, wrapped in a governance framework that neither Washington nor Beijing fully controls.

VI.B. PQC Algorithm Fragmentation

The most immediate operational consequence is algorithm fragmentation across jurisdictions. US-aligned systems will mandate NIST-track algorithms. PRC-aligned systems will mandate domestic alternatives. EU systems will likely accept NIST algorithms but may require additional algorithm support for sovereignty compliance or interoperability with EuroQCI.

For multinational CI operators — global banks, energy companies with cross-border grid interconnections, satellite operators serving multiple jurisdictions, cloud providers with data centers in all three blocs — this means running multiple PQC profiles in parallel. The ICIT convergence paper's analysis of crypto-agility as a migration requirement becomes a permanent architectural condition rather than a transitional one: operators must maintain the ability to negotiate different algorithm families on different paths depending on jurisdiction, and must do so without creating the downgrade and fallback vulnerabilities Section III documented.

Algorithm fragmentation also complicates the CBOM and cryptographic telemetry frameworks the ICIT paper identified as essential for PQC migration management. A CBOM that tracks NIST-algorithm coverage is insufficient for a system that must also demonstrate compliance with PRC or EU cryptographic requirements. The inventory, telemetry, and anomaly detection systems described in the ICIT paper's Section 5 must be extended to cover multi-regime cryptographic postures — a complexity increase that current frameworks do not address.

VI.C. QKD Infrastructure Fragmentation

The QKD dimension fragments differently. The PRC is building sovereign QKD infrastructure and exporting it. The EU is building sovereign QKD infrastructure but not exporting it at PRC scale. The US is building neither — but its CI operators are encountering both through commercial relationships, allied interconnections, and the decentralized adoption dynamics Section V.A described.

The result is that QKD-protected segments will appear in multinational CI service delivery chains without a common governance framework for how they interoperate, how their trusted-node architectures are audited across jurisdictions, or how downgrade behavior at the boundaries between QKD-protected and non-QKD segments is monitored. A European bank whose settlement traffic traverses an EuroQCI-protected segment within the EU, a NIST-PQC-protected segment across the Atlantic, and a CN-QCN-adjacent segment for Asia-Pacific operations faces three different cryptographic regimes, three different trust-anchor models, and three different governance expectations — with the seams between them as the points of maximum vulnerability.

The LEO dimension adds a further complication. Constellation operators (Starlink, Kuiper, OneWeb) serve customers across all three blocs from shared orbital infrastructure. If a LEO operator eventually deploys QKD on satellite-to-ground links, the question of which QKD standard, which trusted-node governance model, and which jurisdiction's key-management requirements apply to a satellite that serves US, EU, and PRC-adjacent ground stations simultaneously has no current answer. The ICIT convergence paper's

observation that LEO ground-segment systems will need to run multiple PQC profiles in parallel extends directly: they may also need to run multiple QKD governance models, or decline QKD entirely and accept the single-assumption posture as the cost of jurisdictional simplicity.

The fragmentation has a signaling dimension that compounds the technical risk. A government that procures QKD infrastructure from PRC-aligned suppliers is not merely selecting a technology — it is accepting PRC cryptographic standards, PRC trusted-node architecture, and, implicitly, PRC governance over the key material transiting that infrastructure. A CI operator that deploys on EuroQCI is operating within EU sovereignty and data-localization frameworks. An operator that declines QKD entirely and relies on NIST-track PQC is aligning, by default, with the US single-assumption posture. In a fragmented landscape, the choice of quantum-era cryptographic tooling — PQC algorithm family, QKD infrastructure provider, key-governance model — functions as a geopolitical alignment signal whether the operator intends it as one or not. For multinational CI operators and LEO constellation providers serving customers across all three blocs, there is no neutral selection. Every architectural decision positions the operator within a regime, and adversaries, allies, and regulators will read it accordingly.

VI.D. CALEA as a Structural Constraint on US QKD

CALEA [44] imposes a specific constraint on US telecommunications QKD deployment that has no equivalent in the PRC or EU contexts.

CALEA requires telecommunications carriers to ensure that their systems can accommodate court-authorized lawful interception. QKD's trusted-node architecture, as Sections III.D and V.C documented, is structurally compatible with centralized state access — the trusted node handles key material at a known location, and whoever controls the node can access the traffic. In the PRC governance model, this is a feature. In the US governance model, it creates a tension: a trusted-node QKD deployment on a US telecommunications backbone would place key material at mediation points that CALEA mandates be accessible to law enforcement under court order, but the security argument for QKD rests on the assumption that no intermediary accesses the key material.

The tension is not irresolvable in principle — endpoint-based lawful-access architectures, tightly governed intercept functions at the application layer, and cryptographic logging mechanisms could satisfy CALEA without compromising the transport layer's quantum resistance. But those architectures have not been developed, standardized, or tested for QKD-specific deployments. Until they are, CALEA functions as a structural blocker for QKD deployment on US telecommunications infrastructure — not because the law prohibits QKD, but because the architectural work required to make QKD CALEA-compatible has not been done.

This constraint does not apply to non-telecommunications CI operators (financial institutions, energy utilities) deploying QKD on owned infrastructure, which is why the JPMorgan and EPB deployments can proceed while telecommunications backbone QKD cannot. But it means that the US CI links most likely to

benefit from QKD as a physics-based hedge — long-haul telecommunications backbones carrying Tier-1 traffic — are the links where deployment faces the most significant governance obstacle.

VI.E. Alliance Interoperability

The fragmentation described above has direct implications for allied military and intelligence interoperability. NATO and Five Eyes operations depend on secure communications across national boundaries, and those communications currently ride on cryptographic infrastructure that each nation controls independently.

PQC interoperability among US-aligned allies is relatively tractable — convergence on NIST-track algorithms provides a common baseline, and NATO's own PQC assessment work is oriented toward NIST-family adoption. [35] But QKD interoperability is not addressed in any current allied framework. If the UK deploys QKD on military communications through its SPOQC program, France through its national quantum plan, and the US does not deploy QKD at all, allied communications will cross QKD/non-QKD boundaries with the same seam vulnerabilities Section III.B documented for partial deployment — and with no common telemetry, downgrade policy, or trusted-node governance to manage the transitions.

The EU's EuroQCI adds a further complication. EuroQCI is designed as sovereign European infrastructure, and its governance framework reflects EU digital-sovereignty priorities that may not align with Five Eyes intelligence-sharing requirements or NATO operational needs. Allied quantum communications infrastructure is being built on national and regional trajectories that are not coordinated with collective defense requirements — a fragmentation pattern that compounds the US–PRC divergence into a multilateral governance gap.

The communications interoperability dimension this paper documents is the visible layer of a deeper coordination challenge. Five Eyes partners are making independent PQC migration and QKD deployment decisions that affect not only secure communications interoperability but also coordinated intelligence capabilities whose operational details are appropriately outside the scope of unclassified analysis. The absence of a published allied framework for QKD coordination is visible at the interoperability level; whether equivalent coordination gaps exist at the intelligence cooperation level is a question for classified assessment. The structural conditions this paper identifies — uncoordinated QKD deployment among allies, divergent PQC migration timelines, no common downgrade telemetry or trusted-node governance — suggest that the same fragmentation dynamics documented here for CI interoperability apply to the intelligence coordination dimension as well. A classified supplement to this analysis, examining PQC/QKD transition effects on allied collection coordination and shared access maintenance, would complement the CI resilience findings presented here.

VII. Implications for US CI Risk Frameworks and Deterrence Policy

The preceding sections established three findings that current US risk frameworks do not adequately register. First, PQC and QKD are structurally coupled in ways that make independent risk assessment of either technology analytically incomplete. Second, the US and PRC have made divergent risk bets across that coupled surface — the PRC purchasing defense-in-depth against algorithmic failure at the cost of infrastructure chokepoints and governance complexity; the US purchasing operational simplicity and a narrower transition attack surface at the cost of single-assumption concentration on its highest-value links. Third, the US is not, in practice, operating the clean PQC-only posture its federal policy describes — decentralized adoption is producing emergent dual-deployment complexity without the corresponding algorithmic hedge or centralized visibility to manage the compound risks.

This section translates those findings into specific implications for US CI risk frameworks — what they currently miss, where they need extension, and what a rigorous defense-in-depth assessment for Tier-1 links would require.

VII.A. The Single-Point-of-Assumption Gap in Current Frameworks

The NIST Cybersecurity Framework 2.0, SP 800-53, and CSWP 48's [3] PQC migration mappings provide a structurally sound approach to quantum-era cryptographic risk for the general case. They drive cryptographic inventory, crypto-agility, telemetry, key management, and vendor engagement into existing governance and compliance cycles. For the overwhelming majority of CI links — enterprise, cloud, OT, and communications infrastructure where the risk is HNDL exposure on traffic with finite confidentiality horizons — PQC-only is the correct response, and the NIST/NCCoE framework is the correct implementation model.

What these frameworks do not do is differentiate risk treatment by confidentiality horizon and consequence severity in a way that surfaces single-assumption concentration as a distinct risk category. A Tier-1 CI link carrying nuclear command data with a permanent confidentiality requirement and a routine enterprise VPN carrying business communications with a five-year sensitivity horizon receive the same PQC migration treatment under current frameworks: migrate to NIST-approved algorithms, enforce crypto-agility, monitor for downgrade. The frameworks do not ask whether the Tier-1 link's consequence profile — catastrophic and irreversible if the algorithmic assumption fails — warrants a different hedging posture than the enterprise VPN's.

In every other domain where the US stakes national security on layered defense, single-point-of-failure analysis is foundational. Nuclear deterrence rests on a triad, not a monad. [1] Missile defense is layered across boost, midcourse, and terminal phases. Cyber defense doctrine — including the ICIT convergence paper's own framework — assumes defense-in-depth as a structural principle. Quantum-era cryptographic resilience is the domain where the US has chosen to rely on a single class of assumption without conducting,

or at least without publishing, an explicit assessment of whether that concentration is acceptable for the links where failure is catastrophic.

The gap is not that the frameworks are wrong. It is that they are undifferentiated. A risk framework that treats all CI links as candidates for the same PQC-only remediation will correctly serve 99 percent of the estate and systematically underassess the residual risk on the fraction of a percent where the consequences of algorithmic failure are permanent and unrecoverable.

VII.B. Authentication Dependency as an Audit Requirement

Section III.A established that any QKD deployment — PRC, EU, or the emergent US commercial experiments — inherits PQC vulnerability at the authentication layer. This has a direct implication for US CI risk assessment that applies regardless of US QKD policy: US CI operators who interact with QKD-protected systems (through allied interconnections, multinational financial networks, or supply chain relationships with EU or PRC-adjacent infrastructure) need to assess whether those QKD systems have completed PQC migration at their authentication layers.

Current risk frameworks do not require this assessment. A US financial institution whose settlement traffic transits an EuroQCI-protected segment has no framework-mandated mechanism for determining whether that segment's QKD authentication is PQC-hardened, classical, or in transition. The segment may be labeled "quantum-secure" while its authentication layer inherits the Valley of Death exposure that the ICIT convergence paper documented. The false confidence problem Section III.A described is not hypothetical — it is an operational condition that will emerge as EuroQCI and PRC QKD infrastructure appear in multinational CI service delivery chains.

The implication is that PQC migration audits for US CI operators need to extend beyond the operator's own cryptographic estate to include authentication-layer assessments for QKD-protected segments in connected infrastructure. This is an extension of the CBOM and supply-chain visibility logic the ICIT convergence paper advocated — applied to the QKD dimension that the convergence paper identified but did not develop.

VII.C. Downgrade Telemetry Across Organizational Boundaries

Section III.B documented the partial-deployment downgrade problem — dual-stack fallback creating silent reversion to classical on critical paths. Section V.A documented how decentralized US QKD adoption compounds this into a fragmented landscape where downgrade surfaces exist at organizational boundaries between CI operators, hyperscalers, backhaul providers, and counterparties.

Current cryptographic telemetry requirements — to the extent they exist in CSWP 48 and NCCoE guidance — are scoped to the operator's own estate. They do not address cross-organizational downgrade visibility. A CI operator can monitor its own PQC enforcement and detect fallback within its managed infrastructure. It cannot, under current frameworks, monitor whether the hyperscaler backbone carrying its Tier-1 traffic

has silently downgraded from PQC to classical under load, or whether the LEO backhaul segment has fallen back from PQC to pre-migration classical because the constellation operator's migration timeline lags the CI operator's.

For Tier-1 links, this gap is consequential. The end-to-end cryptographic posture of a transaction is bounded by its weakest segment, and the weakest segment is likely to be one the CI operator does not control and cannot currently monitor. Extending cryptographic telemetry requirements to include cross-organizational downgrade reporting — through contractual SLAs, sector-level agreements, or regulatory mandate — is a prerequisite for any Tier-1 risk assessment that claims to reflect actual rather than assumed cryptographic posture.

VII.D. The Defense-in-Depth Evaluation for Tier-1 Links

The most consequential implication is the one the paper has been building toward: whether US Tier-1 CI links warrant a formal defense-in-depth evaluation that considers physics-based hedging options alongside PQC.

This is not a recommendation to deploy QKD. It is a recommendation to conduct the assessment. The assessment would need to address at minimum:

The threat model for Tier-1 links specifically — distinguishing permanent-confidentiality requirements (nuclear command, certain financial settlement, bulk power control authentication) from the general CI estate where PQC-only is appropriate. The assessment should quantify the residual risk of single-assumption concentration against the specific confidentiality horizons and consequence severities of Tier-1 traffic, not against the general case.

The conditional utility of QKD as a PQC-failure hedge — acknowledging the authentication dependency, the partial-weakening scenario analysis from Section V.A, and the conditions under which QKD provides genuine fallback versus false confidence. An assessment that treats QKD as unconditionally useful is as analytically wrong as one that treats it as categorically impractical.

The architectural options that differ from PRC-style terrestrial trusted-node deployment — including LEO constellation-based approaches with their different risk profiles on the concentrated-node dimension, limited point-to-point deployments on identified Tier-1 segments rather than continental backbone buildout, and hybrid approaches that use QKD for key establishment on a small number of catastrophic-consequence links while PQC covers the general estate.

The governance requirements — CALEA compatibility for any telecommunications-adjacent deployment, oversight mechanisms for trusted-node access, and the coordination frameworks needed to manage cross-organizational cryptographic posture on Tier-1 service delivery chains.

The cost and complexity tradeoffs — including the LOTL-surface expansion and transition complexity costs Section III.F documented, measured against the specific risk reduction on Tier-1 links rather than against a generic cost-benefit for the full CI estate.

The point is not that this assessment will necessarily conclude in favor of QKD deployment. It may conclude that the single-assumption posture is acceptable for Tier-1 links after explicit evaluation — and that would be a defensible outcome. What is not defensible is the current condition: a single-assumption posture on the nation's highest-value CI links that is the product of general-purpose technology assessment rather than Tier-1-specific risk analysis, adopted by institutional inertia rather than examined risk acceptance.

VII.E. Persistent Competition and Long-Horizon Resilience

The risk-posture divergence between the US and PRC has strategic consequences that extend beyond CI resilience into the domain of persistent cyber competition. Those consequences are poorly served by classical deterrence framing and require a different analytical lens.

Deterrence models assume a decision-theoretic structure: an adversary contemplates a discrete action, calculates the probability of success against the defender's posture, weighs the expected return against the cost and risk of retaliation, and decides whether to act. Applied to cryptographic posture, the deterrence argument would hold that a defense-in-depth posture raises the adversary's perceived cost of achieving Tier-1 access, thereby discouraging the attempt. That logic is structurally mismatched with the operational environment this paper describes.

The cyberspace operations environment — as the persistent engagement doctrine [41][42], CYBERCOM's defend-forward posture, and the operational record of campaigns like Volt Typhoon [43] make clear — is not a domain of discrete attacks separated by periods of peace that a deterrent threat can prevent. It is a domain of continuous contact, where adversaries are already present inside critical infrastructure networks, already harvesting encrypted traffic under HNDL assumptions, and already pre-positioned on management planes through LOTL techniques for future exploitation. These activities are not prospective threats amenable to deterrence. They are baseline conditions of the competitive environment.

Cryptographic posture does not deter any of them. Passive harvest of encrypted traffic is undetectable and deniable; no defender posture prevents it. LOTL pre-positioning operates below the threshold where the defender's cryptographic architecture is relevant to adversary access — the adversary is already inside, using legitimate tools and credentials, and the question is not whether they maintain access but what that access yields over time. The PRC's dual posture does not deter collection against PRC systems. The US single-assumption posture does not invite collection that would not otherwise occur. Both states' Tier-1 traffic is being harvested now, and both states' management planes are contested terrain now.

What cryptographic posture does determine — and what the US–PRC divergence makes consequential — is the long-horizon payoff structure of the adversary's persistent presence. This is a resilience question, not a deterrence question, and it is the question that matters for strategic competition in a persistent-engagement environment.

Under a single-assumption posture, the adversary's return on persistent engagement is concentrated. A single event — the weakening or failure of the PQC algorithm family protecting Tier-1 traffic — converts the entire HNDL corpus harvested over years of patient collection into readable material. The same event

converts persistent management-plane access into a viable path for real-time cryptographic compromise of the defender's highest-value communications. Years of positioning, collection, and waiting produce a concentrated payoff at a single algorithmic threshold. The adversary's investment in persistent presence has a binary return structure: below the threshold, the harvested corpus and pre-positioned access have latent but unrealized value; above it, they unlock comprehensive Tier-1 access simultaneously.

Under a defense-in-depth posture, the adversary's return on persistent engagement is distributed and compound. The same algorithmic advance that would unlock the entire single-assumption estate yields only partial returns against a dual-layer defender: PQC-protected traffic is exposed, but key material exchanged over QKD channels remains information-theoretically secure — its protection does not degrade with advances in cryptanalysis or quantum computation. To achieve the same comprehensive Tier-1 access that a single algorithmic event provides against a single-assumption defender, the adversary must combine the algorithmic advance with independent compromise of the QKD infrastructure — trusted-node exploitation, authentication-layer attacks, or supply-chain compromise of quantum hardware. Neither component is deterred, but the compound requirement means that no single capability advance, however dramatic, is sufficient to unlock the full return on persistent engagement.

The strategic implication is not about deterrence — it is about how each state's cryptographic architecture shapes the adversary's planning horizon and capability requirements for achieving strategic access. A single-assumption defender creates a condition where the adversary can plan against a single technical milestone: the day PQC weakens. All collection, all pre-positioning, all patient infrastructure mapping converges on that single event. A defense-in-depth defender denies the adversary that convergence point. The adversary must plan against multiple independent technical milestones — algorithmic advance plus infrastructure compromise — with no guarantee that achieving one delivers the other.

In a persistent competition environment where adversaries are already present and already collecting, the defender cannot control whether persistent engagement occurs. What the defender can control is whether the long-term architecture of its most consequential links allows a single capability advance to convert years of adversary patience into comprehensive strategic access — or whether the architecture imposes compound requirements that bound the return on any single advance. That is the resilience dimension of the US–PRC divergence, and it is the dimension that a Tier-1 defense-in-depth evaluation should assess.

The United States currently maintains a cryptographic architecture on its Tier-1 CI links where the adversary's entire persistent-engagement investment converges on a single algorithmic threshold. The PRC maintains an architecture where it does not. Whether that asymmetry is acceptable is a strategic judgment. That it should be made explicitly — rather than inherited as a byproduct of general-purpose technology assessment — is the minimum standard that a persistent-competition environment demands.

VIII. Findings and Conclusion: The Question Defense-in-Depth Requires

This paper began with a specific question: does the United States apply the same defense-in-depth rigor to quantum-era cryptographic resilience that it applies to nuclear deterrence, missile defense, cyber architecture, and every other domain where it stakes national security on layered protection? The analysis supports a clear answer: it does not — and the gap between doctrine and practice is consequential for the small number of CI links where the cost of single-assumption failure would be catastrophic and irreversible.

Findings

PQC and QKD are structurally coupled, not independent alternatives. The five coupling mechanisms documented in Section III — authentication dependency, partial-deployment downgrade, hardware maturity gaps, concentrated-node vulnerability, and compound interaction effects during simultaneous migration — establish that no risk framework can evaluate either technology in isolation. QKD inherits PQC's failure modes at the authentication layer. PQC-only postures are affected by QKD's global emergence through allied interconnections, multinational service delivery chains, and decentralized commercial adoption. Treating PQC and QKD as separable technology choices — the framing that currently governs US policy — systematically misestimates quantum-era CI exposure.

The US and PRC have made structurally divergent risk bets, neither of which dominates across all dimensions. The PRC has purchased defense-in-depth against algorithmic failure — a physics-based cryptographic fallback on Tier-1 links whose security does not depend on computational hardness assumptions — at the cost of 145 permanent classical infrastructure chokepoints, governance entanglement between cryptographic assurance and state access, and elevated LOTL-exploitable complexity during the transition period. The United States has purchased operational simplicity and a narrower transitional attack surface at the cost of concentrating all Tier-1 CI protection on a single class of mathematical assumption. Each state accepts risks the other avoids. The divergence is in risk allocation, not deployment scale.

The US is not operating the PQC-only posture its federal policy describes. Decentralized QKD adoption by financial institutions, DOE-funded demonstrations on operational grid infrastructure, and commercial quantum network development are producing an emergent patchwork in which some Tier-1 traffic segments carry QKD protection, others carry PQC, and others remain classical during transition. This patchwork reproduces the compound failure modes of a deliberate dual deployment — partial-deployment downgrade at organizational boundaries, expanded management plane complexity, priority inversion where QKD presence fingerprints high-value traffic — without the corresponding algorithmic-failure hedge and without centralized visibility into the end-to-end cryptographic posture. The US is drifting toward the complexity costs of both postures while securing the resilience benefits of neither.

In a persistent-competition environment, single-assumption concentration creates a convergent payoff structure for adversary persistent engagement. Adversaries are not waiting to be deterred from collection or pre-positioning; they are already present in CI networks, already harvesting encrypted traffic, and already dwelling on management planes through LOTL techniques. Cryptographic posture determines not whether these activities occur but what they yield over time. A single-assumption architecture means a single algorithmic advance converts the adversary's entire HNDL corpus and persistent access into comprehensive Tier-1 penetration. A defense-in-depth architecture imposes compound requirements — algorithmic advance plus independent infrastructure compromise — that deny the adversary a single convergence point. The US currently maintains the architecture that permits convergence. The PRC does not.

Current US risk frameworks are undifferentiated in their treatment of quantum-era cryptographic risk. NIST CSF 2.0, SP 800-53, CSWP 48 [3], and NCCoE's PQC migration guidance [30] correctly serve the general CI estate — the 99 percent of links where PQC-only is the appropriate, cost-effective response. They do not differentiate risk treatment by confidentiality horizon and consequence severity in a way that surfaces single-assumption concentration as a distinct risk category for Tier-1 links with permanent confidentiality requirements and catastrophic failure consequences. A framework that treats all CI links as candidates for the same remediation will correctly serve the general case and systematically underassess the residual risk on the links where failure is unrecoverable.

Conclusion

This paper does not conclude that the United States should deploy QKD. That determination requires a Tier-1-specific risk assessment that does not currently exist, informed by classified threat timelines, operational cost modeling, architectural analysis of options that differ from PRC-style terrestrial trusted-node deployment, and governance engineering that resolves the CALEA and oversight tensions Section VI documented.

What this paper concludes is that the assessment must be conducted.

The PRC's posture has the characteristics of an examined risk acceptance: the investment is deliberate, the costs are structural and visible, and the tradeoffs are embedded in the architecture by design. The US posture has the characteristics of an unexamined default: NSA and NIST guidance correctly identifies QKD's limitations for general-purpose deployment, and the policy apparatus has treated that general-purpose assessment as dispositive for all CI links — including the Tier-1 links where the risk calculus is structurally different from the general case. The question of whether a physics-based fallback is warranted on nuclear command segments, financial settlement backbones, and bulk power control networks has not been asked with the specificity those assets demand.

The ICIT convergence paper [1] established that the window to embed cryptographic resilience into AI and LEO infrastructure closes in the early 2030s. This paper extends that finding: the same window applies to the defense-in-depth question. If the assessment is deferred until after PQC migration is substantially complete, the infrastructure decisions that would have enabled a layered posture — fiber allocation, trusted-node or

constellation architecture, authentication-layer diversification, governance frameworks — will have been foreclosed by the same capex-cycle lock-in the ICIT paper warned against for PQC.

The United States applies defense-in-depth as a foundational principle in every domain where it stakes national security on layered protection. Whether quantum-era cryptographic resilience for Tier-1 critical infrastructure warrants the same treatment is a question that deserves an explicit answer — not the implicit one that institutional inertia has provided.

Bibliography

References

1. Mussington, D. (2026). *Quantum-Resilient Convergence: The Shared Defense of AI, Space, and Critical Infrastructure*. Institute for Critical Infrastructure Technology. February 2026.
2. NIST. (2024). *Transition to Post-Quantum Cryptography Standards* (IR 8547). U.S. Department of Commerce.
3. NIST. (2025). *Mappings of Migration to PQC Project Capabilities to Risk Management Frameworks* (CSWP 48).
4. NSA. (2024). *The Commercial National Security Algorithm Suite 2.0 and Quantum Computing FAQ*. Ver. 2.1, December 2024.
5. Chen, Y.-A. et al. (2025). "Implementation of carrier-grade quantum communication networks over 10,000 km." *npj Quantum Information*. August 2025.
6. U.S.-China Economic and Security Review Commission. (2025). *Vying for Quantum Supremacy: U.S.-China Competition in Quantum Technologies*.
7. Chinese Academy of Sciences. (2025). "Chinese-led Team Achieves World's First 10,000km Quantum-secured Communication." March 20, 2025.
8. Liao, S.-K. et al. (2018). "Satellite-Relayed Intercontinental Quantum Network." *Physical Review Letters* 120, 030501.
9. Sander, O. et al. (2025). "South Africa and China set up a quantum communication link." *The Conversation*. August 11, 2025.
10. European Commission. (2024). *European Quantum Communication Infrastructure* (EuroQCI) Initiative.
11. European Commission. (2026). *Proposed Amendment to NIS2 Directive: Post-Quantum Cryptography Requirements*. January 2026.
12. NSA. (2020). "Quantum Key Distribution (QKD) and Quantum Cryptography (QC)." Cybersecurity Perspectives.
13. DoD Chief Information Officer. (2025). *Post-Quantum Cryptography Directive*. November 2025.
14. OMB. (2022). *M-23-02: Migrating to Post-Quantum Cryptography*.
15. The White House. (2025). *Executive Order 14144: Strengthening and Promoting Innovation in the Nation's Cybersecurity*. January 2025.
16. DOE. (2025). "Energy Department Announces $625 Million to Advance the Next Phase of National Quantum Information Science Research Centers." November 4, 2025.
17. DARPA. (2025). "Quantum mechanics, classical backbone: DARPA's QuANET advances practical quantum networking." August 2025.
18. DARPA. (2023). "A Network Security Revolution Enhanced by Quantum Communication." QuANET Program Announcement.
19. IonQ. (2025). "IonQ Announces $22M Deal with EPB Establishing Chattanooga, Tennessee as the First Quantum Computing & Networking Hub in the U.S." April 25, 2025.

20. EPB. (2025). “EPB and IonQ Partner to Establish Chattanooga as the First Quantum Computing and Networking Hub in the U.S.” April 25, 2025.
21. The Quantum Insider. (2026). “Inside Tennessee’s Growing Quantum Ecosystem and its Federal Impact.” March 6, 2026.
22. JPMorgan Chase. (2024). “JPMorgan Chase establishes quantum-secured crypto-agile network.” Presented at OFC’24.
23. ID Quantique. (2024). “IDQ enables successful 100Gbps IPSec VPN demo at JPMorgan Chase.” September 10, 2024.
24. ID Quantique. (2024). “Why the financial sector is embracing a dual quantum-safe strategy.” December 12, 2024.
25. JPMorgan Chase, Toshiba, and Ciena. (2022). “First Quantum Key Distribution Network Used to Secure Mission-Critical Blockchain Application.” February 17, 2022.
26. ISACA Journal. (2025). “Building Resilient Security in the Age of Quantum Computing.” Volume 6, 2025.
27. NIST. (2024). *FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard*.
28. NIST. (2024). *FIPS 204: Module-Lattice-Based Digital Signature Algorithm Standard*.
29. NIST. (2024). *FIPS 205: Stateless Hash-Based Digital Signature Algorithm Standard*.
30. NCCoE. (2024). *Migration to Post-Quantum Cryptography* (SP 1800-38).
31. GSA. (2025). *Post Quantum Cryptography Buyer’s Guide*.
32. CISA, NSA, & NIST. (2023). *Quantum-Readiness: Migration to Post-Quantum Cryptography*.
33. Kampanakis, P. et al. (2025). “Post-Quantum Hybrid Key Exchange in TLS 1.3.” IETF Internet-Draft.
34. White House. (2024). *National Security Memorandum on Critical Infrastructure Security and Resilience* (NSM-22).
35. RAND Corporation. (2025). “U.S.-Allied Militaries Must Prepare for the Quantum Threat to Cryptography.” June 2025.
36. ANSSI, BSI, NLNCSA, SWEDNCSA. (2024). *Position Paper on Quantum Key Distribution*.
37. Yin, J. et al. (2017). “Satellite-based entanglement distribution over 1200 kilometers.” *Science 356*, 1140–1144.
38. ORNL/LANL/EPB. (2019–2025). QKD for grid security research program. DOE CESER.
39. Quantum Zeitgeist. (2026). “The Ultimate Guide To China Quantum Computing Companies In 2026.” March 2026.
40. NSA. (2024). *CNSA 2.0 FAQ*. Ver. 2.1. “NSA does not generally consider QKD a practical security solution for protecting NSS.”
41. Nakasone, P. (2019). “A Cyber Force for Persistent Operations.” *Joint Force Quarterly 92*.
42. U.S. Cyberspace Solarium Commission. (2020). *Final Report*. March 2020.
43. CISA. (2024). *PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure*. Joint Advisory AA24-038A (Volt Typhoon).
44. Communications Assistance for Law Enforcement Act (CALEA). 47 U.S.C. §§ 1001–1010.

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