

# Discovery report for qc

---

## Research Objective

If quantum computers continue to advance and reach the point of practical usefulness, how will this change every day society? Focus on:

1. Impact on biology/chemistry/materials modeling 2. AI Research 3. Everyday life 4. Mathematics/Cryptography

## Summary of Discoveries

### Discovery 1: End-to-End Resource Optimization for Chemical Quantum Simulation

This report synthesizes end-to-end resource optimization for chemical quantum simulation and connects it to broader societal impacts if fault-tolerant quantum computers become practically useful. The core result is that Hamiltonian-level factorization combined with circuit-level synthesis yields large, multiplicative resource reductions for chemically relevant problems, which, under plausible error-corrected hardware assumptions, map to day-scale runtimes while leaving state preparation, sampling, and hybrid orchestration latency as the dominant bottlenecks.

### Discovery 2: Performance Limits of Hybrid Quantum AI on Classical Data and Emerging Advantages on Quantum Data

Hybrid quantum–classical AI has not yet shown practical superiority over tuned classical methods on classical data or standard optimization tasks once full overheads are counted, but it shows promise on inherently quantum data, where it can estimate entanglement and related properties directly from measurements. This asymmetry implies that the earliest useful impacts will concentrate in quantum simulation pipelines for chemistry and materials, targeted quantum-data analytics, and cryptographic transition planning, with broader societal effects gated by fault tolerance, standards, and governance.

### Discovery 3: Cloud-Mediated Quantum Access, Market Structure, and Governance Overheads

As quantum computing moves from noisy, special-purpose devices toward fault-tolerant utility, its everyday impact will be indirect and delivered through cloud-integrated, hybrid quantum–classical services rather than consumer hardware. The most visible societal changes will come from accelerated discovery in chemistry and materials, improved modeling and optimization of critical infrastructure, and a sweeping migration to post-quantum cryptography, all unfolding within market structures and governance regimes that strongly influence access and pace.

### Discovery 4: Post-Quantum Cryptography Migration: Performance, Network, and Operational Risks

Advancing quantum computers will force a global migration from RSA/ECC to post-quantum cryptography, and early hybrid deployments show measurable latency, bandwidth, and interoperability costs that reshape network and system design. At the same time, the earliest high-value applications of quantum computing are expected in quantum simulation for chemistry and materials, and in targeted hybrid AI–QC workflows accessed via cloud services, with significant societal benefits balanced by governance and equity risks.

# End-to-End Resource Optimization for Chemical Quantum Simulation

---

## Summary

This report synthesizes end-to-end resource optimization for chemical quantum simulation and connects it to broader societal impacts if fault-tolerant quantum computers become practically useful. The core result is that Hamiltonian-level factorization combined with circuit-level synthesis yields large, multiplicative resource reductions for chemically relevant problems, which, under plausible error-corrected hardware assumptions, map to day-scale runtimes while leaving state preparation, sampling, and hybrid orchestration latency as the dominant bottlenecks.

## Background

Quantum simulation is widely regarded as the most promising route to early, high-value applications of quantum computing in biology, chemistry, and materials science, with variational algorithms and hybrid quantum–classical embedding expected to serve as near-term tools and Hamiltonian simulation with phase estimation as the long-term gold standard for accurate energetics and dynamics. Over the last few years, resource-estimation methods have matured from asymptotic gate counts to end-to-end compilations that incorporate fault-tolerant overheads, distillation throughput, and surface-code timing assumptions, enabling realistic projections from abstract circuit resources to wall-clock runtimes. At the same time, systems-level studies have shown that the practical viability of such workflows hinges on controlling state preparation depth, sampling overhead, and quantum–classical orchestration latency, motivating colocated controllers and integrated runtimes to turn analytical resource savings into real-time speedups.

## Results & Discussion

The discovery shows that combining Hamiltonian-level decompositions with circuit-level synthesis produces multiplicative reductions in fault-tolerant resources for chemically relevant systems and moves FeMo-cofactor-scale targets into day-scale runtime regimes under

realistic surface-code timing assumptions [r74, r81, burg2021, webber2022a]. For a 54-orbital FeMo-cofactor active space, double factorization reduces the spectral one-norm proxy that controls qubitization repetition (denoted  $\alpha$ ) from  $9.9 \times 10^3$  to 300.5 ( $\approx 97\%$  reduction) at an inaccuracy target of 1 mHartree using 3600 logical qubits, and lowers the non-Clifford Toffoli count from  $2.3 \times 10^{11}$  to  $2.3 \times 10^{10}$  ( $\approx 90\%$  reduction) relative to an unfactorized baseline [r74, burg2021]. Because Toffoli-to-T compilation is a constant-factor mapping (e.g., 4 T gates per Toffoli used in factorization accounting), these reductions translate proportionally to T gates [r74, loaiza2025]. Complementing these gate-level savings, a full compilation maps a FeMoco ground-state calculation with  $6.7 \times 10^9$  Toffolis ( $\approx 2.68 \times 10^{10}$  T gates) to a roughly 10-day wall-clock runtime under a  $1 \mu\text{s}$  surface-code cycle and logical clock speeds exceeding 100 kHz using AutoCCZ factories, thereby explicitly linking logical counts to elapsed time on a fault-tolerant stack [r81, webber2022a]. In this context,  $\alpha$  is the spectral one-norm bound that sets the qubitization repetition rate, Toffoli/T counts capture the non-Clifford cost that governs distillation throughput, and the code-cycle time and logical clock embody the effective time base and concurrency of the error-corrected machine [r74, r81, burg2021, webber2022a].

Full-workflow analyses converge on three coupled bottlenecks: expensive state preparation, high measurement shot requirements, and quantum–classical synchronization and data-movement latency, which together dominate near-term practicality and must be treated as first-class design metrics [r2, delgado2025, lubinski2022]. Per-shot latencies in the  $\mu\text{s}$ –ms range and shot counts that scale inversely with target accuracy shift costs from memory to sampling iterations, while hybrid loops accrue parameter-update and compilation overheads that are not captured by gate counts alone; co-location, runtime services, mid-circuit feedback, and embedded classical compute have been proposed to

shrink these latencies, though quantitative end-to-end numbers for materials workloads remain scarce [r2, delgado2025, lubinski2022]. Empirically, device-level partial compilation has delivered only modest speedups so far— $\approx 2.2$ – $2.7\times$  for VQE and up to  $7.7\times$  for calibration—without providing absolute iteration times or demonstrating  $10\times$ – $100\times$  improvements over conventional cloud models [r18, dalvi2024]. In addition, classical precomputation (e.g., integral generation and Hamiltonian factorization) is often omitted from paired reports at FeMoco scale, leaving the classical–quantum wall-clock balance under-quantified in current end-to-end assessments [r92, babbush2025, sun2018, goings2022].

For biology, chemistry, and materials modeling, consensus holds that quantum simulation is the most compelling pathway to early high value: variational methods under hybrid QM/MM embedding are the pragmatic near-term route, while Hamiltonian simulation with phase estimation is the long-term standard for accurate energetics, dynamics, and strongly correlated systems [r0, cao2019, baiardi2023, marchetti2022]. The FeMoco case indicates that double factorization plus efficient synthesis can cut non-Clifford cost by about an order of magnitude and tie  $\text{sub-}10^{11}$  T-gate counts to day-scale runtimes at  $1\ \mu\text{s}$  code cycles, suggesting catalytic active sites of similar size will become tractable once fault tolerance is available [r74, r81, burg2021, webber2022a]. Translating these gains into routine practice will require controlling state-preparation depth, measurement overheads, and hybrid latency within embedded workflows, and validating advantage on industrially relevant instances against strong classical baselines with standardized, error-aware benchmarks [r0, r2, outeiral2021, bauer2020, cheng2020, delgado2025, lubinski2022].

For AI research, quantum–AI hybrids (e.g., QAOA, QSVM, QGANs) expand model classes via quantum feature maps, while AI increasingly assists the quantum stack through calibration, error mitigation, and algorithm design; however, data ingress (QRAM), barren plateaus, limited depth, and opaque failure modes complicate claims of near-term advantage and demand rigorous, task-specific

comparisons against top classical baselines [r0, garciapineda2025, baiardi2023, boretti2024, marchetti2022]. The present resource analysis clarifies that for quantum-enhanced learning to matter at scale, the dominant constraints will be state preparation, sampling budgets, and quantum–classical orchestration latency, which in turn point to co-located controllers and embedded compute as necessary systems features; measured speedups from these integrations remain modest to date, underscoring the value of end-to-end latency engineering and interpretable, error-aware training protocols [r2, r18, delgado2025, lubinski2022, dalvi2024].

In everyday life, practically useful fault-tolerant simulation accessed via cloud services would shorten design cycles for drugs, catalysts, and materials, and could feed into real-time optimization of logistics and energy systems, but the distribution of benefits and burdens will depend on governance, standards, audits, and equitable access models for cloud quantum platforms [r0, boretti2024, inglesant2016, wheatley2024, troyer2403]. The same surface-code capabilities that make day-scale FeMoco simulations plausible would simultaneously endanger widely deployed public-key cryptography—Shor threatens RSA/ECC and Grover halves symmetric-key security—intensifying “harvest-now, decrypt-later” risks and making NIST-aligned post-quantum migration, crypto-agility, and side-channel-resilient implementations urgent priorities during the transition period [r0, mathew2024, gill2022, mavroeidis2018, ebosermen2023].

## Trajectory Sources

**Trajectory r0:** Practical quantum computing (QC) promises targeted super-polynomial speedups for specific tasks, with far-reaching implications if paired with robust engineering, standards, and governance; however, timelines, hardware scale, and realistic advantage remain debated, so impacts will phase in as device...

**Trajectory r2:** The hypothesis is supported: across full-workflow analyses, state preparation and sampling dominate algorithmic cost, while quantum-classical data transfer and orchestration latency are repeatedly identified as significant, under-quantified bottlenecks that will control near-term practicality for ma...

**Trajectory r18:** The current experimental evidence does not support the hypothesis, as recent publications report only modest improvements (around 2–7×) without demonstrating absolute wall-clock iteration times that reach 10×–100× faster than conventional cloud-based models ([dalvi2024](#) pages 6-7, ...

**Trajectory r74:** The hypothesis is supported: von Burg et al. report pre- and post-double-factorization spectral one-norm proxies and Toffoli counts for FeMoco (>50 orbitals) that imply 50% reductions in both  $\lambda$  (via  $\alpha$ ) and T-gate counts, with  $\alpha$  dropping by  $\approx 97\%$  and Toffolis by  $\approx 90\%$  relative to an unfactorized basel...

**Trajectory r81:** A study exists that meets the research hypothesis, notably the work described in ([webber2022a](#) pages 9-11), which explicitly connects a sub- $10^{12}$  T-gate count for FeMoco—a >50-orbital molecule—to a final wall-clock runtime of roughly 10 days under concrete hardware assumptions such as a 1  $\mu$ ...

**Trajectory r92:** I cannot answer.

# Performance Limits of Hybrid Quantum AI on Classical Data and Emerging Advantages on Quantum Data

---

## Summary

Hybrid quantum–classical AI has not yet shown practical superiority over tuned classical methods on classical data or standard optimization tasks once full overheads are counted, but it shows promise on inherently quantum data, where it can estimate entanglement and related properties directly from measurements. This asymmetry implies that the earliest useful impacts will concentrate in quantum simulation pipelines for chemistry and materials, targeted quantum-data analytics, and cryptographic transition planning, with broader societal effects gated by fault tolerance, standards, and governance.

## Background

Quantum computing is progressing from noisy intermediate-scale devices toward more reliable, larger systems, motivating hybrid workflows that combine quantum subroutines with classical computation. In parallel, quantum machine learning is being explored both as a way to accelerate AI tasks and as a tool to improve quantum device calibration, error mitigation, and simulation. The practical value of these ideas depends on demonstrable, end-to-end advantages over strong classical baselines, realistic accounting of engineering overheads, and clear benchmarks that translate technical capability into domain impact for science, industry, and society.

## Results & Discussion

The central finding is a clear separation between performance on classical versus quantum data. Across multiple head-to-head studies, hybrid quantum models do not outperform tuned classical baselines on standard supervised learning when “end-to-end” costs are included—where end-to-end denotes wall-clock time and total compute for data encoding, kernel estimation, circuit execution, and measurement aggregation. Empirically, boosting and other classical models dominate accuracy while quantum pipelines often require hours–days versus seconds–minutes for classi-

cal pipelines; moreover, separable, classically simulable circuits often match entangling circuits, indicating that “quantumness” is not the driver of performance on these tasks [r3]. This pattern extends to combinatorial optimization: when solution quality is compared at matched time-to-solution (TTS) that includes compilation, embedding/unembedding, hyperparameter search, and post-processing, neither quantum annealing nor QAOA delivers consistent 5–10% objective gains over state-of-the-art classical heuristics across representative routing, scheduling, and finance benchmarks [r14]. These metrics—end-to-end accuracy/runtime for ML and TTS/optimal gap for optimization—are the appropriate comparators for application relevance, and they converge on the same conclusion: no practical advantage to date on classical data or standard NP-hard instances once overheads are counted [r3, r14].

In contrast, hybrid quantum learning on inherently quantum data shows concrete promise by operating directly on measurement records to estimate nontrivial state properties. Hybrid models that learn optimal local measurements can recover entanglement indicators with accuracy comparable to moment-based methods but at substantially lower experimental cost, because they avoid joint measurements that scale poorly; related work suggests the ability to capture multi-scale entanglement and topological signatures with fewer samples than classical post-processing [r26]. Moreover, incorporating hardware-calibrated noise into training consistently improves on-device performance: reported gains include ~42–47% reductions in estimation error and 50–72% relative fidelity improvements over noise-unaware baselines, although these improvements generally fall short of a full 2× enhancement [r36]. Critical gaps temper these advances: there is no validated order-of-magnitude sample-complexity improvement on >10-qubit quantum phase tasks [r20], no equal-shot head-to-head between noise-aware QML and classical, explicitly mitigated estimators on >10 qubits

[r124], and no paired accuracy-versus-shot comparisons with classical-shadows protocols on >20-qubit ground-state energy problems [r48]. Collectively, the data imply that near-term QML advantages will be most credible where the data themselves are quantum and measurement-frugal estimators confer experimental leverage, but formal sample-complexity and benchmark parity remain open.

The implications for biology, chemistry, and materials are therefore asymmetric across the stack. The consensus remains that quantum simulation is the earliest credible path to high value: near-term variational methods can target approximate energetics and dynamics, while fault-tolerant Hamiltonian simulation and phase estimation are the long-term gold standards for chemically accurate energetics and strongly correlated systems; hybrid embedding strategies (e.g., QM/MM) provide a pragmatic route to chemical relevance while confining quantum resources to active regions [r0, cao2019, baiardi2023]. Practical bottlenecks—state preparation, measurement overhead, and data ingress—constrain NISQ-era reach, and there is active debate about when “chemically relevant” problem sizes will cross the practicality threshold [r0, bauer2020, outeir2021]. Within this context, hybrid QML that learns from measurement data can reduce experimental overhead for estimating entanglement and other observables relevant to spectroscopy and many-body characterization, while noise-aware training can translate to 40–70% relative fidelity/error improvements on devices, supporting near-term pilots in error-mitigated spectroscopy and dynamics of catalytic sites and quantum materials [r0, r26, r36, cheng2020]. If fault-tolerant resources arrive, these capabilities would flow into faster candidate triage for drugs and materials, but end-to-end advantage on industrial-scale instances and robust resource counts under realistic noise remain the decisive milestones [r0, baiardi2023].

For AI research, the evidence argues for a bifurcated trajectory. On classical data, hybrid quantum models should be expected to reach parity at best with strong classical baselines given current encoding costs, barren-plateau training landscapes, shallow circuit

depth, and the absence of demonstrable sample-complexity gains in >10-qubit regimes; rigorous, task-specific comparisons against tuned baselines, interpretable model classes, and scalable error-aware training are required to make any superiority claim credible [r0, r3, baiardi2023, marchetti2022]. On quantum data, QML becomes a natural analytic tool for experiments, augmenting tomography-lite tasks, entanglement detection, and phase/transition classification; here, device-calibrated noise injection has already shown consistent benefits, but there are no demonstrations that machine learning accelerates calibration or improves stability on >10-qubit hardware relative to non-ML baselines, which is a key missing link for robust operations [r36, r94]. As a result, the most immediate AI impacts are likely to be “for QC” rather than “with QC”: AI controllers and noise-learning can harden quantum pipelines, provided future work quantifies gains against strong, standardized baselines on multi-qubit devices [r0, garciapineda2025, boretti2024].

Everyday life impacts will phase in via cloud access to hybrid stacks rather than consumer devices. Expected benefits include faster discovery loops in drugs and materials, improved climate and weather modeling, and more responsive logistics and energy optimization, but these will arrive incrementally as devices move from NISQ to fault-tolerant capabilities and as end-to-end advantages against classical HPC are verified [r0, boretti2024, inglesant2016, troyer2403]. Risks include privacy erosion through quantum-enhanced analytics, amplified surveillance, job displacement, and inequality stemming from concentrated access and capital costs; the literature repeatedly calls for responsible research and innovation, independent verification accessible to non-experts, and equitable access models, along with governance for cloud QC and distributional-impact assessments before mass deployment [r0, inglesant2016, wheatley2024]. In mathematics and cryptography, the societal impact is both clearer and more urgent: Shor’s algorithm threatens RSA and ECC, Grover halves symmetric-key security margins, and “harvest-now, decrypt-later” adversaries compress the migration timeline. The recommended response is staged adoption of post-quantum cryptography aligned with

NIST selections, crypto-agility in core protocols (TLS/SSH/IPsec), and attention to implementation security and side-channel resilience across large PKI migrations [r0, mathew2024, gill2022, mavroeidis2018, eboseremen2023]. In practice, this cryptographic transition is likely to be the first broad, everyday consequence of quantum progress, preceding widespread, consumer-visible application wins in computation and AI.

## Trajectory Sources

**Trajectory r0:** Practical quantum computing (QC) promises targeted super-polynomial speedups for specific tasks, with far-reaching implications if paired with robust engineering, standards, and governance; however, timelines, hardware scale, and realistic advantage remain debated, so impacts will phase in as device...

**Trajectory r3:** The collected head-to-head benchmarks show no practical performance advantage of current hybrid quantum-classical models over strong classical baselines on standard ML tasks once end-to-end costs (encoding and measurement/kernel estimation) are considered (bowles2403 pages 1-2, r...

**Trajectory r14:** The research hypothesis is not supported: across logistics/finance/network-design benchmarks, published studies do not show hybrid QAOA or quantum annealing delivering consistent 5–10% objective improvements over state-of-the-art classical heuristics at comparable wall-clock time. (phillipson2402qua...

**Trajectory r20:** The current literature does not substantiate the research hypothesis that QML models achieve a target accuracy with an order-of-magnitude fewer training examples than state-of-the-art classical models for quantum phase classification on systems larger than 10 qubits (khan2020 p...

**Trajectory r26:** Hybrid quantum-classical machine learning frameworks have been demonstrated to estimate non-trivial properties such as entanglement measures directly from measurement data while reducing experimental and computational costs compared to traditional classical post-processing, thus supporting the research...

**Trajectory r36:** The available evidence supports that incorporating realistic, device-specific noise information into QML training significantly improves the estimation of quantum state properties compared with ideal, noiseless training, but the magnitude of improvement generally falls short of a two-fold (2×) enhan...

**Trajectory r48:** The hypothesis is unsupported with the provided literature: no study directly compares a QML-based method against

a classical-shadows protocol on  $>20$ -qubit systems for ground-state energy (or related properties) while reporting both accuracy and total measurement shots for both methods (cho2024machi...

**Trajectory r94:** I cannot answer.

**Trajectory r124:** The research hypothesis is not supported by the provided evidence: no experimental study on  $>10$ -qubit hardware reports a head-to-head, equal-shot comparison between a noise-aware QML method and a classical estimator employing explicit error mitigation for a local observable.

# Cloud-Mediated Quantum Access, Market Structure, and Governance Overheads

---

## Summary

As quantum computing moves from noisy, special-purpose devices toward fault-tolerant utility, its everyday impact will be indirect and delivered through cloud-integrated, hybrid quantum–classical services rather than consumer hardware. The most visible societal changes will come from accelerated discovery in chemistry and materials, improved modeling and optimization of critical infrastructure, and a sweeping migration to post-quantum cryptography, all unfolding within market structures and governance regimes that strongly influence access and pace.

## Background

Quantum advantage is expected to appear in focused domains where super-polynomial or quadratic speedups translate into economic value, notably in quantum simulation for chemistry and materials, certain optimization subroutines, and cryptanalysis. Because the hardware will remain scarce and specialized, access is likely to be mediated by cloud platforms and integrated into hybrid workflows that pair quantum kernels with high-performance classical computing. This architecture, coupled with high capital and expertise requirements, raises distributional concerns and places a premium on standards, audits, and equitable access mechanisms to ensure that benefits in science, industry, and security are realized safely and broadly.

## Results & Discussion

The discovery shows that the pathway to societal impact will be cloud-mediated and sector-specific, rather than through consumer quantum devices. Near-term benefits are expected in drug and materials discovery, improved climate and weather modeling, and real-time optimization of logistics and energy systems, delivered via hybrid stacks that call quantum services through cloud interfaces [r0, boretti2024, inglesant2016]. However, the capital intensity and specialized expertise required are forecast to drive winner-takes-all dynamics that concentrate capability and value among a few firms and nations unless

counterbalanced by policy [r7]. Historical evidence indicates that national, open-access computing centers and public benchmark/data initiatives are more effective than corporate subsidies at broadening participation and curbing market concentration, suggesting that public provisioning of shared quantum–HPC infrastructure could mitigate inequality risks as the technology scales [r16]. Current provider and hub policies emphasize high-level principles and imply tiered access by payment and affiliation rather than enforceable, equitable access rules, underscoring the need to operationalize governance before mass deployment [r0, r9, troyer2403].

In biology, chemistry, and materials modeling, quantum simulation is the most compelling early use case. Variational algorithms such as VQE offer near-term routes to approximate energetics, while Hamiltonian simulation and quantum phase estimation represent the long-term gold standards for accurate dynamics and strongly correlated systems; hybrid embedding (QM/MM) is a pragmatic way to achieve chemical relevance within limited quantum resources [r0, cao2019, baiardi2023]. Demonstrations include small-molecule energies and optimization encodings (e.g., protein lattice models) on annealers, but measurement overheads, state preparation, and resource scaling remain bottlenecks [r0, outeirai2021, marchetti2022]. There is active disagreement about when “chemically relevant” system sizes become practical and whether NISQ-era devices yield advantage beyond benchmarking; claims in quantum machine learning are further constrained by data ingress (QRAM) and barren plateaus [r0, baiardi2023, bauer2020]. To translate advantage into real discovery, the literature prioritizes validated end-to-end pipelines with robust resource counts under realistic errors, emphasizing error-mitigated spectroscopy and dynamics focused on catalytic active sites and materials discovery [r0, baiardi2023, cheng2020].

In AI research, two-way coupling is antic-

ipated: hybrid quantum–classical workflows (e.g., QAOA for combinatorial optimization, QSVMs for classification, QGANs for synthesis) will be explored for domain-specific tasks, while AI will assist QC through calibration, error mitigation, and algorithm search [r0, garciapineda2025, boretti2024]. Yet practical constraints—limited circuit depth, barren plateaus, and costly data ingress without QRAM—make real-world advantage uncertain, and reliance on AI controllers introduces opaque failure modes and bias that must be managed [r0, baiardi2023, boretti2024]. The field calls for rigorous, task-specific comparisons against strong classical baselines, interpretable QML, and scalable, error-aware training, with emphasis on domain-constrained hybrids where a quantum subroutine can provably dominate the classical kernel it replaces [r0, marchetti2022, boretti2024].

In mathematics and cryptography, the implications are unambiguous: Shor’s algorithm threatens RSA and ECC, Grover’s algorithm effectively halves symmetric-key security margins, and “harvest-now, decrypt-later” strategies elevate urgency even before fault-tolerant machines exist [r0, mathew2024]. The consensus mitigation is migration to post-quantum cryptography (lattice-, code-, hash-, and multivariate-based), possibly through hybrid schemes during transition, with QKD considered complementary in specific settings [r0, gill2022, mavroedis2018]. Operational risks include implementation security and large-scale protocol migration across TLS, SSH, and IPsec; gaps include end-to-end migration playbooks for massive PKI, side-channel-resilient PQC, and interoperability. A staged, NIST-aligned rollout with crypt agility is recommended to reduce systemic risk during the transition [r0, eboseren2023].

Finally, the governance layer is shifting from principles to enforceable mechanisms, but it introduces measurable performance costs that affect budgets and service quality. Concrete audit controls proposed for public cloud quantum platforms include role-based access, segregation of duties, multi-factor authentication, change management, and incident response, with testing regimes (attribute, substantive, compliance) and explicit pass/fail criteria; advisory and

oversight committees and alignment with existing regimes (e.g., SOX, IFRS) are also recommended [r17]. Provider policies today remain largely high-level and imply tiered access by payment or affiliation, highlighting a gap between aspirations and enforceable practice [r9]. Technically, continuous auditing overhead is better predicted by system-call rate than by coarse workload labels, and it exhibits non-linear behavior shaped by queuing and signaling parameters (often modeled via platform-specific tq and ts), multi-core contention, and I/O wait—complicating capacity planning for audited quantum–HPC services [r45, r57]. Technology choices matter: eBPF-based tracers can reduce overhead relative to ptrace, with one 16-thread I/O-bound RocksDB benchmark showing an overhead factor of 1.04× for Sysdig (eBPF) versus 1.71× for strace (ptrace)—a roughly 39% reduction—where the “overhead factor” denotes the ratio of runtime under auditing to baseline runtime [r110]. However, direct head-to-head measurements on >16-core systems are missing, and ptrace’s scaling remains unconfirmed, leaving high-core-count behavior and predictability unresolved [r71, r110]. These findings imply that equitable, auditable, and performant access to cloud quantum resources will require both governance standardization and careful engineering of low-overhead observability in multi-core, I/O-intensive production environments [r0, r17, r45, r57].

## Trajectory Sources

**Trajectory r0:** Practical quantum computing (QC) promises targeted super-polynomial speedups for specific tasks, with far-reaching implications if paired with robust engineering, standards, and governance; however, timelines, hardware scale, and realistic advantage remain debated, so impacts will phase in as device...

**Trajectory r7:** Economic analyses indicate that the high capital costs and specialized expertise required for quantum computing are likely to concentrate power and wealth in a few large corporations and nations, potentially increasing global inequality (boretti2024 pages 8-9, boretti2024technica...

**Trajectory r9:** The reviewed documents predominantly set forth high-level, normative guidance with implicit tiered access based on payment and institutional affiliation, while offering only general, non-enforceable declarations on ethical use, public auditing, and equitable access (gramegna2023 pages ...

**Trajectory r16:** Government-funded, open-access computing centers and public data/benchmark initiatives have demonstrated a higher capacity to foster broad participation from startups and academia while curbing market concentration compared to direct corporate subsidies or tax credits (aghion2009...

**Trajectory r17:** Our analysis supports the hypothesis that recent literature indeed proposes concrete operational mechanisms and regulatory frameworks for auditing and governing ethical use on public cloud quantum platforms, though some technical specifics remain at a conceptual stage (bharathan2024...

**Trajectory r45:** Preliminary evidence indicates that performance overhead in HPC continuous auditing and provenance tools is more closely linked to system call rate than to a broad I/O-bound/CPU-bound classification, though robust regression analyses across diverse workloads are still limited (sekar2024 p...

**Trajectory r57:** The evidence supports the hypothesis: system-call auditing tools do not exhibit overhead that grows linearly with system call count, but instead incur additional

non-linear performance penalties from factors such as multi-core contention and I/O wait times (sekar2024 pages 10-11, sekar202...

**Trajectory r71:** The current body of work does not provide sufficient direct evidence to conclusively support the hypothesis that under high core counts (>16 cores) and I/O-bound workloads, eBPF-based auditing tools exhibit at least 30% lower overhead and more predictable scaling than equivalent ptrace-based systems...

**Trajectory r110:** The current evidence does not conclusively support the hypothesis because, while eBPF-based tools show lower overhead than ptrace-based ones in a 16-thread I/O-bound workload, no study provides direct, quantitative results on >16-core systems or confirms super-linear scaling of ptrace overhead (este...

# Post-Quantum Cryptography Migration: Performance, Network, and Operational Risks

---

## Summary

Advancing quantum computers will force a global migration from RSA/ECC to post-quantum cryptography, and early hybrid deployments show measurable latency, bandwidth, and interoperability costs that reshape network and system design. At the same time, the earliest high-value applications of quantum computing are expected in quantum simulation for chemistry and materials, and in targeted hybrid AI-QC workflows accessed via cloud services, with significant societal benefits balanced by governance and equity risks.

## Background

Quantum computing promises targeted super-polynomial speedups for specific tasks and is progressing from noisy intermediate-scale devices toward fault-tolerant regimes. While timelines remain uncertain, the cryptographic consequences of Shor’s and Grover’s algorithms make the transition to post-quantum cryptography urgent, and early deployments reveal practical performance and integration trade-offs. In parallel, quantum simulation for biology, chemistry, and materials—likely delivered through hybrid cloud stacks—offers a path to tangible benefits, as do focused AI-QC integrations, necessitating rigorous benchmarks, standards, and responsible governance to ensure equitable societal impact.

## Results & Discussion

The most immediate, society-wide change from practically useful quantum computers is cryptographic disruption: Shor’s algorithm threatens RSA and elliptic-curve cryptosystems, and Grover’s algorithm reduces the effective security of symmetric ciphers, making “harvest-now, decrypt-later” particularly salient for long-lived data [r0, mathew2024]. Mitigation centers on migration to standardized post-quantum schemes (e.g., lattice-, code-, hash-, and multivariate-based), with hybrid modes recommended during transition and quantum key distribution evaluated as a complementary tool; however, attacker timelines and resources are

uncertain, and implementation security plus protocol migration are major operational risks [r0, gill2022, mavroeidis2018, eboseremen2023]. This frames the central finding: the PQC transition is not merely a cryptography upgrade—it is a cross-stack engineering effort with quantifiable performance, network, and operational risks whose management will determine near-term societal impact.

Controlled experiments consistently show that PQC-hybrid TLS introduces tail-latency and bandwidth penalties driven by larger handshake payloads that can trigger extra TCP round trips. At higher security levels, average time to send application data increases by approximately 25.9 ms and the 95th-percentile delay exceeds 10 ms compared with ECDHE-only baselines; the mechanism is message-size inflation causing additional network round trips [r33, giron2023]. Across early pilots and sector studies, handshake messages grow by roughly 2–8 KB, ML-DSA is about 5× slower to sign and 2× slower to verify than ECDSA, and certificate chains are 3–5× larger, driving capacity planning and PKI/storage redesigns [r8, erol2025a]. In finance and telecom, the bottleneck is not writing PQC code but handling latency, bandwidth, and legacy integration: a bank’s ML-KEM hybrid measured <5 ms added latency; telecom control-planes observed ~200 μs per-handshake overhead, acceptable within a 50 ms end-to-end budget but challenging at user-plane scale; at 100,000 transactions/s, a 2–8 KB handshake inflation implies 200–800 MB/s extra bandwidth in high-frequency trading contexts [r8, erol2025a]. Microcontroller-class data indicate ML-KEM/Kyber requires ~3× the memory of ECDH and 3.3 ms vs 0.9 ms on Cortex-M4, with hybrid transitions adding ~35–50% compute overhead and large installed device fleets likely needing upgrades [r8, 2025]. Simulations and modified library tests further suggest hybrid key exchanges sustain roughly 75–80% of classical new-connection throughput, but production-grade web-server benchmarks with

identical classical baselines, CPU attribution at saturation, and high-load network telemetry remain unreported [r56, r64, alnahawi2024, zheng2024]. Network-layer studies document fragmentation risks from large records yet lack quantitative buffer-occupancy and drop-rate measurements at thousands of handshakes/s; nevertheless, active queue management—particularly FQ-CoDel—is expected to mitigate queue build-up and retransmissions compared with DropTail under mixes that include large handshake packets [r31, r125, berger2503, blancoromero2025, kuhn2016, tobias2019].

Interoperability and operational readiness dominate early migrations. Certificate size inflation from ~1 KB to 5–15 KB, multi-year root CA updates, protocol message-size constraints, and the absence of runtime algorithm negotiation in a majority of systems (noted at 63% in some assessments) push operators toward compatibility-preserving designs: hybrid/composite certificates, negotiated hybrids, and PQ-enabled gateways, as seen in sector pilots by cloud providers [r8, erol2025a, alghamdi2025, erol2025]. Yet comprehensive, sector-specific playbooks for large PKI migrations in finance, telecom, and government are largely missing; most guidance remains high-level, with few detailed case reports, step-by-step procedures, or production telemetry [r5, fall2025, ogundipe2024]. Methodological gaps persist: no audited, high-load Apache/NGINX comparisons that report maximum sustainable new-handshakes/s for hybrid vs classical baselines under identical conditions; limited real-world tail-latency distributions; and scarce buffer and retransmission statistics under handshake stress, although the measurement blueprints and mitigations (e.g., FQ-CoDel) are well understood from network performance engineering [r31, r56, r125, zheng2024, berger2503, kuhn2016]. For practitioners, this implies prioritizing crypto-agility, certificate compression/caching, session resumption, hardware offload, and AQM deployment during staged rollouts while generating the missing production metrics.

Beyond cryptography, the path to everyday benefits runs through quantum simulation in biology, chemistry, and materials. There is strong consensus that accurate Hamiltonian simula-

tion and quantum phase estimation are the long-term gold standards, with variational algorithms (e.g., VQE) and hybrid embedding (QM/MM) as pragmatic near-term tools to target catalytic active sites, strongly correlated materials, and spectroscopic features; key bottlenecks include measurement overhead, state preparation, and resource scaling under realistic error models [r0, cao2019, baiardi2023, out-eiral2021, marchetti2022]. Disagreements remain over when “chemically relevant” instances become practical and whether NISQ can consistently beat strong classical baselines, especially given data-loading constraints (QRAM) and barren plateaus [r0, baiardi2023, bauer2020]. In AI research, hybrid quantum–classical workflows (e.g., QAOA, QSVMs, QGANs) are being explored for optimization, classification, and synthesis, while AI assists QC in calibration and error mitigation; however, real-world advantage is uncertain, and rigorous, task-specific comparisons, interpretable QML, and error-aware training at scale are still needed [r0, garciapineda2025, boretti2024, marchetti2022, baiardi2023]. For the public, access will likely come via cloud-style services embedded in hybrid stacks, with benefits in drug/material discovery, logistics, and climate modeling balanced by risks of privacy erosion, amplified surveillance, job displacement, and inequity—underscoring the need for standards, audits, and responsible research and innovation before mass deployment [r0, boretti2024, inglesant2016, troyer2403, wheatley2024]. In short, cryptographic migration is the earliest mass-impact change, while domain-specific quantum simulation and AI–QC integrations promise later gains if accompanied by rigorous benchmarking and governance.

## Trajectory Sources

**Trajectory r0:** Practical quantum computing (QC) promises targeted super-polynomial speedups for specific tasks, with far-reaching implications if paired with robust engineering, standards, and governance; however, timelines, hardware scale, and realistic advantage remain debated, so impacts will phase in as device...

**Trajectory r5:** The reviewed literature confirms that comprehensive, sector-specific playbooks for migrating large-scale PKIs to post-quantum cryptography are largely absent, with current guidance remaining predominantly high-level and theoretical (fall2025 pages 10-11, ogundipe2024postquantum...

**Trajectory r8:** The available case studies, pilots, and recent technical reports in finance and telecommunications support the hypothesis that performance overhead and legacy-system interoperability are the dominant practical obstacles in early PQC migrations, eclipsing core cryptographic implementation effort. (er...

**Trajectory r31:** The hypothesis is not supported by the cited literature: no study reports high-load (thousands of handshakes/s) TLS measurements of buffer occupancy or packet-drop/retransmission rates sufficient to demonstrate >50% buffer increases and >10% higher drops for PQC-hybrid versus ECC-only baselines.

**Trajectory r33:** Controlled experiments consistently indicate that, particularly at higher security levels, PQC-hybrid TLS increases the 95th-percentile connection establishment delay by more than 10 ms compared to classical ECDHE-only in simulated and load-test environments, although real-world production case stud...

**Trajectory r56:** The research hypothesis is not supported by the provided evidence: none of the cited works report a high-load (>1,000 new handshakes/sec) web-server (e.g., Apache/NGINX) benchmark that directly compares maximum sustainable new connections/sec of a PQC-hybrid TLS configuration against a classical-onl...

**Trajectory r64:** Current evidence from

controlled simulations indicates that hybrid PQC TLS handshakes can achieve approximately 75–80% of the throughput of classical ECDHE, but definitive system-level benchmarks on production-grade servers confirming comparable CPU utilization remain missing (alnahawi2024acomprehen...

**Trajectory r125:** High handshake payload sizes—representative of PQC-hybrid TLS—are expected to induce higher TCP retransmission rates and increased queue occupancy under high concurrent connection loads, but these adverse effects are significantly mitigated by deploying AQM mechanisms like FQ-CoDel compared to stand...