

Quantum-Safe Communication Architecture for Synchronized Optical Data Center Interconnects

Stan Mater

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.

stan.work@ku.edu

Euan McMurrough

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.

mcmurrough1987@buffalo.edu

Laoise McGrehan

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.

contactlaoise@oregonstate.edu

Abstract

This paper examines the architectural, infrastructural, and governance requirements for quantum-safe communication in synchronized optical data center interconnects. The central problem is not merely the replacement of vulnerable public-key primitives with quantum key distribution or post-quantum cryptography, but the integration of those mechanisms into a high-precision optical fabric in which time synchronization, wavelength switching, telemetry, and control traffic are tightly coupled. We argue that synchronization should be treated as a first-class security asset rather than a background service, because the integrity of key exchange, authentication, replay protection, and fault localization depends on temporal coherence across interconnected facilities. The paper proposes a layered reference architecture that separates quantum-safe key management, optical transport, synchronization, policy control, and sustainability governance while preserving interoperability across vendor domains. It analyzes the trade-offs between dedicated quantum channels, coexistence with classical dense wavelength-division multiplexed traffic, and post-quantum cryptographic fallback. Particular attention is given to hollow-core fiber and wavelength-switched intra-data-center interconnects, which reduce latency but also create new dependencies and attack surfaces. The discussion extends to deployment staging, standardization, fairness in multi-tenant environments, resilience under partial compromise, and the environmental cost of cryptographic redundancy. The conclusion is that quantum-safe optical data center interconnect design must be understood as a socio-technical problem in which protocol choice, physical-layer engineering, operational policy, and long-term infrastructure investment are mutually constitutive. Without coordinated governance, quantum-safe islands may emerge that are secure internally but fragile externally.

Keywords

quantum-safe communication; optical data center interconnects; time synchronization; quantum key distribution; post-quantum cryptography; hollow-core fiber; wavelength switching; infrastructure governance; sustainability; resilience.

1. Introduction

Data center interconnection is entering a period of simultaneous expansion and cryptographic transition. Hyperscale operators, financial exchanges, research networks, and sovereign cloud providers increasingly depend on optical links that must carry enormous volumes of traffic with deterministic latency and precise timing. At the same time, the long-term risk posed by quantum computing to widely deployed public-key cryptography has moved from speculative concern to infrastructure planning assumption. Shor’s algorithm established that sufficiently capable quantum computers could break factoring and discrete-logarithm systems, while Grover’s algorithm alters the effective security margins of symmetric primitives [1], [2]. For optical data center interconnects, this creates a dual challenge: the transport layer must become faster, more synchronized, and more flexible, while the security layer must become resistant to adversaries with quantum capabilities. The result is not a simple upgrade path but an architectural reconfiguration.

Quantum key distribution has long been proposed as a physically grounded response to key exhaustion and long-term secrecy risks [3], [4]. Its security model, however, depends on assumptions about device behavior, channel integrity, authentication, and operational discipline that are difficult to sustain in complex multi-tenant environments [5], [6], [7]. Post-quantum cryptography offers a software-oriented alternative that can be deployed across existing protocols, but it introduces new questions about implementation assurance, migration cost, and cryptographic agility [12], [13]. A quantum-safe communication architecture for synchronized optical data center interconnects must therefore combine physical-layer innovation with cryptographic diversity, operational governance, and careful timing design. The contribution of this paper is a system-level analysis of that combination. It does not propose a single protocol or hardware blueprint; instead, it identifies structural trade-offs and offers a reference orientation for researchers, operators, and policymakers.

2. Quantum-Safe Foundations and Threat Model

The threat model for optical data center interconnects is broader than the canonical store-now-decrypt-later scenario. An adversary may intercept optical signals, exploit misconfigured key management, manipulate synchronization messages, induce traffic-analysis patterns, or compromise a tenant’s virtualized control plane. Quantum key distribution can provide information-theoretic security under idealized conditions, but practical systems must contend with detector imperfections, finite-key effects, side channels, and denial-of-service pressure [6], [8]. Reviews of quantum cryptography emphasize that security proofs and deployment realities are not identical [5], [9]. Field demonstrations, including metropolitan and wide-area networks, have shown that quantum key distribution can operate over installed fiber, but they also reveal the operational burden of trusted nodes, calibration, and dedicated maintenance [10], [11]. These lessons matter for data center interconnects, where uptime requirements are extreme and where cryptographic failure can cascade into service outages.

Post-quantum cryptography complements rather than replaces quantum key distribution. Lattice-based, hash-based, code-based, and multivariate schemes offer different performance and assurance profiles, and standardization efforts have accelerated the transition from research to procurement [12], [13]. In an optical data center interconnect, post-quantum cryptography is likely to protect authentication, control-plane messaging, and fallback key establishment, while quantum key distribution may be reserved for high-value links or long-lived secrets. The architecture must therefore support cryptographic heterogeneity. A link may use quantum key distribution for bulk key material, post-quantum signatures for authentication, and classical symmetric encryption for data encapsulation. This heterogeneity

is not a weakness if it is governed explicitly; it becomes a weakness when key lifecycles, revocation, and audit trails are fragmented across domains. The threat model must also include insider risk, supply-chain compromise, and policy failure, because quantum-safe algorithms cannot compensate for weak operational governance [20].

3. Synchronization as a First-Class Architectural Resource

Time synchronization is often treated as a supporting service, yet in synchronized optical data center interconnects it is an architectural resource with security, performance, and economic consequences. Protocols such as the Network Time Protocol and IEEE 1588 provide mechanisms for distributing time, but their accuracy depends on network asymmetry, timestamping quality, and path stability [14], [15]. In optical fabrics using wavelength switching, synchronization messages may travel over different paths from data traffic, creating transient skew during reconfiguration. If synchronization is disrupted, key exchange windows may expire, replay protection may fail, and telemetry may become unreliable. An adversary that cannot break encryption may still degrade synchronization to force fallback to weaker modes or to cause service-level violations. Quantum-safe design must therefore treat synchronization as a protected control plane, not as best-effort metadata.

A synchronized optical interconnect can use time as a shared secret substrate. Precise timestamps can bind key identifiers to specific switching epochs, making replay and relay attacks more difficult. They can also support temporal isolation between tenants, allowing policy engines to verify that a key was generated and consumed within an authorized window. However, this creates a circular dependency: synchronization protects key management, and key management protects synchronization. The architecture must break that circle through redundant timing sources, authenticated time distribution, and hardware-rooted trust. It must also avoid over-centralizing time authority, because a single compromised grandmaster could become a systemic risk. The design implication is that synchronization should be distributed, verifiable, and auditable. It should be engineered with the same rigor as encryption, including independent monitoring, graceful degradation, and clear recovery procedures.

4. Layered Architecture and Key Management

A quantum-safe communication architecture for synchronized optical data center interconnects can be organized into several interacting layers. The physical layer includes fiber, amplifiers, transceivers, wavelength-selective switches, and quantum channels. The synchronization layer distributes precise time and frequency references across nodes. The key management layer negotiates, stores, rotates, and retires cryptographic material. The control layer enforces policy, routing, and fault recovery. The service layer exposes secure connectivity to tenants and applications. Finally, the governance layer defines standards, audits, liability, and sustainability constraints. This layering is analytical rather than rigid; in practice, cross-layer optimization is essential. For example, wavelength switching decisions affect synchronization stability, which affects key freshness, which affects the acceptable latency for rekeying. A purely layered model that ignores these dependencies will produce brittle systems.

Key management is the most politically and operationally sensitive layer. Quantum key distribution systems often generate keys continuously, but applications consume them at variable rates. Buffering, pooling, and prioritization are therefore necessary. In a multi-tenant data center, key material must be isolated by tenant, by service class, and by legal jurisdiction. Some tenants may require sovereign key custody, while others may accept operator-managed

keys with strong audit. The architecture must support policy-based key routing, including the ability to deny key use when a link's synchronization confidence falls below a threshold. Post-quantum cryptography can serve as a fallback when quantum key distribution is unavailable, but fallback should not be silent. Silent downgrade is a classic security failure mode because it hides degradation from operators and regulators. A quantum-safe architecture should make fallback visible, logged, and subject to policy approval.

5. Optical Transport, Wavelength Switching, and Hollow-Core Fiber

Optical interconnects for data centers have evolved from simple point-to-point links toward switched fabrics that can reconfigure capacity on demand [16]. Wavelength switching offers low latency and high throughput, but it also introduces transient impairments during reconfiguration. In quantum-safe systems, these transients matter because quantum channels are sensitive to loss, phase drift, and crosstalk. Coexistence of quantum and classical signals in the same fiber is attractive economically, but it requires careful power management, filtering, and wavelength planning. Space-division multiplexing and advanced fiber designs may increase capacity and isolation, though they also raise cost and operational complexity [18]. The broader trajectory of fiber-optic networking suggests that capacity growth will continue, but the limiting factors will increasingly be control, synchronization, and security rather than raw bandwidth alone [19].

Hollow-core fiber is particularly relevant because it can reduce latency and nonlinear impairment, which benefits both classical data and quantum signals. Recent work on low-latency wavelength-switched clock-synchronized intra-data-center interconnects enabled by hollow-core fiber demonstrates how physical-layer advances can be coupled with synchronization and switching [17]. That work is important not because it solves all quantum-safe problems, but because it shows that latency, clock distribution, and wavelength agility can be co-designed. In a quantum-safe architecture, hollow-core fiber may improve key generation rates by reducing loss and may improve synchronization by reducing propagation delay. However, it also changes failure modes: air-core structures may have different environmental sensitivities, splice losses, and long-term reliability profiles than conventional silica fiber. Deployment must therefore be staged, with telemetry that compares predicted and observed latency, loss, and key yield.

6. Deployment, Governance, and Standardization

Deployment of quantum-safe optical data center interconnects will be uneven. Hyperscale operators may build private quantum channels between selected sites, while smaller enterprises may rely on post-quantum cryptography delivered through cloud services. This unevenness creates interoperability risks. If one facility uses quantum key distribution and another uses only post-quantum cryptography, the interconnection point becomes a policy translation zone. Standards must define how key identifiers, synchronization confidence, and fallback events are signaled across domains. Governance must also address liability. If a quantum-safe link fails because of a synchronization attack, is the fiber operator, the equipment vendor, the key management provider, or the tenant responsible? Clear contracts and audit regimes are as important as protocol design. The ethics of information infrastructure further suggest that security should not be treated as a purely private good when critical services depend on it [20].

Standardization should avoid premature lock-in while ensuring baseline security. Open interfaces for key management, time distribution, and telemetry can allow competition and

innovation. At the same time, regulators may need to mandate minimum cryptographic agility, disclosure of fallback events, and retention of security logs. Data center interconnects often cross jurisdictional boundaries, so policy harmonization is difficult. A regional approach may emerge, with stricter rules for financial and governmental traffic and lighter rules for best-effort commercial traffic. The architectural implication is that policy must be encoded in the control plane, not merely written in contracts. A link that cannot enforce a jurisdictional key custody rule should not be allowed to carry regulated traffic. This requires a tight integration between identity, policy, and optical path selection.

7. Robustness, Fairness, and Sustainability

Robustness in quantum-safe optical interconnects requires redundancy across cryptographic, optical, and synchronization domains. A link should be able to continue operating, perhaps at reduced assurance, when one domain is degraded. For example, if quantum key distribution fails, post-quantum cryptography may sustain authentication while operators investigate. If synchronization degrades, the system may switch to a holdover clock and restrict high-assurance key use. If an optical path fails, wavelength switching may reroute traffic, but the new path must be revalidated for latency, loss, and synchronization. These transitions should be rehearsed and automated, because manual recovery under attack is slow and error-prone. Robustness also means resisting traffic analysis. Quantum-safe encryption does not hide metadata such as flow size, timing, and destination. In synchronized fabrics, timing patterns can leak information about application behavior. Padding, shaping, and tenant isolation may be necessary, though they consume bandwidth and energy.

Fairness is a less obvious but important concern. Quantum-safe infrastructure is expensive, and early deployment may favor large operators and high-value tenants. Smaller providers may be left with weaker fallback mechanisms, creating a two-tier security ecosystem. Multi-tenant data centers must ensure that one tenant's quantum key distribution usage does not starve another's synchronization or bandwidth. Fairness also applies to maintenance windows and fault recovery. A tenant with a large share of traffic may receive priority for rekeying, while smaller tenants experience delays. Policy should define fairness metrics and enforce them through scheduling and key allocation. Sustainability adds another dimension. Quantum key distribution requires dedicated lasers, detectors, cooling, and sometimes trusted nodes, all of which consume energy. Post-quantum cryptography increases computational load, especially during migration when both classical and quantum-safe algorithms may run in parallel. The environmental cost of cryptographic redundancy must be measured against the risk reduction it provides. In some cases, a well-managed post-quantum deployment may be more sustainable than a poorly utilized quantum channel.

8. Case Illustrations and Comparative Perspectives

Consider a financial data center interconnect between two trading facilities. Latency is critical, and synchronization is tightly regulated. A quantum-safe architecture might use hollow-core fiber for low-latency transport and wavelength switching for rapid reconfiguration [17]. Quantum key distribution could protect long-lived audit logs, while post-quantum cryptography protects session authentication. Synchronization would be distributed through redundant grandmasters, with cryptographic authentication of time messages. If a switching event causes skew beyond a defined threshold, the control plane would restrict new key generation until timing confidence is restored. This case illustrates that quantum safety is not only about algorithms; it is about maintaining a chain of trust across time, light, and policy.

A contrasting case is a multi-tenant cloud region serving many enterprises. Here, tenant isolation and cost sharing dominate. Quantum key distribution may be offered as a premium service, while post-quantum cryptography is the default. The architecture must prevent cross-tenant timing side channels and ensure that premium key material is not observable through metadata. Governance must define what happens when a tenant requests sovereign key custody that conflicts with the operator's key management design. A third case is a research network connecting universities and national laboratories. It may prioritize experimentation, openness, and interoperability over commercial latency. These cases show that there is no single optimal quantum-safe architecture. The design space is shaped by threat tolerance, regulatory context, traffic economics, and sustainability constraints.

9. Conclusion

Quantum-safe communication for synchronized optical data center interconnects is a systems problem that cannot be solved by selecting a single cryptographic primitive. The architecture must integrate quantum key distribution, post-quantum cryptography, precise synchronization, optical switching, and policy governance into a coherent operational fabric. Synchronization should be treated as a protected resource because it underpins key freshness, replay protection, and fault recovery. Optical innovations such as hollow-core fiber and wavelength switching can improve performance, but they also create new dependencies that must be modeled and monitored. Deployment should be staged, with visible fallback, cryptographic agility, and clear audit trails. Governance must address liability, fairness, and sustainability, because otherwise quantum-safe islands may form that are secure internally but fragile at their boundaries. The most durable approach is one that treats security, timing, and optical transport as co-designed layers, supported by standards that reward transparency and interoperability. Future research should develop quantitative models for synchronization-confidence-aware key management, cross-domain policy enforcement, and lifecycle assessment of quantum-safe infrastructure. Without such work, the transition to quantum-safe optical interconnects may be technically impressive but operationally incomplete.

References

1. Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5), 1484–1509.
2. Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, 212–219.
3. Bennett, C. H., & Brassard, G. (1984). Quantum cryptography: Public key distribution and coin tossing. *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, 175–179.
4. Ekert, A. K. (1991). Quantum cryptography based on Bell's theorem. *Physical Review Letters*, 67(6), 661–663.
5. Gisin, N., Ribordy, G., Tittel, W., & Zbinden, H. (2002). Quantum cryptography. *Reviews of Modern Physics*, 74(1), 145–195.
6. Scarani, V., Bechmann-Pasquinucci, H., Cerf, N. J., Dušek, M., Lütkenhaus, N., & Peev, M. (2009). The security of practical quantum key distribution. *Reviews of Modern Physics*, 81(3), 1301–1350.

7. Lo, H.-K., Curty, M., & Tamaki, K. (2014). Secure quantum key distribution. *Nature Photonics*, 8(8), 595–604.
8. Diamanti, E., Lo, H.-K., Qi, B., & Yuan, Z. (2016). Practical challenges in quantum key distribution. *npj Quantum Information*, 2, 16025.
9. Pirandola, S., Andersen, U. L., Banchi, L., Berta, M., Bunandar, D., Colbeck, R., Englund, D., Gehring, T., Lupo, C., Ottaviani, C., Pereira, J. L., Razavi, M., Shamsul Shaari, J., Tomamichel, M., Usenko, V. C., Vallone, G., Villoresi, P., & Wallden, P. (2020). Advances in quantum cryptography. *Advances in Optics and Photonics*, 12(4), 1012–1236.
10. Chen, J.-P., Zhang, C., Liu, Y., Jiang, C., Zhang, W., Hu, X.-L., Guan, J.-Y., Yu, Z.-W., Xu, H., Lin, J., Li, M.-J., Chen, H., Li, H., You, L., Wang, Z., Wang, X.-B., Zhang, Q., & Pan, J.-W. (2021). Twin-field quantum key distribution over a 511 km optical fibre linking two distant metropolitan areas. *Nature Photonics*, 15(8), 570–575.
11. Wang, S., Chen, W., Yin, Z.-Q., Li, H.-W., He, D.-Y., Li, Y.-H., Zhou, Z., Song, X.-T., Li, F.-Y., Wang, D., Chen, H., Han, Y.-G., Huang, J.-Z., Guo, J.-F., Hao, P.-L., Li, M., Zhang, C.-M., Liu, D., ... Pan, J.-W. (2014). Field and long-term demonstration of a wide area quantum key distribution network. *Optics Express*, 22(18), 21739–21756.
12. Bernstein, D. J., & Lange, T. (2017). Post-quantum cryptography. *Nature*, 549(7671), 188–194.
13. Chen, L., Jordan, S., Liu, Y.-K., Moody, D., Peralta, R., Perlner, R., & Smith-Tone, D. (2016). Report on post-quantum cryptography (NISTIR 8105). National Institute of Standards and Technology.
14. Mills, D. L. (1991). Internet time synchronization: The network time protocol. *IEEE Transactions on Communications*, 39(10), 1482–1493.
15. Eidson, J. C., & Kang, L. (2016). IEEE 1588 standard for a precision clock synchronization protocol for networked measurement and control systems. *Proceedings of the IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication*, 1–6.
16. Kachris, C., & Tomkos, I. (2012). A survey on optical interconnects for data centers. *IEEE Communications Surveys & Tutorials*, 14(4), 1021–1036.
17. Zhou, Zichuan, et al. "Low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber." *Optics Express* 31.15 (2023): 24739-24748.
18. Puttnam, B. J., Rademacher, G., & Luís, R. S. (2021). Space-division multiplexing for optical fiber communications. *Optica*, 8(9), 1186–1203.
19. Winzer, P. J., Neilson, D. T., & Chraplyvy, A. R. (2018). Fiber-optic transmission and networking: The previous 20 and the next 20 years. *Optics Express*, 26(18), 24190–24239.
20. Floridi, L. (2013). *The ethics of information*. Oxford University Press.