Mind the Gap: Policy vs Reality in Post-Quantum TLS Deployment

Nimesha Wickramasinghe
n.wickramasinghe@unsw.edu.au
University of New South Wales
Sydney, NSW, Australia

Frank Li
frankli@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

Sanjay Jha
sanjay.jha@unsw.edu.au
University of New South Wales
Sydney, NSW, Australia

Arash Shaghaghi
a.shaghaghi@unsw.edu.au
University of New South Wales
Sydney, NSW, Australia

Abstract

Post-quantum cryptography (PQC) has evolved from a long-term planning concern into an operational priority. Following NIST's standardization of PQC algorithms, governments and standards bodies published transition roadmaps outlining migration timelines, priority sectors, and deployment strategies. However, our survey of these policies reveals substantial divergence in technical prescriptions and urgency. It remains unclear how widely PQC has been adopted in practice and how policy differences translate into observable deployment outcomes.

To quantify trends in public-facing post-quantum TLS (PQ-TLS) adoption, we conduct three measurement rounds between July 2025 and March 2026, covering one million public HTTPS endpoints. Across these rounds, we establish more than two billion TLS 1.3 handshakes from 11 globally distributed vantage points to characterize cryptographic negotiation behavior. We observe strong convergence in hybrid key exchange, with every observed post-quantum negotiation selecting `X25519MLKEM768`, but find no verified deployment of post-quantum signature algorithms for authentication. Hybrid key-exchange adoption is highly concentrated among endpoints attributed to managed infrastructure providers. Country and sector comparisons show limited correspondence between early deployment patterns and published transition timelines. Contrary to early experimental studies suggesting measurable overhead, our experiments show no meaningful latency increase for hybrid key exchange in Internet settings. We further observe that PQC adoption frequently coexists with legacy TLS configurations.

Together, these findings highlight a gap between policy expectations and early deployment reality, and provide empirical insight to inform more grounded PQ-TLS transition.

CCS Concepts

• **Security and privacy** → **Network security**; **Cryptography**; • **Networks** → **Network protocols**; **Network measurement**.

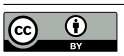

IMC '26, Karlsruhe, Germany

ACM ISBN 979-8-4007-2327-8/2026/10
https://doi.org/10.1145/3777912.3839797

Keywords

Post-Quantum Cryptography, Post-Quantum TLS, Internet Measurement, TLS Deployment, PQC Migration

ACM Reference Format:
Nimesha Wickramasinghe, Frank Li, Sanjay Jha, and Arash Shaghaghi. 2026. Mind the Gap: Policy vs Reality in Post-Quantum TLS Deployment. In *Proceedings of the 2026 ACM Internet Measurement Conference (IMC '26), October 12–16, 2026, Karlsruhe, Germany*. ACM, New York, NY, USA, 20 pages. https://doi.org/10.1145/3777912.3839797

1 Introduction

The continued progress of quantum computing fundamentally challenges the security of long-standing public-key cryptography [73]. In response, post-quantum cryptography (PQC) seeks to replace these vulnerable schemes with algorithms that remain secure against cryptographically relevant quantum computers (CRQCs) [2].

The transition to PQC has been shaped by a global standardization effort led by the National Institute of Standards and Technology (NIST). After eight years of an open, multi-round evaluation process, NIST finalized its first set of PQC algorithms in August 2024 [57–59]. These algorithms now form the basis for protocol development within standards bodies such as, the IETF [11, 64, 76] and ISO [46], are adopted by widely used cryptographic libraries [3], and are increasingly supported by major infrastructure providers [4, 72]. Following these standardization efforts, governments and standards bodies worldwide have published transition roadmaps and policy guidelines to promote awareness and adoption of PQC (e.g., [2, 8, 34]).

Across these policies, however, we observe a fragmented landscape, with substantial divergence in technical prescriptions, urgency, and deployment guidance. At the same time, the extent of PQC adoption in practice and its alignment with differing policy commitments remain largely unknown.

To address this gap, we focus on TLS as deployed in HTTPS, the secure form of HTTP, which carries a dominant share of application-layer traffic on the Internet [30]. TLS exposes cryptographic negotiation behavior directly to clients, making deployment decisions observable at Internet scale. Post-Quantum TLS (PQ-TLS) in HTTPS therefore provides a practical vantage point for examining how standardized algorithms are adopted and deployed in real-world systems.

In this paper, we present the first Internet-scale measurement study of PQ-TLS adoption following NIST's standardization of PQC

algorithms. Our study comprises three measurement rounds conducted in July 2025, November 2025, and March 2026, with approximately four months between each round. We employ a custom TLS 1.3 measurement client to conduct controlled handshake probing from 11 globally distributed vantage points. In total, we analyze cryptographic negotiation outcomes of more than 2 billion TLS handshakes across one million public HTTPS endpoints, enabling us to characterize early changes in PQ-TLS adoption.

Our measurements reveal that, despite diverse policy guidance and a broad standardized design space, operational deployment is strongly concentrated on a single hybrid key-exchange configuration, `X25519MLKEM768`. Adoption is also highly concentrated among endpoints attributed to managed infrastructure providers such as CDNs, with Cloudflare and Fastly accounting for nearly 70% of observed hybrid deployments. In contrast, we find no verified deployment of the post-quantum signature schemes for authentication. National timelines and sectoral priorities show limited correspondence with early deployment patterns. Contrary to expectations from controlled experiments, our matched measurements show no meaningful latency increase for `X25519MLKEM768` relative to `X25519` in Internet settings. We also find that PQC adoption frequently coexists with support for legacy TLS configurations.

Overall, these findings highlight a disconnect between policy expectations and early deployment realities, while offering empirical evidence to inform more grounded PQC transition strategies.

2 Background

This section summarizes the cryptographic foundations of TLS, the quantum threat to its public-key components, and the standardization efforts that motivate the transition to PQC.

2.1 Public-Key Cryptography in TLS

Transport Layer Security (TLS) combines public-key and symmetric-key cryptography to provide authentication, integrity, and confidentiality for Internet communication. Public-key cryptography is used during the handshake for key establishment and authentication, while symmetric cryptography protects application data once a secure channel is established. Consequently, the security of the resulting session keys depends directly on the public-key mechanisms used during the handshake.

In TLS, key establishment is performed using ephemeral Diffie–Hellman (DH) exchanges. In practice, this is predominantly realized using elliptic-curve Diffie–Hellman (ECDH), most commonly over X25519 and, less frequently, X448. The exchange produces a shared secret from which symmetric session keys are derived. The use of ephemeral keys provides forward secrecy under the assumption that the underlying elliptic-curve problem remains hard to solve [66].

Authentication in TLS is provided through digital signatures embedded in X.509 certificates. Today, the dominant signature algorithms are ECDSA over secp256r1 and secp384r1, and RSA using probabilistic signature schemes such as RSA-PSS. These signatures allow clients to authenticate servers and verify possession of the corresponding private keys [66].

2.2 CRQC and Public-Key Risk

A CRQC is a quantum system capable of executing Shor's algorithm [73] at a scale sufficient to break widely deployed public-key cryptographic schemes. Such a capability would compromise both ECDH-based key exchange (e.g., X25519) and classical digital signature algorithms (e.g., ECDSA and RSA).

This threat enables two classes of attacks against TLS. First, an adversary may record encrypted TLS traffic and later decrypt it by recovering session keys once the key establishment mechanism is broken, commonly referred to as a *Harvest Now, Decrypt Later (HNDL)* attack. Second, an adversary may forge digital signatures, enabling server impersonation or man-in-the-middle attacks during connection establishment [54].

By contrast, symmetric cryptography and hash functions are not directly broken by Shor's algorithm and generally require only parameter changes. As a result, the quantum threat to TLS is concentrated on the public-key primitives used during the handshake.

2.3 NIST's PQC Standardization

To address these risks, in August 2024, NIST finalized its first set of PQC standards, publishing FIPS 203 [58] for key encapsulation, FIPS 204 [59] for digital signatures, and FIPS 205 [57] for stateless hash-based signatures.

FIPS 203 standardizes ML-KEM, a lattice-based key encapsulation mechanism intended to serve as a PQ alternative to the key agreement functionality of classical Diffie–Hellman. FIPS 204 and FIPS 205 standardize ML-DSA and SLH-DSA for authentication. Together, these standards define PQ replacements for public-key primitives used by protocols, such as TLS.

2.4 International PQC Standardization Efforts

Although NIST's PQC standards form the primary interoperability baseline for Internet protocols, several national and regional standardization bodies have also proposed or evaluated alternative PQ algorithms. For example, the German Federal Office for Information Security (BSI) has introduced key establishment mechanisms such as FrodoKEM and Classic McEliece [34]. Similarly, the State Cryptography Administration (SCA) of China has supported domestic candidates including LAC, a lattice-based key establishment mechanism, as well as signature schemes such as AAGL and PQC-Sign [18]. In South Korea, the Telecommunications Technology Association (TTA) has proposed SOLMAE [67], a lattice-based key establishment mechanism designed for PQ key exchange.

2.5 Integration of PQC into TLS

Despite these parallel efforts, NIST-selected algorithms form the principal basis for current PQC integration into TLS. NIST standardizes the underlying cryptographic primitives, while the IETF specifies how they are negotiated and used within the protocol.

For TLS 1.3, RFC 10024 standardizes three hybrid key-exchange groups: `X25519MLKEM768`, `SecP256r1MLKEM768`, and `SecP384r1MLKEM1024` [50]. As of August 2026, specifications for standalone ML-KEM groups, ML-DSA authentication, and broader recommendations for PQC deployment in TLS-based applications remain within the IETF publication process [24, 42, 65].

Implementation support has developed alongside standardization and protocol specifications. NIST-selected mechanisms are supported by major infrastructure providers, including Cloudflare [63] and Google [37], as well as contemporary cryptographic libraries [3]. ISO/IEC has also initiated related standardization activities based on NIST-selected algorithms [46]. Together, protocol specifications and implementation support provide an interoperability foundation for early PQ-TLS deployment.

2.6 PQC Deployment Urgency

Although both key establishment and authentication are vulnerable to quantum attacks, their mitigation differs in urgency and deployment feasibility. Compromise of key establishment enables retrospective decryption of recorded traffic, meaning that each quantum-vulnerable handshake permanently increases future confidentiality risk. By contrast, signature forgery attacks require active exploitation with a CRQC at the time of connection and do not retroactively compromise previously captured sessions.

For this reason, PQ key establishment has emerged as the primary near-term focus of PQC transition efforts for TLS.

3 Related Work

Internet-scale measurement has long provided an empirical basis for understanding how TLS evolves on the public Internet. Early work mapped the SSL/TLS ecosystem using active and passive measurements, revealing concentration and misconfiguration in the Web PKI [43]. Subsequent studies provided a longitudinal view of the HTTPS certificate ecosystem [27], examined the evolution of HTTPS deployment at scale [5], and showed that weak cryptographic keys remained widespread in network devices years after their disclosure [41]. More recent work has investigated long-term TLS adoption [49], the rollout and centralization of TLS 1.3 [44, 51], and geographic variation in negotiated TLS behavior [52]. Our work builds on these measurement methodologies and applies them to the emerging transition to PQ-TLS in HTTPS.

Research on PQC readiness spans several complementary perspectives. Experimental deployments such as CECPQ1 in 2016 and CECPQ2 in 2019 demonstrated the feasibility of hybrid PQ key exchange before standardization [36, 38]. Subsequent work has examined PQC support across cryptographic libraries [3] and documented deployment within individual provider ecosystems [21]. Controlled studies have also evaluated PQ-TLS performance in simulated environments and testbeds [68, 74, 75, 80]. These studies report that lattice-based key exchange increases handshake message size but can incur modest computational and latency costs in their evaluated settings. Kempf et al. similarly report limited handshake overhead for several PQ mechanisms in QUIC, while larger signature constructions introduce greater costs [48]. Their protocols, algorithms, implementations, and experimental environments differ from our public HTTPS-over-TCP setting. Our measurements complement these controlled studies by evaluating TLS handshake performance across operational Internet endpoints under real-world network conditions.

Operational telemetry provides another perspective on PQC deployment. For example, Cloudflare Radar reports the request-weighted share of HTTPS connections to Cloudflare that negotiate PQ encryption and uses daily scans to test `X25519MLKEM768` support among Cloudflare customer origins [22, 63]. Our study instead uses a fixed population of public domains and external measurements spanning multiple infrastructure providers and geographic vantage points. This design allows us to distinguish the group negotiated under a consistent client configuration from explicit algorithm support, providing a complementary view based on a different target population, denominator, and measurement approach.

Overall, our contribution is a three-round, Internet-scale measurement of PQ-TLS deployment across a fixed population of public HTTPS endpoints. The study combines cross-provider negotiation and support measurements with infrastructure attribution, probable TLS-management classification, and country, sector, security, and latency analyses. This allows us to characterize not only whether PQC is deployed, but also where deployment is concentrated and how it varies across the measured ecosystem.

4 Research Questions

To contextualize our measurement study in comparable policy contexts, we survey publicly available national and regional PQC transition documents.

We construct our policy set by selecting countries with demonstrated engagement in PQC, including participants in NIST's standardization process [16, 29], and major economic regions [39]. For each country or region, we collect published policies through systematic web-based searches and retain documents from authoritative sources, including government publications and national standards bodies. Full selection methodology and inclusion criteria are provided in Appendix A.

Table 1 summarizes the selected policies and the dimensions most relevant to PQ-TLS deployment, including recommended algorithms, guidance on hybrid constructions, sectoral scope, and transition timelines. While not exhaustive, our survey of eight countries captures a cross-section of PQC transition approaches whose differing requirements may place varying demands on protocol implementers, infrastructure providers, and service operators.

To assess how these policy expectations translate into observable deployment behavior, and whether they converge or diverge at the protocol level, we formulate a set of targeted research questions aligned with each policy dimension.

4.1 Algorithm Recommendations

Across the surveyed policies, there is clear convergence on ML-KEM as the primary PQ replacement for classical key-exchange in TLS. However, this consensus masks meaningful divergence in deployment guidance. France and Germany additionally reference alternative KEMs such as Frodo-KEM or Classic McEliece, while other jurisdictions focus exclusively on ML-KEM. Parameter recommendations also vary: Australia, India, and the United States of America (USA) explicitly prefer higher security levels (e.g., ML-KEM-1024), Canada permits a broader range of parameters, and France and Germany leave parameter selection largely unspecified.

Guidance on PQ signature algorithms is less prescriptive. Although ML-DSA and SLH-DSA are commonly referenced, some policies additionally mention XMSS or LMS.

Table 1: Comparative summary of national and regional PQC transition policies relevant to TLS

Country / Region	Recommended PQC Algorithms		Hybrid Stance	Requirement & Sector	PQC Transition Timeline		Classical Crypto. Timeline	
	KEM	SA			Complete by	Default by	Depreciate by	Disallow by
Australia [8, 9]	ML-KEM - 768/1024 Prefer 1024	ML-DSA - 65/87 65 only till 2030	Allowed Not recommended	● ■ △	2030	2030	2030	2030
Canada [32, 33]	ML-KEM - 512/768/1024	ML-DSA - 44/65/87 SLH-DSA	NS	○ □ △	2031	2035	2031	2035
Europe [23]	Promote international cooperation	Promote international cooperation	Recommended	Advisory for member states	2030	2035	2035	NS
France [7]	Kyber (ML-KEM) Frodo-KEM	Dilithium (ML-DSA) FN-DSA, XMSS, LMS	Recommended	○ □ △	2030	NS	NS	NS
Germany [34, 35]	ML-KEM, Frodo-KEM Classic McEliece	ML-DSA - 65/87 SLH-DSA, XMSS, LMS	Recommended	○ □ △	2030	NS	NS	NS
India [17, 56]	ML-KEM - 768/1024 Prefer 1024	ML-DSA - 65/87 SLH-DSA	Recommended	○ □ △	NS	2035	NS	NS
United Kingdom [14, 15]	ML-KEM - 768	ML-DSA - 65 SLH-DSA, XMSS, LMS	Transitional	○ □ △	2035	2035	NS	NS
USA [2, 45]	ML-KEM - 1024	ML-DSA - 87	Allowed Not recommended	● ■ △	2035	NS	2030	2035

● / ○ **(circle)**: Government / National Security; ■ / □ **(square)**: Critical / high-risk / long-lived systems; ▲ / △ **(triangle)**: Industry / organisations;
Filled (● ■ ▲): Mandatory; **Hollow** (○ □ △): Advisory; **NS**: Not Specified

Hybrid constructions, which combine a classical cryptographic algorithm (e.g., X25519) with a PQ algorithm (e.g., ML-KEM), are widely discussed but positioned differently across jurisdictions. Australia and the USA permit hybrid deployment without explicitly endorsing it, the United Kingdom frames hybrid use as transitional, and other policies recommend hybrid constructions without constraints.

Taken together, the diversity in algorithms, security levels, and hybrid constructions motivates an empirical examination of how PQ-TLS is deployed in practice.

> **RQ1 – Configuration Landscape:** *Which standardized PQ-TLS configurations are deployed in practice, and how does the observed configuration landscape compare to the diversity articulated in national transition policies?*

4.2 Infrastructure Concentration

National PQC transition policies are framed with limited explicit consideration of the infrastructure platforms. In practice, however, a substantial fraction of Internet traffic is served by large-scale CDNs and managed hosting providers [40, 44, 51, 82].

Deployment decisions made at the infrastructure layer can therefore exert disproportionate influence over observed adoption patterns, security posture, and performance characteristics. Distinguishing between infrastructure-driven rollout and independent operator migration is essential for understanding whether observed PQC adoption reflects systemic readiness or infrastructure-level defaults.

> **RQ2 - Deployment Agency:** *How much of observed PQ-TLS adoption is driven by infrastructure-managed deployment versus potentially owner-managed migration?*

4.3 Policy Scope and Timelines

The surveyed policies differ in both scope and transition timelines, defining which sectors should prioritize PQC adoption and when migration should be completed. Some jurisdictions, including Australia and the USA, explicitly mandate migration for government systems, national security, and critical infrastructure, while others provide broader advisory guidance spanning public and private sectors. These differences reflect varying national strategic priorities, particularly for systems handling sensitive or long-lived data.

Timelines likewise diverge. As summarized in Table 1, several jurisdictions, including Australia, Canada, and the European Union, target completion of PQC migration by 2030–2031, whereas others extend transition to 2035.

Such variation creates measurable expectations: earlier deadlines and narrower sectoral mandates could plausibly correspond to higher present-day readiness.

Internet-facing services span diverse sectors, including finance, healthcare, telecommunications, government services, and consumer platforms, each with differing regulatory exposure and constraints. Examining PQ-TLS deployment across countries and sectors therefore enables an empirical assessment of whether policy-defined scope and timelines manifest in observable deployment differences.

> **RQ3 – Policy Alignment:** *To what extent do national transition timelines and sectoral prioritizations correspond to observable differences in PQ-TLS deployment?*

4.4 Operational Characteristics

Early ecosystem-specific measurements have suggested that PQ-TLS introduces measurable performance overhead during deployment. For example, Google reported that the larger key share size of Kyber resulted in a 4% median increase in TLS handshake latency in Chrome on desktop [13]. Such findings highlight operational costs associated with larger key sizes and handshake messages in PQ constructions [61].

National transition policies, however, provide limited discussion of deployment constraints such as performance overhead, middlebox compatibility, or handshake reliability at Internet scale. Implicit

in many policy roadmaps is the assumption that standardized PQ algorithms can be integrated into operational systems without prohibitive impact.

For Internet-facing security protocols, operational viability encompasses not only cryptographic computation cost but also handshake latency, message size expansion, and negotiation reliability across heterogeneous infrastructures. Empirical measurement of these characteristics is therefore essential to ensure that policy recommendations are compatible with real-world deployment constraints.

> **RQ4 – Operational Viability:** *Does PQ-TLS introduce measurable performance, reliability, or compatibility constraints in operational Internet deployments?*

4.5 Security Co-evolution

While the surveyed policies provide guidance on *which* cryptographic algorithms to adopt and *when* to transition, they focus primarily on primitive-level replacement. *How* these algorithms should be integrated into Internet-facing protocols such as TLS is typically left implicit or deferred to existing security baselines and operational practices, which may evolve independently of PQ transition planning (e.g., [77, 81]).

However, prior Internet-scale measurement studies have consistently shown that protocol-level security depends not only on cryptographic primitives but also on configuration hygiene, including supported protocol versions, cipher suites, certificate management practices, and exposure to legacy vulnerabilities [1, 12, 53].

As PQ-TLS deployment progresses, an open question is whether PQ adoption occurs alongside broader security modernization, or whether PQ mechanisms are layered onto existing configurations without corresponding improvements in overall protocol security. Empirical evaluation of this relationship is necessary to understand whether observable PQ adoption reflects holistic security advancement or isolated cryptographic substitution.

> **RQ5 – Security Co-evolution:** *Does PQ-TLS adoption coincide with broader improvements in protocol-level security posture?*

5 Methodology

This section outlines the domain selection process (§ 5.1), measurement infrastructure (§ 5.2), TLS probing techniques (§ 5.3), TLS management attribution (§ 5.4), and the three measurement rounds (§ 5.5) used to analyze PQ-TLS adoption.

5.1 Target Selection and Domain Coverage

Our measurements are based on DomCop's Top 10M Domains list [1], which is primarily derived from Open PageRank data. The open-source initiative, Open PageRank, estimates the relative importance of domains using web graph data from Common Crawl [25].

For configuration landscape (RQ1), deployment agency (RQ2), operational viability (RQ4), and security co-evolution (RQ5) analyses, we focus on the Top 1M domains. This subset captures a broad cross-section of widely deployed Internet services while maintaining feasibility for repeated measurements.

[1] https://www.domcop.com/top-10-million-domains

To support country level policy alignment analysis (RQ3), we additionally draw from the broader Top 10M corpus to construct per-country samples, including Top 10K nationally relevant domains and Top 1K government-operated domains.

We select DomCop for two methodological reasons. First, many country-specific government and public-sector domains do not appear in global Top 1M lists (e.g., Tranco [62]), despite being operationally significant within their jurisdictions. Access to a larger corpus therefore enables the construction of more representative national samples. Although the Google Chrome UX Report (CrUX) [19] provides broader coverage, it reports domains in coarse rank-order magnitude buckets (i.e., Top 1K, 10K, 100K, 1M, and >1M) rather than exact rankings, limiting the ability to systematically select top domains [70]. Second, using a single, consistent ranking source across both global and country level measurements ensures internal consistency across our analyses.

To assess the robustness of our findings, we replicate PQ-TLS adoption measurements using widely used top 1M domain rankings, including Tranco and CrUX. The adoption rates are comparable and are reported in Appendix B.

5.2 Regional Scanning Infrastructure

Content delivery networks, geo-replicated cloud services, and regulatory environments can result in region-dependent TLS behavior [52]. To capture this diversity, we conducted measurements from 11 geographically distributed vantage points: North America West (USA, California), North America Central (Canada, Montreal), North America East (USA, Virginia), South America (Brazil, São Paulo), Africa (South Africa, Cape Town), United Kingdom (England, London), Europe (Germany, Berlin), Middle East (United Arab Emirates, Dubai), South Asia (India, Mumbai), East Asia (Japan, Tokyo), and Oceania (Australia, Sydney).

All vantage points were hosted on cloud infrastructure and operated independently, allowing us to observe region-specific negotiation outcomes.

To ensure consistency and mitigate measurement bias, all scans were conducted from dedicated cloud instances provisioned with identical CPU, memory, and network configurations. Scanning tools were deployed from identical precompiled binaries to avoid client-side fingerprinting discrepancies and ensure uniform TLS behavior across vantage points. Measurements were executed within shared time windows to reduce temporal variance and limit the impact of configuration changes during the measurement period.

5.2.1 DNS Resolution and Regional Targeting. Many domains are served by CDNs or geo-replicated infrastructures that return different IP addresses depending on client location. To account for this behavior, we perform DNS resolution locally at each vantage point prior to scanning. For each domain, we resolve all available A and AAAA records and apply a latency-based IP selection strategy following prior work [51]. Specifically, we measure the round-trip time to each candidate IP using TCP connection attempts on port 443 and select the address with the lowest observed latency among responsive endpoints. We then perform TLS scans against that endpoint, using the domain name in the SNI extension. This approach ensures that our measurements reflect the endpoint most representative of the user experience in each region.

5.3 Scanning Tools and Measurements

We develop a custom TLS scanning client built on OpenSSL 3.5.0 (released in April 2025), which provides native support for NIST-standardized PQC algorithms. The client uses TLS 1.3 exclusively and otherwise retains OpenSSL's default group configuration. Under this setting, the ClientHello lists the hybrid group `X25519MLKEM768` first in the `supported-groups` extension and includes two predicted `KeyShareEntry` values, one for `X25519MLKEM768` and another for `X25519`. The hybrid entry concatenates the client's ML-KEM-768 encapsulation key with its ephemeral X25519 public share. At the time of our measurements, prioritizing `X25519MLKEM768` was broadly consistent with several widely deployed contemporary implementations, such as Chromium/BoringSSL [20], Firefox/NSS [55], Go's crypto/tls [79], rustls [69], and AWS s2n-tls [6].

For each target domain, we establish two reference TLS 1.3 handshakes. The first, *default-client handshake*, leaves OpenSSL 3.5.0's default supported-group and signature-scheme configurations unchanged. We record the negotiated key-exchange group and signature scheme as *default_KE* and *default_SA*, respectively. The second, *classical-only handshake*, restricts the advertised groups and signature schemes to classical algorithms. We record the resulting selections as *classical_KE* and *classical_SA*, providing a classical reference when PQ options are unavailable.

Using these preferences as anchors, we probe the server's **supported** cryptographic capabilities without enumerating the full Cartesian product of key exchange and signature algorithms. To this end, we fix the key exchange group to (*default_KE*) and iterate over supported PQ signature algorithms. Conversely, we fix the signature algorithm to (*default_SA*), we iterate over all key exchange groups supported by the client to determine the server's supported key establishment mechanisms.

We measure the operational impact of PQC components by applying a differential measurement strategy that isolates the performance contribution of key exchange and signature algorithms independently. To assess the performance impact of PQ *signature algorithms*, we restrict our analysis to domains that advertise support for at least one PQ signature scheme. For each such domain, we fix the client supported key exchange group to *classical_KE*, and iterate over the server's supported PQ signature algorithms. For every evaluated configuration, we immediately perform a corresponding baseline handshake using *classical_KE* and *classical_SA*, allowing any observed differences in latency or message size to be attributed solely to the signature algorithm. Similarly, to assess the performance impact of PQ key exchange mechanisms, we restrict measurements to domains that support hybrid or pure PQ key exchange. For these domains, we fix the signature algorithm to *classical_SA* and iterate over the server's supported hybrid and pure key exchange groups. For hybrid key exchange groups, we further control for the classical component by selecting the corresponding classical key exchange as the baseline (e.g., comparing `X25519MLKEM768` against `X25519`), provided the classical counterpart is supported by the server. This design isolates the performance impact attributable to the PQ component of each construction. For pure PQ key exchanges, we compare against an immediate baseline handshake using *classical_KE* and *classical_SA*.

For each PQ–enabled configuration, we record the TLS handshake time as well as the total number of bytes transmitted and received. Let $T_{\text{pqc}}(D, c)$ denote the handshake time for domain D under a PQ configuration c, and $T_{\text{classical}}(D, c')$ the time for the corresponding classical baseline c'. We define the latency delta as $\Delta_T(D, c) = T_{\text{pqc}}(D, c) - T_{\text{classical}}(D, c')$. Analogously, we compute byte-level deltas for transmitted and received handshake data as $\Delta_{B^{\text{tx}}}(D, c)$ and $\Delta_{B^{\text{rx}}}(D, c)$.

For each evaluated handshake configuration, the client establishes five matched PQ and baseline TLS handshake pairs using identical parameters, spacing successive pairs by four hours. Reporting results across multiple handshakes reduces sensitivity to transient network jitter and short-lived congestion effects. The enforced four-hour spacing further minimizes the influence of localized traffic spikes, routing instability, or load-balancing artifacts, ensuring that observed differences primarily reflect protocol behavior rather than measurement noise. Accordingly, we restrict measurements to at most one scan pair (PQ and baseline-classical) per domain every four hours for any given configuration.

5.4 TLS Management Attribution

To examine whether observed PQ-TLS adoption reflects independent operator migration or infrastructure-driven deployment, we attribute each domain to the entity responsible for terminating its TLS connections and subsequently infer who is potentially responsible for managing its TLS configuration.

We first identify the infrastructure provider behind each domain using a layered resolution strategy. We begin by matching the resolved IP address against known CDN and cloud provider ranges (e.g., Cloudflare[2], Fastly[3]). If no match is found, we analyze DNS CNAME chains, reverse DNS (PTR) lookups, HTTP(S) response metadata, and WHOIS records for provider-specific patterns. When these signals are inconclusive, we use autonomous system (AS) ownership information as a fallback.

Given the inferred infrastructure provider, we then estimate whether TLS configuration is managed by the provider or by the domain operator following [51]. Domains associated with known managed infrastructure platforms (e.g., CDNs), associated with platform-as-a-service offerings (e.g., github.io, blogspot.com), or exhibiting characteristics of shared hosting such as high IP fanout and known infrastructure-linked metadata, are labeled as *potentially infrastructure provider managed*. In contrast, domains with characteristics of dedicated hosting, including low fanout, domain-aligned certificates are labeled as *potentially owner managed*. Domains that cannot be reliably classified are labeled as *unknown*.

This framework provides probabilistic attribution, enabling us to distinguish between deployments *potentially* driven by infrastructure platforms and those managed directly by service operators. Further details and validation are provided in Appendix C.

5.5 Adoption Tracking

To enable continuous analysis of PQ-TLS deployment, we anchor all measurements to DomCop Top 10M domains obtained in July 2025. This identical domain set is reused for scans conducted in July

[2] https://www.cloudflare.com/ips-v4
[3] https://api.fastly.com/public-ip-list

Table 2: Key-exchange groups negotiated under the default configuration and detected through constrained support probes (n = 684,494)

PQC KEM	Default			Support		
	Jul-25	Nov-25	Mar-26	Jul-25	Nov-25	Mar-26
Hybrid KEMs						
X25519MLKEM768	31.26% (214,020)	47.37% (324,276)	49.22% (336,919)	31.27% (214,093)	47.37% (324,288)	49.22% (336,921)
SecP256r1MLKEM768	0	0	0	1.02% (6,981)	3.64% (24,937)	5.53% (37,831)
SecP384r1MLKEM1024	0	0	0	0.01% (83)	0.07% (477)	1.95% (13,348)
Pure PQC KEMs						
ML-KEM-512	0	0	0	0	≈ 0.00% (2)	≈ 0.00% (3)
ML-KEM-768	0	0	0	0	≈ 0.00% (3)	≈ 0.00% (6)
ML-KEM-1024	0	0	0	0	≈ 0.00% (3)	0.66% (4,592)

Default: Indicates the KEM selected during the unconstrained, default TLS 1.3 handshake.

Support: Indicates KEMs that successfully negotiated when explicitly requested.

2025, November 2025, and March 2026, allowing direct observation of changes in default cryptographic negotiation behavior over time without confounding effects from ranking churn or domain turnover.

A domain is treated as unsuccessful if it fails to complete a TLS handshake after three retry attempts.

6 Evaluation

This section characterizes PQ-TLS deployment across the measured public HTTPS endpoints and descriptively compares country and sector patterns with the policy context surveyed in § 4.

TLS Handshake Outcomes

We define a *successful scan* as a domain that completes an unconstrained TLS 1.3 handshake from every measurement vantage point in all three measurement rounds (July 2025, November 2025, and March 2026). Under this definition, 684,494 domains (68.44%) consistently negotiated TLS 1.3 handshakes across all regions and time points. These domains constitute the **stable panel** used for the deployment analysis.

The remaining 315,506 domains (31.55%) failed to complete an unconstrained TLS 1.3 handshake from at least one vantage point in at least one measurement round, despite three retry attempts. The majority of these failures consistently stem from TLS version mismatches, reflecting servers that do not successfully negotiate TLS 1.3, as well as DNS resolution, and transport-layer issues. We further examine these failure modes and their underlying causes in § 6.4.2.

6.1 Configuration Landscape

Across the stable panel, **we do not observe any successful negotiation of PQ signature algorithms**. This absence is consistent across all vantage points and measurement rounds, indicating that PQ signature deployment remains negligible in publicly reachable TLS services during the study period. This finding aligns with the urgent transition focus on PQ key establishment relative to authentication (Appendix 2.6). Consequently, the remainder of our analysis focuses on PQ key establishment, which represent the only

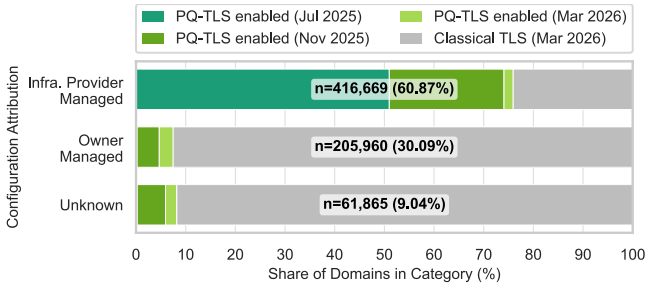

Figure 1: PQ-TLS adoption by management attribution.

observable form of PQC deployment in operational TLS endpoints at present.

Table 2 characterizes PQ key establishment in the stable panel under the default-client handshake and the explicit capability probes. The adoption increases markedly over the eight-month study period. The share selecting a hybrid group under the default-client handshake increases from 31.26% in July 2025 to 47.37% in November 2025, an increase of 16.11 percentage points. It then reaches 49.22% in March 2026, a further increase of 1.85 percentage points. By the end of the measurement period, **nearly half of the stable panel selects hybrid key exchange**, while the remainder selects classical groups.

Despite this rapid adoption, the configuration landscape exhibits minimal diversity. Across all measurement rounds, **every endpoint that negotiates PQ key exchange selects the same hybrid construction, `X25519MLKEM768`**. Alternative hybrid constructions were negotiated only when it is offered in isolation. By March 2026, 5.53% of domains advertise support for `secp256r1MLKEM768` and 1.95% for `secp384r1MLKEM1024`.

Examining the number of PQ configurations supported per domain further highlights this concentration. By March 2026, 336,919 domains advertise support for at least one PQ key exchange mechanism. Among these, 88.76% support exactly one PQ configuration, exclusively `X25519MLKEM768`. An additional 7.28% support two configurations and 3.96% support three, but in all such cases the supported options remain hybrid constructions. **Pure `ML-KEM-*` key exchange groups are not deployed in isolation**, and only three domains advertise a broader set of PQ options that includes both hybrid and pure constructions. None of these alternative hybrid or pure groups is selected under the default-client handshake.

Beyond aggregate adoption, the three measurement rounds also reveal a high degree of deployment stability. In early measurements we observe a small number of domains that advertise support for PQ key exchange but continue to negotiate classical groups. Specifically, in July 2025, 214,093 domains advertise PQ capability while 214,020 select it under the default-client handshake, leaving 73 domains that retain PQ support without prioritizing it. However, this gap disappears in subsequent measurement rounds. By March 2026, only two domains exhibit this behavior. Moreover, we do not observe any instances where domains that previously negotiated PQ key exchange revert to classical-only configurations. These observations suggest that **PQ-TLS rollout proceeds in a largely monotonic manner: once enabled and promoted to the default configuration, deployments tend to remain stable**.

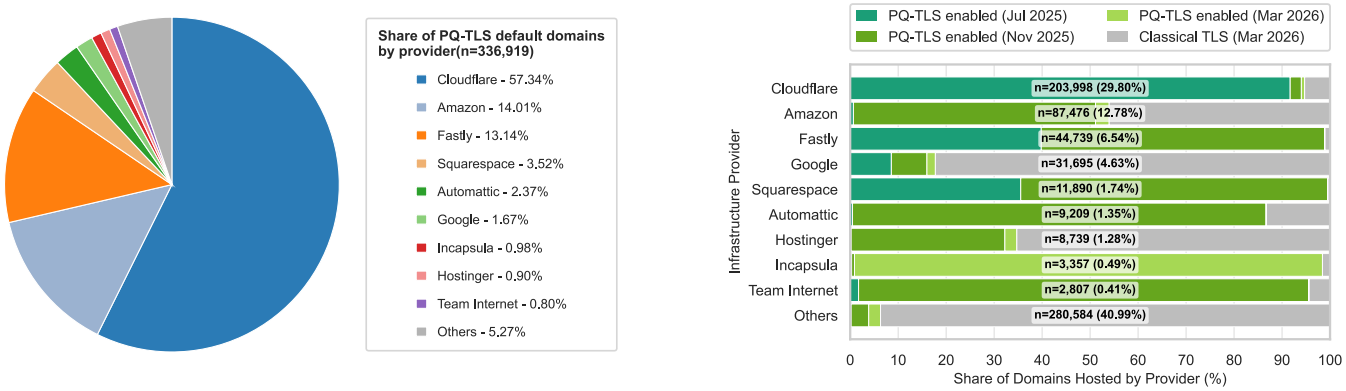

(a) Distribution of domains defaulting to PQ-TLS across infrastructure providers.

(b) PQC adoption across the infrastructure providers hosting most PQ-TLS domains.

Each bar in Subfigure (b) is annotated with the absolute number of served domains and the corresponding fraction of all successful scans (n = 684,494).

Figure 2: Infrastructure concentration of PQ-TLS deployment by March 2026.

> **Key Takeaways for RQ1:** *Despite diverse policy recommendations, operational PQ-TLS has effectively standardized on a single hybrid key exchange construction while signature algorithms are absent in deployment. Observed adoption reflects a narrow set of interoperable default configurations, rather than the broader range of algorithmic options described in PQC transition guidelines.*

6.2 Deployment Agency

To assess whether observed PQ-TLS adoption reflects independent operator migration or infrastructure-driven deployment, we apply the TLS management attribution methodology described in § 5.4. Based on this analysis, we classify each domain as *potentially infrastructure provider managed*, *potentially owner managed*, or *unknown*, reflecting the inferred responsibility for TLS configuration.

Among the 684,494 domains in the stable panel, 416,669 (60.87%) are attributed to potentially infrastructure provider managed deployments, 205,960 (30.09%) are potentially owner managed, and 61,865 (9.04%) remain unattributed.

By March 2026, **93.92% of observed PQ-TLS deployments originate from infrastructure provider managed configurations**. In contrast, **only 4.58% of observed PQ-TLS deployments were owner managed**, with the majority continuing to negotiate classical TLS.

Figure 1 further illustrates this imbalance within attribution categories. More than 75% of infrastructure provider managed domains support PQ-TLS, compared to less than 10% of owner managed domains. Taken together, these disparities indicate that the **majority of observable PQ-TLS adoption is driven by platform-level configuration changes rather than deliberate migration by individual service operators**.

To further examine the role of infrastructure providers, Figure 2-(a) presents the distribution of PQ-TLS–negotiated domains within the providers responsible for terminating their TLS connections. PQ-TLS deployments are highly concentrated among a small number of platforms. In particular, **Cloudflare and Fastly together account for nearly 70% of all observed PQ-TLS adoption**.

Figure 2-(b) shows how PQ-TLS adoption evolves within these infrastructure providers over the course of the measurement period. Early adoption is primarily driven by Cloudflare, Fastly and Squarespace, with a substantial portion of their hosted domains already negotiating PQ-TLS in July 2025. Subsequent measurement rounds show incremental adoption across additional platforms, including Amazon through CloudFront, Squarespace, Automattic, and others. Given the active role of these providers in experimenting with and deploying PQ-TLS, their prominence among PQ-enabled domains is expected. However, **even within infrastructure providers, a considerable fraction of served domains continue to negotiate classical-TLS** as of March 2026, indicating that provider-level rollouts currently remain in progress.

> **Key Takeaways for RQ2:** Observed PQ-TLS adoption is largely driven by infrastructure-provider deployment rather than deliberate migration by individual domain operators. A small number of large-scale platforms account for the majority of PQ-enabled domains, highlighting the central role of infrastructure providers in translating policy expectations into observable deployment.

6.3 Policy Alignment

We assess whether observed PQ-TLS deployment aligns with the priorities and timelines outlined in PQC transition strategies by analyzing trends across policy-relevant sectors (§ 6.3.1) and countries (§ 6.3.2).

6.3.1 Sectoral Adoption. To analyze deployment patterns across sectors, we first categorize domains in the stable panel using SafeDNS domain classification [71]. This taxonomy assigns domains to functional service categories based on the primary content or services they host.

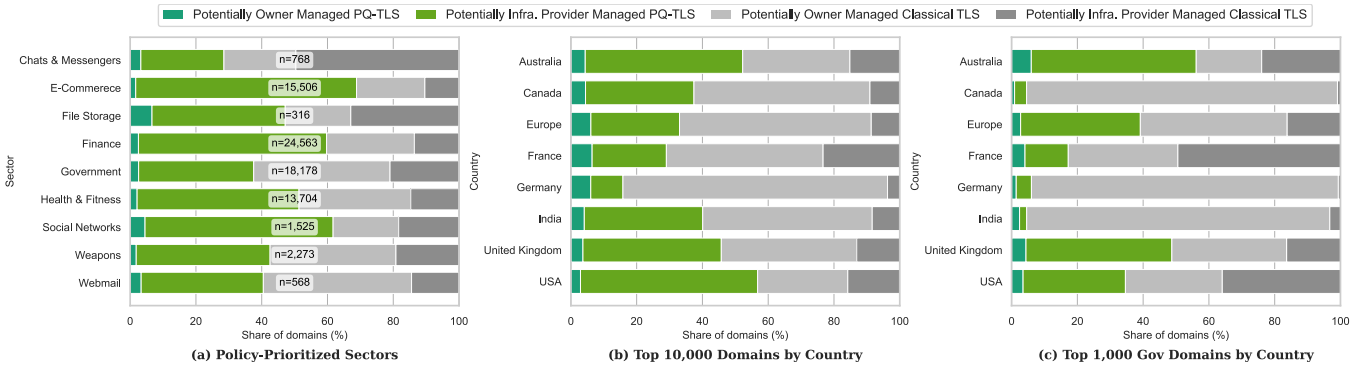

(a) Policy-Prioritized Sectors (b) Top 10,000 Domains by Country (c) Top 1,000 Gov Domains by Country

Figure 3: PQ-TLS adoption across sensitive sectors and Jurisdictions (March 2026)

From this categorization, we focus on sectors that commonly host services exchanging sensitive or high-value information with long confidentiality lifetimes including, government services, banking and financial transactions, healthcare-related services, personal communications (e.g., webmail) and platforms hosting private or regulated content.

Across these policy-relevant sectors, PQ-TLS adoption varies substantially. By March 2026, e-commerce (n = 15,506) and social networks (n = 1,525) exhibit the highest overall adoption, with approximately 68% and 62% of domains negotiating hybrid key exchange under the default-client configuration. Finance (n = 24,563) and healthcare-related services (n = 13,704) follow, with adoption levels of 59% and 51%, respectively. File storage (n = 316) and webmail (n = 568) show moderate uptake at around 46% and 41%. In contrast, government services (n = 18,178) and chats & messengers (n = 768) lag behind, with adoption levels of 38% and 27%, respectively. These results indicate that, **even in policy-prioritized sectors that handle sensitive or long-lived data, PQ deployment remains partial and uneven** across the ecosystem.

While our analysis focuses on policy-relevant sectors, Appendix-D reports PQ-TLS adoption across all SafeDNS categories for completeness.

6.3.2 National Adoption. To characterize national and regional variation in PQ-TLS deployment, we analyze country-specific domain samples drawn from jurisdictions that have published national or regional PQC transition strategies. For each country or region listed in Table 1, we construct domain samples based on country-code top-level domains (ccTLDs). Specifically, we scan the Top 10,000 nationally relevant domains under the corresponding ccTLD (e.g., .au for Australia) and the Top 1,000 government-operated domains under the associated government subdomain (e.g., .gov.au), yielding 11,000 targets per jurisdiction.

Across the general domain sample, PQ-TLS adoption differs markedly across countries. By March 2026, the USA and Australia exhibit the highest adoption levels, with approximately 57% and 53% of domains, respectively, negotiating hybrid key exchange under the default-client configuration. The United Kingdom and India follow, with adoption levels around 45% and 40%. In contrast, France and Germany show substantially lower adoption, at 27% and 16%, respectively.

A similar but more pronounced pattern is observed among government-operated domains. By March 2026, Australia and the United Kingdom again exhibit the highest adoption levels, at 57% and 47%, respectively, although these values remain modest in absolute terms. **In countries such as Canada, Germany, and India, adoption remains extremely limited, with fewer than 7% of government domains defaulting to PQ-TLS.**

6.3.3 Infrastructure Mediation. To understand the drivers of sectoral and national adoption patterns, we analyze PQ-TLS deployment through the lens of deployment agency (§ 5.4).

Figure 3 shows that PQ-TLS adoption across sectors, national top domains, and government services is strongly mediated by the underlying infrastructure. Across all three views, domains served via potentially infrastructure-managed platforms exhibit consistently high adoption rates, typically ranging from approximately 50% to over 80%, while potentially owner-managed deployments remain markedly low, often below 10% and, in many government contexts, below 3%. These patterns indicate that **even in sectors prioritized by national policies, observable PQ-TLS adoption is concentrated among domains that inherit PQ-TLS through managed infrastructure providers**.

6.3.4 Adoption Alignment. Given this strong infrastructure influence, we examine whether national transition timelines correspond to observable differences in PQ-TLS adoption. However, comparing these observations with national transition timelines reveals no clear evidence that earlier policy targets correspond to higher observable PQ-TLS deployment. As Figure 3-(b) shows, **jurisdictions with earlier transition targets (2030–2031) exhibit lower mean adoption rate (32.4%) than those with later horizons (45.7%)**, suggesting counterintuitive alignment with policy timelines.

This apparent reversal is largely explained by deployment composition. When restricting the analysis to potentially owner-managed domains, adoption remains uniformly low across both groups (9.6% vs. 8.4%). Likewise, within Cloudflare-hosted domains, which provide the largest common managed platform across jurisdictions, adoption rates are highly similar across countries (93.5% vs. 94.1%).

Taken together, these results indicate that **cross-sector and cross-jurisdiction variation in PQ-TLS adoption aligns more**

closely with managed infrastructure composition than with sectoral priorities or policy timelines.

> **Key Takeaways for RQ3:** *Observable PQ-TLS deployment shows limited alignment with the sectoral priorities and timelines outlined in national PQC policies. Across sectors and jurisdictions, PQ-enabled deployments are predominantly associated with domains served via managed infrastructure platforms, while owner-managed domains, including government-operated services, remain largely classical.*

6.4 Operational Viability

To provide empirical feedback on the performance and compatibility implications of PQ key exchange, we examine the PQ-TLS deployment along three dimensions: handshake performance (§ 6.4.1), negotiation reliability (§ 6.4.2), and geographic consistency (§ 6.4.3).

6.4.1 Handshake performance. We evaluate the performance impact of PQ key exchange using the differential measurement methodology described in § 5, reporting averages over repeated handshakes to mitigate the network-induced variability. Our primary comparison focuses on the hybrid key exchange group `X25519MLKEM768` against its classical counterpart `X25519`, isolating the incremental cost of the PQ component.

Our measurements across 328,290 domains that successfully negotiated both configurations show that **PQ key exchange does not result in a meaningful increase in TLS handshake latency**. As shown in Figure 4, the median delta is 0 ms, with an interquartile range (IQR) of $[-5, +5]$ ms. Nearly half of the domains (45.90%) exhibit a negative delta, while an additional 31.26% fall within 0–5 ms. Only 13.02% of domains exhibit a delta exceeding 10 ms. These negative deltas may not necessarily imply that PQ key exchange is intrinsically faster; rather, they indicate that the overhead of the hybrid construction is effectively masked by normal network latency variation and connection setup noise.

Absolute handshake time distributions further support this observation. For classical `X25519`, the median handshake time is 32 ms, with an IQR of $[26.0, 61.0]$ ms. Corresponding values for `X25519MLKEM768` are nearly identical, with the same median of 32 ms and an IQR of $[26.0, 59.0]$ ms. Even at the tail, the 90th-percentile handshake times differ by only 5 ms (456 ms for `X25519` vs. 451 ms for `X25519MLKEM768`), indicating that hybrid PQ key exchange does not materially increase end-to-end handshake latency.

While latency impact is negligible, **hybrid PQ key exchange introduces a predictable increase in handshake message size**. Relative to the classical baseline, `X25519MLKEM768` increases client-to-server traffic by a median of 1,176 bytes and server-to-client traffic by 1,088 bytes. This overhead is highly consistent across domains, reflecting the deterministic size of ML-KEM public keys.

6.4.2 Negotiation reliability and failure modes. To assess whether increased handshake message sizes or altered negotiation structure introduce compatibility issues, we analyze TLS handshake failures across the full Top 1M domain set. Of the 315,506 domains that fail to complete a TLS handshake, 57.08% fail before any cryptographic negotiation occurs: DNS resolution and reachability errors account for 105,189 domains (33.34%), while TCP and transport-layer failures account for an additional 74,901 domains (23.74%). These failures typically manifest as *connection refusals*, *unexpected EOF*, or write errors (e.g., `errno=104`).

A further 74,585 domains (23.64%) explicitly reject TLS 1.3 connections due to protocol version or policy constraint related issues. As our client advertises only TLS 1.3 to enable PQ key exchange, these failures reflect a lack of support for contemporary TLS versions rather than incompatibility with PQC algorithms.

Only 60,829 domains (19.28%) fail during the TLS handshake phase where cryptographic parameters are negotiated. Even within this subset, failures are dominated by generic negotiation errors such as *no common cipher suite*, which account for 49,089 cases (80.7% of handshake-phase failures) and typically reflect restrictive classical configurations. Other handshake failures are comparatively rare and include SNI-related rejections (e.g., *unrecognized name*), server internal errors, malformed message handling (e.g., decode errors), and packet-level violations. Collectively, these account for less than 4% of handshake-phase failures.

While certain low-frequency errors, such as *decode errors* or *packet length violations*, could theoretically arise from increased handshake message sizes, their extremely limited occurrence (less than 0.16% of handshake-phase failures), combined with the absence of correlated performance degradation, suggests that **MTU-related or fragmentation-induced failures are not a significant barrier to PQ key exchange deployment in current Internet practice**.

6.4.3 Geographic consistency. In the July 2025 measurement round, we observe 468 domains (0.05% of the stable panel) that negotiate different key exchange mechanisms depending on client location. For these domains, some regions expose hybrid PQ key exchange while others continue to negotiate classical groups.

Many of these domains are hosted on managed infrastructure platforms deployed via globally distributed CDNs (e.g., Amazon CloudFront), where configuration updates may propagate gradually across geographically distributed endpoints. As a result, different regions may reflect different stages of PQ-TLS deployment at a given point in time.

However, in subsequent measurement rounds, the number of domains exhibiting region-dependent negotiation decreases substantially, and by March 2026 most previously inconsistent domains converge to a uniform configuration across all vantage points. This convergence suggests that early PQ-TLS rollouts may initially appear regionally uneven due to gradual configuration propagation across distributed infrastructure, but tend to stabilize as deployments mature.

Across all measurement rounds, we do not observe regional differences in handshake latency, message size overhead, or failure characteristics once PQ-TLS is negotiated. These results indicate that **while early PQ deployments temporarily expose inconsistent configurations across regions, the operational behavior of PQ-TLS remains consistent once configurations are uniformly deployed**.

> **Key Takeaways for RQ4:** *PQ-TLS deployment exhibits strong operational viability across performance, reliability, and consistency dimensions. At Internet scale, hybrid PQ key exchange introduces*

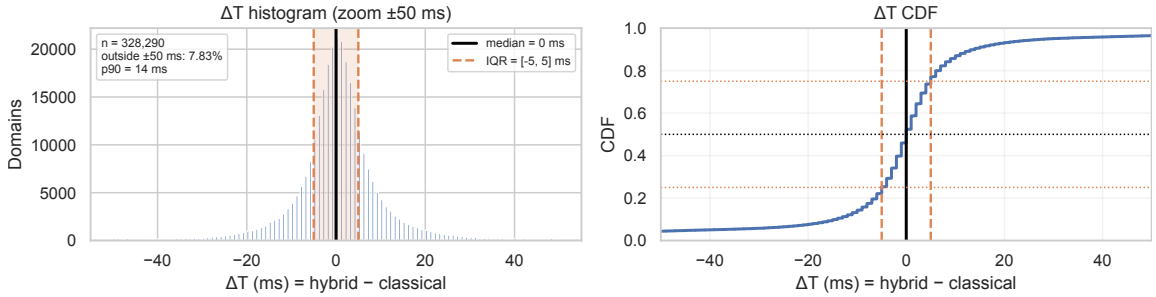

Left: Histogram of Δ_T zoomed to ±50 ms, showing strong concentration around 0 ms with median 0 ms and IQR $[-5, 5]$ ms.
Right: Empirical CDF of Δ_T, indicating that over 90% of domains experience a latency difference of at most 14 ms

Figure 4: Per-domain handshake latency difference $\Delta_T = T_{\text{hybrid}} - T_{\text{classical}}$ for `X25519MLKEM768` relative to `X25519`.

no meaningful increase in TLS handshake latency. Compatibility issues are largely unrelated to PQ mechanisms, instead arising from legacy protocol constraints and general network conditions. PQ-TLS behavior remains consistent across regions once uniformly deployed.

6.5 Security Co-evolution

To examine whether PQ adoption coincides with broader improvements in protocol-level security posture, we compare PQ and classical TLS domains while explicitly accounting for deployment agency.

Our analysis evaluates four dimensions of protocol-level security: support for legacy protocol versions (e.g., TLS 1.0, TLS 1.1), deprecated cipher suites (e.g., 3DES, IDEA, and obsoleted constructions), indicators associated with historically weak TLS configurations (e.g., BEAST, BREACH, Lucky13, Sweet32, and RC4 support), and certificate configuration issues (e.g., hostname mismatch and chain-of-trust errors). Full labeling criteria are provided in Appendix E.

We evaluate these properties using `testssl.sh` [78]. To minimize operational impact, we restrict measurements to non-intrusive tests and enforce a delay of five minutes between successive scans of the same domain, following [10].

6.5.1 Provider-managed domains. For *potentially infrastructure provider managed* domains, we compare PQ-TLS and classical TLS within the same providers, retaining those with at least 1,000 domains in each class and sampling evenly.

Under this provider-balanced comparison, **PQ-TLS domains more frequently retain legacy compatibility features**. Approximately 18.0% support TLS 1.0/TLS 1.1 compared to 14.7% of matched classical domains, and 73.6% advertise at least one deprecated cipher suite versus 49.9%. This pattern extends to vulnerability indicators: PQ-TLS domains more frequently retain legacy cryptographic features, such as CBC-mode ciphers, TLS 1.0-era record handling, and compression mechanisms, which corresponds to a higher prevalence of indicators associated with BEAST (17.2% vs. 14.3%), BREACH (50.6% vs. 31.7%), while support for stream ciphers such as, RC4 is absent in both classes.

Certificate-management signals are stronger among PQ-TLS deployments. OCSP stapling is more common (18.2% vs. 3.9%), hostname mismatches are less frequent (0.7% vs. 10.4%), and chain-of-trust issues are rare in both groups.

These patterns vary across providers, reflecting platform-specific configurations. For example, Cloudflare exhibits higher legacy protocol and cipher support among PQ-TLS domains, Fastly shows lower legacy protocol support but higher deprecated cipher usage, while Amazon CloudFront presents a mixed profile. These differences indicate that **managed PQ-TLS deployments largely inherit the TLS compatibility profile of the underlying platform**.

6.5.2 Owner-managed domains. For domains classified as *potentially owner managed*, we compare randomly sample PQ-TLS and classical TLS domains (10,000 each).

In contrast, **PQ-TLS adoption among owner-managed domains is associated with modestly stricter configurations**, with lower support for legacy protocol versions (10.5% vs. 11.9%) and deprecated cipher suites (46.9% vs. 52.6%). This trend is also reflected in a lower prevalence of vulnerability indicators associated with BEAST (8.0% vs. 10.0%), BREACH (40.0% vs. 43.1%), and Lucky13 (46.6% vs. 52.3%), along with reduced support for deprecated stream ciphers such as RC4 (0.04% vs. 0.30%), although indicators linked to 64-bit block cipher usage (Sweet32) remain more common (6.8% vs. 4.5%).

Certificate-related signals align similarly, with higher OCSP stapling rates (12.9% vs. 6.3%), fewer hostname mismatches (5.0% vs. 7.5%) and chain-of-trust issues (2.9% vs. 5.9%).

Taken together, these results indicate that **security co-evolution during early PQ-TLS deployment is strongly shaped by *deployment agency***. Among provider-managed domains, PQ-TLS is typically introduced through platforms that combine strong certificate automation with broad backward compatibility. Among owner-managed domains, PQ-TLS adoption more closely reflects deliberate protocol modernization.

Key Takeaways for RQ5: *PQ-TLS adoption does not imply a uniform improvement in protocol-level security. Provider-managed deployments inherit platform-level configurations that combine stronger certificate practices with greater retention of legacy protocol features, while owner-managed deployments are associated with modestly stricter TLS configurations.*

7 Discussion

This section synthesizes the key findings from our measurements and examines their implications for ongoing PQC transition efforts. By contextualizing policy expectations with observed deployment behavior, we highlight where current strategies align with operational realities and where adjustments may be necessary to better support the transition.

Configuration Convergence in Practice. Although national PQC transition policies frequently emphasize algorithm agility and recommend multiple candidate constructions, our measurements show that operational TLS deployments in HTTPS have effectively converged on a single hybrid configuration: `X25519MLKEM768`. Alternative constructions appear only rarely as secondary capabilities and are almost never prioritized by servers as negotiation choices. This convergence reflects practical deployment constraints, including interoperability requirements, browser defaults, and infrastructure-provider support, which naturally steer the ecosystem toward a small set of widely supported configurations. Policy guidance that emphasizes broad algorithmic diversity without corresponding deployment incentives risks becoming more aspirational than actionable. Policymakers may therefore achieve greater impact by recognizing de facto convergence and focusing guidance on the deployment, monitoring, and evolution of dominant constructions while keeping the recommended design space minimal and interoperable.

Infrastructure Providers as Primary Deployment Actors. Our analysis shows that the majority of observable PQ-TLS deployment in HTTPS is driven by a small number of managed infrastructure providers rather than independent migration decisions by individual domain operators. Content delivery networks and cloud platforms effectively act as deployment multipliers: once a provider enables PQ-TLS, thousands of served domains inherit the capability simultaneously. This architecture significantly accelerates visible adoption and also concentrates influence over the transition within a small number of actors. Engaging infrastructure providers directly through guidance, coordination, and transparency mechanisms may therefore offer greater leverage, given their central role in shaping deployment across the Internet.

Understanding National Adoption Differences. Our sectoral analysis indicates that PQ-TLS deployment in HTTPS does not uniformly follow the sectoral priorities articulated in many national transition strategies. Sectors commonly highlighted in policy discussions, such as government, exhibit comparatively low levels of adoption, while several platform-driven consumer services show substantially higher deployment. This disparity reflects structural differences in service provisioning: many consumer-facing platforms inherit cryptographic capabilities through shared managed infrastructure, whereas sectors such as government and healthcare face procurement cycles, legacy system dependencies, and regulatory constraints that slow the adoption of new cryptographic mechanisms [31]. These findings suggest that policies may benefit from distinguishing between sectors that can passively inherit security properties through shared infrastructure and those that require targeted support, funding, or technical assistance to enable migration within more constrained operational environments.

Operational Feasibility of PQ-TLS. Concerns about performance, compatibility, and reliability often appear implicitly in PQC transition discussions, yet our measurements show that PQ-TLS key exchange does not introduce a meaningful latency penalty and causes no systemic negotiation failures at Internet scale. Observed failures are dominated by pre-existing DNS, transport, and TLS version issues rather than PQC itself. These results support the technical feasibility assumptions underlying many policies, while underscoring the value of empirical validation. Embedding continuous measurement into transition planning can help policymakers distinguish genuine from perceived deployment risks and adjust their guidance accordingly as the ecosystem evolves.

PQC Readiness and Broader TLS Security Practices. Our analysis shows that PQ-TLS adoption in HTTPS does not correspond to a uniform shift toward stricter TLS configurations; instead, its security implications depend strongly on deployment context. Provider-managed domains show strong certificate automation with continued support for legacy protocol versions and deprecated cipher suites, reflecting a compatibility-oriented deployment model. In contrast, PQ-TLS adoption among owner-managed domains is associated with modestly stricter configurations, including reduced support for legacy features. These patterns indicate that PQC transition is not a simple cryptographic substitution, but is shaped by how and where deployment occurs. More broadly, they highlight that effective security outcomes depend not only on adopting new cryptographic mechanisms, but also on the surrounding protocol configuration and operational practices in which they are embedded.

Implications for PQC Transition Efforts. Taken together, these findings suggest that the PQC transition is currently progressing primarily through infrastructure-mediated deployment rather than through coordinated migration by independently managed systems. While large CDNs and cloud platforms can rapidly introduce PQ-TLS across thousands of hosted domains, most owner-managed services depend on operating system distributions and packaged software stacks for cryptographic capabilities. Support for PQ-enabled TLS has already been incorporated into widely used cryptographic libraries and server software for roughly a year. For example, OpenSSL 3.5.0, released in April 2025, introduced built-in support for the NIST-standardized PQC algorithms [60]. Yet our measurements show that adoption outside large infrastructure platforms remains limited. This gap indicates that software availability alone is insufficient to drive deployment. Supporting the broader transition will therefore require more than technical standardization of *which* algorithms to adopt by *when*. Practical guidance, deployment tooling, and operator education will be important to help organizations understand *how* to enable PQ mechanisms within their infrastructure and integrate them safely into existing operational environments.

Aligning Policy with Protocol Standardization. Effective deployment guidance must also distinguish between standardizing cryptographic algorithms and specifying how those algorithms are used within Internet protocols. NIST standardization defines the underlying cryptographic primitives, while IETF specifications define how they are represented and negotiated within TLS. RFC 10024, published after our measurement period, now specifies the three hybrid TLS 1.3 groups examined in our capability measurements

[50]. Standalone ML-KEM key exchange and ML-DSA authentication remain in the IETF publication process [24, 42], while broader application guidance is also being developed [64]. This difference in protocol maturity helps explain why hybrid key exchange has progressed more rapidly than PQ authentication. Future policy updates could engage more directly with the IETF process by distinguishing deployment-ready protocol profiles from algorithms whose protocol integration is still evolving. Aligning policy milestones with protocol-standardization milestones could make transition guidance more actionable and support coordinated implementation.

Acceleration and Limits of Security Transitions. Historical transitions in TLS illustrate both the accelerating pace and the limits of security adoption at Internet scale. The shift from TLS 1.1 to TLS 1.2 required nearly five years to reach 15% deployment, reflecting a largely operator-driven upgrade process. In contrast, driven by a small number of major infrastructure providers, TLS 1.3 reached nearly 48% adoption within two years and four months following its standardization in August 2018 [51]. We observe a similar but even more pronounced trend for PQ-TLS in HTTPS, which achieves comparable adoption levels (47.37%) in approximately one year and three months (November 2025), driven by the same infrastructure-mediated deployment model. However, accelerated early adoption does not necessarily imply complete transition. Even more than seven years after standardization, TLS 1.3 has not yet reached universal deployment (§ 6.4.2), reflecting persistent long-tail inertia among systems. For PQC, this long tail represents a growing security risk. Successful PQC transition will therefore require sustained effort to mitigate inertia through early preparation, stronger coordination across stakeholders, and continued investment to support secure and widespread deployment.

8 Limitations

We use the surveyed policies to contextualize the measured country and sector patterns rather than to estimate a causal policy effect. The documents vary in authority, scope, and implementation timeline, and most milestones occur several years after our final measurement round. Consequently, limited correspondence between policy guidance and measured deployment should not be interpreted as evidence of policy ineffectiveness or noncompliance. Instead, our study highlights the current gap between early, externally observable PQC deployment and the broader transition objectives articulated in national policy guidance, providing a baseline against which future progress can be assessed.

Our measurements rely on active scanning of public-facing domains drawn from the Top 1M, and therefore capture only the *public Internet*. We do not observe PQ deployments within private intranets, VPNs, or other internal networks not exposed to the public Internet. While some national PQC strategies prioritize national or classified systems (e.g., CNSA 2.0 [2]), many published transition roadmaps (e.g., BSI [34]) address a broader set of systems, including Internet-facing public services. Measuring the public web therefore remains the most appropriate proxy for evaluating the externally observable component of these policy mandates.

Our measurements focus exclusively on TLS as used for HTTPS. While TLS is also employed by other application protocols (e.g., email, messaging, and custom services), HTTPS dominates Internet-facing cryptographic traffic [30]. As such, HTTPS-based TLS provides the most operationally significant surface for evaluating PQ deployment at Internet scale.

All measurements were conducted from cloud-based vantage points using distributed measurement infrastructure. While this setup enables globally comparable observations, it may not fully capture the experience of end users operating from residential networks, enterprise environments, or constrained network conditions. Factors such as last-mile latency, network congestion, and device heterogeneity can influence handshake performance and reliability in ways not reflected in our study. As a result, observed operational viability should be interpreted as representative of well-provisioned network paths rather than all end-user environments.

Taken together, these limitations define the intended scope of our results. We empirically assess PQ-TLS readiness in HTTPS, configuration, and operational viability on the public web, in line with contemporary PQC transition policies.

9 Conclusion

This paper presents the first Internet-scale measurement study of PQ-TLS deployment in HTTPS. By analyzing more than 2 billion TLS handshakes across one million domains, collected from 11 globally distributed vantage points, we provide an empirical view of how PQC is transitioning from standardization into real-world deployment.

Our measurements reveal limited correspondence between national policy expectations and early deployment patterns. While nearly half of the stable domain panel prioritize PQ key exchange by March 2026, deployment is narrowly concentrated around a single hybrid construction, `X25519MLKEM768`. In contrast, we find no verified deployment of PQ signature algorithms for TLS authentication. PQ-TLS adoption is overwhelmingly driven by managed infrastructure providers, while owner-managed and government-operated domains largely default to classical. Across sectors and countries, observable deployment shows limited correspondence with policy timelines and priorities.

At the same time, our results demonstrate that PQ-TLS is operationally viable at Internet scale. Hybrid key exchange introduces no meaningful increase in handshake latency, and PQ mechanisms do not result in observable compatibility failures. However, PQ adoption does not necessarily coincide with broader TLS hardening, as domains often retain legacy protocol versions and deprecated cipher support.

These findings suggest that successful PQC transition will require more than algorithm standardization and calendar-based migration targets. Policies should account for the central role of infrastructure providers, support owner-managed systems with concrete deployment guidance and tooling, and incorporate empirical protocol-level measurements into progress tracking. Without this operational focus, PQC migration may become visible in aggregate statistics while remaining uneven across sectors, dependent on a small number of platforms, and only partially aligned with the security outcomes transition policies aim to achieve.

References

[1] David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry, Matthew Green, J. Alex Halderman, Nadia Heninger, Drew Springall, Emmanuel Thomé, Luke Valenta, Benjamin VanderSloot, Eric Wustrow, Santiago Zanella-Béguelin, and Paul Zimmermann. 2015. Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15)*. Association for Computing Machinery. doi:10.1145/2810103.2813707

[2] National Security Agency. 2024. The Commercial National Security Algorithm Suite 2.0 and Quantum Computing FAQ. https://media.defense.gov/2022/Sep/07/2003071836/-1/-1/0/CSI_CNSA_2.0_FAQ_.PDF. [Accessed 27-11-2025].

[3] Nadeem Ahmed, Lei Zhang, and Aryya Gangopadhyay. 2025. A Survey of Post-Quantum Cryptography Support in Cryptographic Libraries. arXiv:2508.16078 [cs.CR] https://arxiv.org/abs/2508.16078

[4] Akamai. 2025. Akamai Enables Post-Quantum Cryptography on the Edge. https://www.akamai.com/blog/security/akamai-enables-post-quantum-cryptography-edge. [Accessed 13-12-2025].

[5] Johanna Amann, Oliver Gasser, Quirin Scheitle, Lexi Brent, Georg Carle, and Ralph Holz. 2017. Mission accomplished? HTTPS security after diginotar. In *Proceedings of the 2017 Internet Measurement Conference (IMC '17)*. Association for Computing Machinery. doi:10.1145/3131365.3131401

[6] Amazon Web Services. 2025. s2n-tls Version 1.6.0 Release Notes. https://github.com/aws/s2n-tls/releases/tag/v1.6.0. Added post-quantum key exchange to the default security policy. Accessed August 14, 2026.

[7] ANSSI. 2023. ANSSI views on the Post-Quantum Cryptography transition. https://messervices.cyber.gouv.fr/documents-guides/follow_up_position_paper_on_post_quantum_cryptography.pdf. [Accessed 11-26-2025].

[8] Australian Signals Directorate and Australian Cyber Security Centre. 2025. Information security manual: Guidelines for cryptography. https://www.cyber.gov.au/sites/default/files/2025-07/22.%20ISM%20-%20Guidelines%20for%20cryptography%20%28June%202025%29.pdf. Accessed: 2025-11-23.

[9] Australian Signals Directorate and Australian Cyber Security Centre. 2025. Planning for post-quantum cryptography. https://www.cyber.gov.au/sites/default/files/2025-09/Planning%20for%20post-quantum%20cryptography%20%28September%202025%29.pdf. Accessed: 2025-11-23.

[10] Michael Bailey, David Dittrich, Erin Kenneally, and Doug Maughan. 2012. The Menlo Report. *IEEE Security and Privacy* (2012). doi:10.1109/MSP.2012.52

[11] Kaveh Bashiri, Scott Fluhrer, Stefan-Lukas Gazdag, Daniel Van Geest, and Stavros Kousidis. 2025. Internet X.509 Public Key Infrastructure – Algorithm Identifiers for the Stateless Hash-Based Digital Signature Algorithm (SLH-DSA). RFC 9909. https://www.rfc-editor.org/info/rfc9909

[12] Benjamin Beurdouche, Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub, and Jean Karim Zinzindohoue. 2015. A Messy State of the Union: Taming the Composite State Machines of TLS. In *2015 IEEE Symposium on Security and Privacy*. doi:10.1109/SP.2015.39

[13] Chromium Blog. [n. d.]. Advancing Our Amazing Bet on Asymmetric Cryptography — blog.chromium.org. https://blog.chromium.org/2024/05/advancing-our-amazing-bet-on-asymmetric.html. [Accessed 23-03-2026].

[14] National Cyber Security Centre. 2024. Next steps in preparing for post-quantum cryptography. https://www.ncsc.gov.uk/whitepaper/next-steps-preparing-for-post-quantum-cryptography. [Accessed 26-11-2025].

[15] National Cyber Security Centre. 2025. Timelines for migration to post-quantum cryptography. https://www.ncsc.gov.uk/guidance/pqc-migration-timelines. [Accessed 26-11-2025].

[16] NIST Computer Security Resource Centre. 2017. Round 1 Submissions - Post-Quantum Cryptography. https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-1-submissions. [Accessed 11-11-2025].

[17] Cert-IN and SISA. 2025. Transition to Quantum Cyber Readiness. https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/jul/doc2025714585911.pdf. [Accessed 27-11-2025].

[18] Institute of Commercial Cryptography Standards China. [n. d.]. Next Generation Commercial Cryptographic Algorithms Program (NGCC). https://www.niccs.org.cn/en/. [Accessed 05-12-2025].

[19] Google Chrome. [n. d.]. Overview of CrUX | Chrome UX Report. https://developer.chrome.com/docs/crux. [Accessed 08-03-2026].

[20] Chromium Authors. 2025. Default TLS Supported Groups and Key Shares in Chromium. https://chromium.googlesource.com/chromium/src/+/0d4ac21b930c1b6f730d0b12b41b7cd3c2904db7/services/network/public/mojom/ssl_config.mojom. Accessed August 14, 2026.

[21] Cloudflare. 2025. State of the post-quantum Internet in 2025. https://blog.cloudflare.com/pq-2025. [Accessed 28-12-2025].

[22] Cloudflare Inc. 2026. Post-Quantum Encryption and Key Transparency on Cloudflare Radar · Changelog. https://developers.cloudflare.com/changelog/post/2026-02-27-radar-pq-key-transparency. [Accessed 13-08-2026].

[23] European Commission. 2025. A Coordinated Implementation Roadmap for the Transition to Post-Quantum Cryptography. https://digital-strategy.ec.europa.eu/en/library/coordinated-implementation-roadmap-transition-post-quantum-cryptography. [Accessed 11-28-2025].

[24] Deirdre Connolly. 2026. *ML-KEM Post-Quantum Key Agreement for TLS 1.3*. Internet-Draft draft-ietf-tls-mlkem-09. Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-ietf-tls-mlkem/09/ Work in Progress.

[25] Common Crawl. [n. d.]. Common Crawl - Open Repository of Web Crawl Data. https://commoncrawl.org. [Accessed 25-03-2026].

[26] Zakir Durumeric, Michael Bailey, and J. Alex Halderman. 2014. An Internet-Wide View of Internet-Wide Scanning. In *23rd USENIX Security Symposium (USENIX Security 14)*. USENIX Association.

[27] Zakir Durumeric, James Kasten, Michael Bailey, and J. Alex Halderman. 2013. Analysis of the HTTPS certificate ecosystem. In *Proceedings of the 2013 Conference on Internet Measurement Conference (IMC '13)*. Association for Computing Machinery. doi:10.1145/2504730.2504755

[28] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-wide Scanning and Its Security Applications. In *22nd USENIX Security Symposium (USENIX Security 13)*. USENIX Association.

[29] NIST Dustin Moody. 2017. THE SHIP HAS SAILED: The NIST Post-Quantum Crypto "Competition". https://csrc.nist.gov/csrc/media/projects/post-quantum-cryptography/documents/asiacrypt-2017-moody-pqc.pdf. [Accessed 11-11-2025].

[30] Adrienne Porter Felt, Richard Barnes, April King, Chris Palmer, Chris Bentzel, and Parisa Tabriz. 2017. Measuring HTTPS Adoption on the Web. In *26th USENIX Security Symposium (USENIX Security 17)*. USENIX Association.

[31] Konstantin Fischer, Ivana Trummová, Phillip Gajland, Yasemin Acar, Sascha Fahl, and Angela Sasse. 2024. The Challenges of Bringing Cryptography from Research Papers to Products: Results from an Interview Study with Experts. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association.

[32] Canadian Centre for Cyber Security. 2025. Cryptographic algorithms for UNCLASSIFIED, PROTECTED A, and PROTECTED B information - ITSP.40.111. https://www.cyber.gc.ca/en/guidance/cryptographic-algorithms-unclassified-protected-protected-b-information-itsp40111. [Accessed 05-12-2025].

[33] Canadian Centre for Cyber Security. 2025. Roadmap for the migration to post-quantum cryptography for the Government of Canada (ITSM.40.001). https://www.cyber.gc.ca/en/guidance/roadmap-migration-post-quantum-cryptography-government-canada-itsm40001. [Accessed 05-12-2025].

[34] Bundesamt für Sicherheit in der Informationstechnik (BSI). [n. d.]. BSI – Technical Guideline. https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/TechGuidelines/TG02102/BSI-TR-02102-1.pdf?__blob=publicationFile&v=7. [Accessed 08-01-2026].

[35] Bundesamt für Sicherheit in der Informationstechnik (BSI). 2024. Securing Tomorrow, Today: Transitioning to Post-Quantum Cryptography. https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Crypto/PQC-joint-statement.pdf?__blob=publicationFile&v=5. [Accessed 26-12-2025].

[36] Google. 2016. Experimenting with Post-Quantum Cryptography. https://security.googleblog.com/2016/07/experimenting-with-post-quantum.html. [Accessed 13-12-2025].

[37] Google. 2025. Post-quantum cryptography (PQC) — cloud.google.com. https://cloud.google.com/security/resources/post-quantum-cryptography. [Accessed 11-03-2026].

[38] CloudFlare Google. 2019. The TLS Post-Quantum Experiment. https://blog.cloudflare.com/the-tls-post-quantum-experiment/. [Accessed 13-12-2025].

[39] World Bank Group. [n. d.]. World Bank Open Data. https://data.worldbank.org/indicator/NY.GDP.MKTP.CD?most_recent_value_desc=true. [Accessed 11-11-2025].

[40] Rumaisa Habib, Kimberly Ruth, Gautam Akiwate, and Zakir Durumeric. 2025. Formalizing Dependence of Web Infrastructure. In *Proceedings of the ACM SIGCOMM 2025 Conference (SIGCOMM '25)*. Association for Computing Machinery.

[41] Marcella Hastings, Joshua Fried, and Nadia Heninger. 2016. Weak Keys Remain Widespread in Network Devices. In *Proceedings of the 2016 Internet Measurement Conference (IMC '16)*. Association for Computing Machinery. doi:10.1145/2987443.2987486

[42] Tim Hollebeek, Sophie Schmieg, and Bas Westerbaan. 2026. *Use of ML-DSA in TLS 1.3*. Internet-Draft draft-ietf-tls-mldsa-05. Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-ietf-tls-mldsa/05/ Work in Progress.

[43] Ralph Holz, Lothar Braun, Nils Kammenhuber, and Georg Carle. 2011. The SSL Landscape: A Thorough Analysis of the X.509 PKI Using Active and Passive Measurements. In *Proceedings of the 11th ACM SIGCOMM Conference on Internet Measurement (IMC '11)*. Association for Computing Machinery.

[44] Ralph Holz, Jens Hiller, Johanna Amann, Abbas Razaghpanah, Thomas Jost, Narseo Vallina-Rodriguez, and Oliver Hohlfeld. 2020. Tracking the deployment of TLS 1.3 on the web: a story of experimentation and centralization. *SIGCOMM Comput. Commun. Rev.* (2020). doi:10.1145/3411740.3411742

[45] The White House. 2022. National Security Memorandum on Promoting United States Leadership in Quantum Computing While Mitigating Risks to Vulnerable Cryptographic Systems. https://irp.fas.org/offdocs/nsm/nsm-10.pdf. [Accessed 27-11-2025].

[46] ISO/IEC. [n. d.]. ISO/IEC JTC 1/SC 27 WG2 PQC work program. https://committee.iso.org/sites/jtc1sc27/home/wg2.html. [Accessed 12-01-2026].

[47] Liz Izhikevich, Renata Teixeira, and Zakir Durumeric. 2022. Predicting IPv4 services across all ports. In *Proceedings of the ACM SIGCOMM 2022 Conference (SIGCOMM '22)*. Association for Computing Machinery. doi:10.1145/3544216.3544249

[48] Marcel Kempf, Nikolas Gauder, Benedikt Jaeger, Johannes Zirngibl, and Georg Carle. 2024. A Quantum of QUIC: Dissecting Cryptography with Post-Quantum Insights. In *2024 IFIP Networking Conference (IFIP Networking)*. 195–203. doi:10.23919/IFIPNetworking62109.2024.10619916

[49] Platon Kotzias, Abbas Razaghpanah, Johanna Amann, Kenneth G. Paterson, Narseo Vallina-Rodriguez, and Juan Caballero. 2018. Coming of Age: A Longitudinal Study of TLS Deployment. In *Proceedings of the Internet Measurement Conference 2018 (IMC '18)*. Association for Computing Machinery. doi:10.1145/3278532.3278568

[50] K. Kwiatkowski, P. Kampanakis, B. E. Westerbaan, and D. Stebila. 2026. Post-Quantum Traditional (PQ/T) Hybrid Key Agreement Mechanisms for TLS 1.3. RFC 10024. doi:10.17487/RFC10024

[51] Hyunwoo Lee, Doowon Kim, and Yonghwi Kwon. 2021. TLS 1.3 in Practice:How TLS 1.3 Contributes to the Internet. In *Proceedings of the Web Conference 2021 (WWW '21)*. Association for Computing Machinery. doi:10.1145/3442381.3450057

[52] Joonhee Lee, Hyunwoo Lee, Jongheon Jeong, Doowon Kim, and Ted Taekyoung Kwon. 2021. Analyzing Spatial Differences in the TLS Security of Delegated Web Services. In *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security (ASIA CCS '21)*. Association for Computing Machinery. doi:10.1145/3433210.3453107

[53] Zane Ma, Aaron Faulkenberry, Thomas Papastergiou, Zakir Durumeric, Michael D. Bailey, Angelos D. Keromytis, Fabian Monrose, and Manos Antonakakis. 2023. Stale TLS Certificates: Investigating Precarious Third-Party Access to Valid TLS Keys. In *Proceedings of the 2023 ACM on Internet Measurement Conference (IMC '23)*. Association for Computing Machinery. doi:10.1145/3618257.3624802

[54] Michele Mosca. 2018. Cybersecurity in an Era with Quantum Computers: Will We Be Ready? *IEEE Security & Privacy* 16, 5 (2018), 38–41. doi:10.1109/MSP.2018.3761723

[55] Mozilla. 2024. Bug 1918484: Replace xyber768d00 with mlkem768x25519. https://bugzilla.mozilla.org/show_bug.cgi?id=1918484. Implemented in Firefox 132. Accessed August 14, 2026.

[56] Ministry of Communications Government of India. 2025. Migration to Post Quantum Cryptography. https://www.tec.gov.in/pdf/TR/Final%20technical%20report%20on%20migration%20to%20PQC%2028-03-25.pdf. [Accessed 27-11-2025].

[57] National Institute of Standards and Technology (NIST). 2024. Stateless Hash-Based Digital Signature Standard. doi:10.6028/NIST.FIPS.205

[58] National Institute of Standards, Technology (NIST), Gorjan Alagic, Quynh Dang, Dustin Moody, Angela Robinson, Hamilton Silberg, and Daniel Smith-Tone. 2024. Module-Lattice-Based Key-Encapsulation Mechanism Standard. doi:10.6028/NIST.FIPS.203

[59] National Institute of Standards, Technology (NIST), Thinh Dang, Jacob Lichtinger, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, and Angela Robinson. 2024. Module-Lattice-Based Digital Signature Standard. doi:10.6028/NIST.FIPS.204

[60] OpenSSL Library. 2025. OpenSSL 3.5 Final Release - Live | OpenSSL Library — openssl-library.org. https://openssl-library.org/post/2025-04-08-openssl-35-final-release/. [Accessed 29-03-2026].

[61] Christian Paquin, Douglas Stebila, and Goutam Tamvada. 2020. Benchmarking post-quantum cryptography in TLS. In *Proc. 11th International Conference on Post-Quantum Cryptography (PQCrypto) 2020 (LNCS)*. Springer.

[62] Victor Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczynski, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. doi:10.14722/ndss.2019.23386

[63] CloudFlare Radar. [n. d.]. Data Explorer | Cloudflare Radar. https://radar.cloudflare.com/explorer?dataSet=http&groupBy=post_quantum. [Accessed 12-12-2025].

[64] Tirumaleswar Reddy.K and Hannes Tschofenig. 2025. *Post-Quantum Cryptography Recommendations for TLS-based Applications.* Internet-Draft draft-ietf-uta-pqc-app-00. Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-ietf-uta-pqc-app/00/ Work in Progress.

[65] Tirumaleswar Reddy.K and Hannes Tschofenig. 2026. *Post-Quantum Cryptography Recommendations for TLS-based Applications.* Internet-Draft draft-ietf-uta-pqc-app-03. Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-ietf-uta-pqc-app/03/ Work in Progress.

[66] Eric Rescorla. 2018. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446.

[67] Quantum resistant cryptography research group. [n. d.]. KPQC. https://kpqc.or.kr/. [Accessed 05-12-2025].

[68] Ruben Rios, José A. Montenegro, Antonio Muoz, and Davide Ferraris. 2025. Toward the Quantum-Safe Web: Benchmarking Post-Quantum TLS. *IEEE Network* (2025).

[69] rustls Project. 2025. rustls 0.23.27 Release Notes. https://github.com/rustls/rustls/releases/tag/v%2F0.23.27. Introduced preference for post-quantum key exchange algorithms by default. Accessed August 14, 2026.

[70] Kimberly Ruth, Deepak Kumar, Brandon Wang, Luke Valenta, and Zakir Durumeric. 2022. Toppling top lists: evaluating the accuracy of popular website lists. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22)*. Association for Computing Machinery.

[71] Safe DNS. [n. d.]. List of SafeDNS Categories — SafeDNS docs. https://docs.safedns.com/books/installation-guides/page/list-of-safedns-categories. [Accessed 25-03-2026].

[72] Amazon Web Services. 2025. Bringing quantum-resistance to AWS services and customers. https://aws.amazon.com/security/post-quantum-cryptography/. [Accessed 12-12-2025].

[73] P.W. Shor. 1994. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*.

[74] Dimitrios Sikeridis, Panos Kampanakis, and Michael Devetsikiotis. 2020. Post-Quantum Authentication in TLS 1.3: A Performance Study. In *Proceedings of the 27th NDSS Symposium (NDSS 2020)*. doi:10.14722/ndss.2020.24203

[75] Markus Sosnowski, Florian Wiedner, Eric Hauser, Lion Steger, Dimitrios Schoinianakis, Sebastian Gallenmüller, and Georg Carle. 2023. The Performance of Post-Quantum TLS 1.3. In *Companion of the 19th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT 2023)*. Association for Computing Machinery. doi:10.1145/3624354.3630585

[76] Douglas Stebila, Scott Fluhrer, and Shay Gueron. 2025. *Hybrid key exchange in TLS 1.3.* Internet-Draft draft-ietf-tls-hybrid-design-14. Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-ietf-tls-hybrid-design/14/ Work in Progress.

[77] Indian Computer Emergency Response Team. [n. d.]. Guidelines on Information Security Practices for Government Entities. https://www.cert-in.org.in/PDF/guidelinesgovtentities.pdf. [Accessed 05-12-2025].

[78] TestSSL. [n. d.]. SSL/TLS tester: testssl.sh — testssl.sh. https://testssl.sh. [Accessed 25-03-2026].

[79] The Go Project. 2025. Go 1.24 Release Notes. https://go.dev/doc/go1.24. Accessed August 14, 2026.

[80] Iraklis Tzinos, Konstantinos Limniotis, and Nicholas Kolokotronis. 2022. Evaluating the performance of post-quantum secure algorithms in the TLS protocol. *Journal of Surveillance, Security and Safety* (2022).

[81] National Cyber Security Centre UK. [n. d.]. Using TLS to protect data. https://www.ncsc.gov.uk/guidance/using-tls-to-protect-data. [Accessed 05-12-2025].

[82] Nimesha Wickramasinghe, Mohamed Nabeel, Kenneth Thilakaratne, Chamath Keppitiyagama, and Kasun De Zoysa. 2021. Uncovering IP Address Hosting Types Behind Malicious Websites. arXiv:2111.00142 [cs.CR] https://arxiv.org/abs/2111.00142

Ethics

Internet-wide scanning is a well-established practice to understand the topology and security posture of the global network [26, 47], with prior work demonstrating its importance for assessing systemic risks (see §2). Our study adheres to this tradition, employing only lightweight TLS handshakes, which we gracefully terminate immediately after handshake completion and before any application-layer data exchange. They do not attempt to exploit vulnerabilities, log in, or access protected resources. To further minimize potential impact, we instrumented our tools to enforce a minimum delay of five minutes between successive handshakes to the same domain. Prior research has shown that exchanging TLS handshake messages is trivial [51] and that the vast majority of network operators do not regard such scanning as a significant threat [26].

We also followed community standards for ethical network measurement. Scans were conducted from dedicated IP addresses with Reverse DNS records pointing to a project webpage that described the study's purpose and provided an opt-out mechanism allowing operators to exclude their infrastructure from future scans. Our approach builds on established practices for minimizing disruption to Internet services and aligns with ethical frameworks for measurement research, such as the Menlo Report [10] and guidelines on good Internet citizenship [28].

Our methodology was reviewed by the Human Ethics Board of our institution, which confirmed that no personally identifiable information (PII) is collected and therefore no human subjects approval was required.

Finally, we emphasize the public-interest motivation of this work. TLS underpins the confidentiality and integrity of modern Internet communication, and understanding its preparedness for post-quantum transition has clear societal and security relevance. We therefore believe that the minimal risks associated with our non-invasive measurements are outweighed by the importance of the insights shared in this study.

A Policy Collection and Inclusion Criteria

We identify countries with demonstrated engagement in PQC, reflecting early technical investment and institutional preparedness for transition. Specifically, we include all countries that contributed submissions to NIST's Round 1 PQC standardization process, which comprised 82 proposals from 25 countries [16, 29].

We further expand this set by incorporating the world's largest economic blocs, as identified by World Bank GDP rankings [39], under the rationale that successful PQC adoption within these regions is likely to exert disproportionate global influence through market scale and infrastructure deployment.

Together, these criteria yield a shortlist of 31 countries, shown in Table 3.

A.1 Search Scope and Language

For each shortlisted country, we conduct web-based searches to identify PQC transition documents using structured keyword queries limited to English-language. These queries combine the *country name* with terms such as "post-quantum cryptography", "post-quantum cryptographic migration", "post-quantum cryptographic transition", and "post-quantum cryptography roadmap".

Table 3: Shortlisted countries and their regions

Region	Countries
North America	United States, Canada, Mexico
Europe	France, Germany, United Kingdom, Netherlands, Switzerland, Belgium, Spain, Italy, Denmark, Norway, Czech Republic, Russia, Ukraine
Asia	China, Japan, South Korea, Singapore, Taiwan, Israel, India, Indonesia, Turkiye, Saudi Arabia
Oceania	Australia, New Zealand
South America	Brazil, Argentina
Africa	Senegal

A.2 Authoritativeness Assessment

Only primary sources published by official institutions were eligible. We assessed authority using the following hierarchy:

(1) Binding law or regulation.
(2) Formal national mandate.
(3) Guidance or directives issued by a national cybersecurity or standards agency.
(4) Government strategy or policy guidance.
(5) Technical guidance published by another official public institution.

Publication by an official institution was not, by itself, treated as evidence that a document was legally binding. Vendor publications, news reports, commentary, and secondary summaries were excluded from the policy corpus.

A.3 Inclusion Criteria

A jurisdiction was included in the policy comparison only when its retained official source set satisfied all three of the following criteria:

(1) Explicitly addressed PQC transition and cryptographic interoperability.
(2) Defined a sectoral scope, deployment scope, or migration priority.
(3) Specified a PQC migration milestone, transition timeline, or date for moving away from classical public-key cryptography.

Requiring all three conditions excludes documents that merely acknowledge the quantum threat, describe individual algorithms, or express general support for future migration without defining an operational scope or timeline. Seven of the 31 shortlisted countries met all three criteria in the retained source corpus. The European Union was included as a separate regional policy context, producing the eight contexts reported in the policy comparison.

Table 4: TLS scanning outcomes for the three domain ranking datasets used in our robustness analysis.

Dataset	Domains Scanned	Successful	Fail
DomCop	1,000,000	714,067	285,933
Tranco	1,000,000	669,810	330,190
CrUX	999,888	773,332	226,556

Table 5: PQ-TLS adoption across domain ranking datasets (March 2026)

PQC KEM	Default			Support		
	DomCop	Tranco	CrUX	DomCop	Tranco	CrUX
Hybrid KEMs						
X25519MLKEM768	49.74%	51.55%	52.75%	49.75%	51.56%	52.75%
SecP256r1MLKEM768	0.00%	0.00%	0.00%	5.44%	4.07%	6.88%
SecP384r1MLKEM1024	0.00%	0.00%	0.00%	1.91%	0.40%	0.24%
Pure PQC KEMs						
ML-KEM-512	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
ML-KEM-768	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
ML-KEM-1024	0.00%	0.00%	0.00%	0.64%	0.12%	0.33%

B PQ-TLS Adoption Across Tranco & CrUX

To assess the robustness of our findings with respect to domain ranking sources, we replicate our measurements using two alternative widely used datasets, Tranco and Google CrUX. Table 4 and Table 5 summarize the scanning outcomes and resulting PQ-TLS adoption measurements across these domain lists.

Across all three datasets, PQ-TLS adoption rates and negotiated configurations remain highly consistent, confirming that our findings are not sensitive to the choice of domain ranking source.

C Configuration Attribution

This appendix documents two related but distinct attribution tasks. First, we associate each measured HTTPS endpoint with an infrastructure entity using network and application-layer evidence (§ C.1). Second, we estimate whether the endpoint's TLS configuration is more likely to be controlled by that infrastructure provider or by the domain operator (§ C.2).

C.1 Infrastructure Provider Attribution

To analyze the role of infrastructure platforms in PQ-TLS deployment, we associate each scanned domain with the infrastructure entity visible at its measured HTTPS endpoint. We apply a layered, first-match procedure using published provider address ranges, DNS records, reverse DNS, HTTP response metadata, WHOIS data, and autonomous system information.

We first test whether the resolved IP address falls within an address range published by a known content delivery network or cloud provider. For example, we use the address ranges published by Cloudflare[4] and Fastly[5]. A matching address is attributed to the provider that publishes the corresponding range.

[4]https://www.cloudflare.com/ips-v4
[5]https://api.fastly.com/public-ip-list

If no address-range match is found, we follow the domain's CNAME chain and compare each hostname and suffix against curated provider patterns. These patterns capture provider-controlled namespaces, such as `akamai.net`, `edgekey.net`, `cloudfront.net`, and `fastly.net`. If the CNAME chain remains inconclusive, we perform a PTR lookup on the resolved IP address and apply the same provider-pattern matching to the resulting hostname.

We next inspect application-layer metadata. We send a `HEAD` request to the resolved IP address using the target domain as the HTTPS SNI value and HTTP `Host` header. We attempt HTTPS first and use HTTP as a fallback when no usable HTTPS response is obtained. We compare the `Server` header and response-header names against curated platform signatures, including markers, such as `Server: cloudflare`, `CF-Ray`, and provider-specific load-balancer identifiers.

When these signals are inconclusive, we query Team Cymru's IP-to-ASN WHOIS service and examine both the returned ASN and organization name. Known ASNs are mapped directly to their associated providers, while organization names are compared against curated provider patterns. As a final fallback, we use the autonomous-system organization reported by the MaxMind GeoLite2 ASN database[6]. Because this fallback identifies the network organization rather than a specific managed service, it provides weaker evidence about the platform serving the domain.

If none of these stages yields an entity, we label the endpoint as *Unknown*. Unknown cases are retained when reporting attribution coverage but excluded from analyses that require a named infrastructure entity. This ordered procedure gives precedence to direct provider signals, such as published address ranges and provider-controlled DNS names, before relying on broader network-ownership evidence. The resulting attribution should therefore be interpreted as the infrastructure entity associated with the measured endpoint, not as definitive evidence of domain ownership or TLS-configuration control.

C.2 TLS Management Attribution Heuristics

Infrastructure attribution (§ C.1) identifies the network or platform associated with a measured HTTPS endpoint, but it does not reveal who controls the endpoint's TLS configuration. A domain using a cloud platform (e.g., Amazon Web Services) may inherit provider defaults, select among provider-supported options, or retain direct control over certificates and cryptographic settings. Because these administrative relationships are not publicly observable at Internet scale, we use a rule-based heuristic to estimate likely control of TLS configuration. We adapt the general approach of [51] and classify each domain as *Potentially Infrastructure Provider Managed*, *Potentially Owner Managed*, or *Unknown*.

The classification combines the domain name, the infrastructure entity identified in § C.1, the DNS CNAME chain, the HTTP `Server` header, and the common name of the leaf certificate. We also calculate two IP-fanout measures across the measured population: the number of distinct domains associated with each IP address and the number of distinct effective top-level domains plus one (eTLD+1) associated with that address. When a domain resolves to

[6]https://dev.maxmind.com/geoip/docs/databases/asn

Table 6: Grouped TLS-management attribution heuristics and final assignments for the stable panel (n = 684,494). Provider and Owner denote *Potentially Infrastructure Provider Managed* and *Potentially Owner Managed*, respectively.

Predicate	Example	Label	Domains
CDN identifier appears in the attributed entity, HTTP server value, certificate name, or CNAME metadata	Entity `CDN77` or CNAME `foo.cdn77.org`	Provider	3,179
Domain, attributed entity, or corroborating service metadata identifies a managed infrastructure platform	Entity `Cloudflare`, Google with domain `foo.blogspot.com`, Amazon with `Server: CloudFront`	Provider	310,056
More than 20 measured domains share an IP address	N/A	Provider	80,163
Tokens in the domain align with the attributed infrastructure entity	Domain `google.com` and entity `Google LLC`	Owner	6,558
HTTP server metadata identifies a managed platform	`Server: Google Frontend`, `Server: Vercel`, `Server: Netlify`	Provider	1,641
More than 10 distinct eTLD+1 names share an IP address	N/A	Provider	15,563
At most two distinct eTLD+1 names share an IP address and no provider identifier is present	N/A	Owner	188,162
CNAME metadata identifies a managed platform	`foo.cloudfront.net`, `foo.edgekey.net`	Provider	1,078
Certificate identity aligns with the domain or attributed entity under low-fanout conditions	`CN=*.example.com` for `www.example.com`	Owner	11,240
HTTP metadata identifies an edge service, load balancer, managed object service, or service mesh	`Server: awselb/2.0`, `Server: istio-envoy`	Provider	2,046
Additional DNS, certificate, or HTTP metadata identifies shared infrastructure	`foo.edgesuite.net`, `foo.trafficmanager.net`, `Server: Akamai Ghost`	Provider	2,943
No predicate provides sufficient management evidence	N/A	Unknown	61,865

multiple addresses, we use the maximum fanout observed across those addresses.

Table 6 summarizes the grouped attribution predicates and the number of domains assigned through each group. Closely related provider-specific conditions are combined to present their common evidentiary rationale without enumerating every provider name, service hostname, or metadata pattern. The examples therefore illustrate representative matches and are not exhaustive.

Provider-oriented predicates include explicit CDN identifiers, recognized infrastructure entities, provider-operated service domains, platform-specific HTTP responses, managed-platform CNAMEs, edge-infrastructure metadata, and high IP fanout. We treat an IP address shared by more than 20 measured domains, or by more than 10 distinct eTLD+1 names, as evidence consistent with shared infrastructure. Several predicates combine infrastructure attribution with DNS, HTTP, certificate, domain, or fanout evidence rather than relying on cloud-network attribution alone.

Owner-oriented predicates capture evidence more consistent with domain-specific operation. These signals include alignment between tokens in the domain and attributed infrastructure entity, low eTLD+1 fanout in the absence of provider identifiers, and domain-specific certificate or origin evidence under low-fanout conditions. These signals indicate weaker evidence of shared-platform management, but do not prove that the domain operator directly administers every TLS setting.

Applying these heuristics to the 684,494-domain stable panel classifies 416,669 domains as *Potentially Infrastructure Provider Managed*, 205,960 as *Potentially Owner Managed*, and 61,865 as *Unknown*. Unknown cases are retained when reporting classification coverage but excluded from comparisons that require a provider-managed or owner-managed label.

These labels represent inferred configuration-management relationships, not verified administrative ownership. The provider-managed label does not imply that the provider fixes every TLS parameter. Likewise, the owner-managed label does not imply that the domain operator owns or operates the underlying infrastructure. We therefore use these labels only to compare aggregate deployment patterns and assess their robustness under alternative attribution assumptions in § C.3.

C.3 Validation and Sensitivity Analysis

C.3.1 Manual Validation. To assess agreement between the automated attribution procedure and manual assessment, we examined separate random samples for infrastructure attribution and TLS-management attribution. For each sampled domain, we reviewed evidence independent of the final classification label, including CNAME records, PTR records, HTTP metadata, certificate identities, ASN and RIR information, published provider address ranges, and provider service documentation.

For infrastructure attribution, the automated and manual assessments agreed for 196 of 200 sampled domains, corresponding to an agreement rate of 98.0% and a Wilson 95% confidence interval of approximately 95.0% to 99.2%. The four disagreements involved ambiguous or overlapping infrastructure identifiers, particularly among smaller regional hosting providers.

For TLS-management attribution, the automated and manual assessments agreed for 189 of 200 sampled domains, corresponding to an agreement rate of 94.5% and a Wilson 95% confidence interval of approximately 90.4% to 96.9%. The disagreements primarily involved shared infrastructure on which customers appeared to

Table 7: Sensitivity of March 2026 default-client hybrid key-exchange selection to TLS-management attribution assumptions. The gap is the difference between provider rate and owner rate in percentage points.

Scenario	Provider n	Provider (%)	Owner n	Owner (%)	Unknown n	Unknown (%)	Gap (pp)
Baseline	416,669	75.94	205,960	7.49	61,865	8.21	68.45
No fanout-dependent predicates	350,770	83.75	8,587	23.78	325,137	12.64	59.97
All provider predicates evaluated before owner predicates	420,570	75.62	202,059	6.84	61,865	8.21	68.78
All owner predicates evaluated before provider predicates	409,644	76.18	212,985	9.29	61,865	8.21	66.89
Fanout thresholds multiplied by 0.5	442,905	72.26	158,168	7.37	83,421	6.27	64.89
Fanout thresholds multiplied by 2	400,507	77.99	243,103	7.93	40,884	12.97	70.06
All unknown cases assigned to owner	416,669	75.94	267,825	7.65	0	N/A	68.29
All unknown cases assigned to provider	478,534	67.18	205,960	7.49	0	N/A	59.70
Cloudflare-attributed endpoints excluded	212,671	57.94	205,960	7.49	61,865	8.21	50.45

retain partial control over TLS configuration. This ambiguity is intrinsic to distinguishing provider-managed defaults from customer-configurable services using externally observable evidence.

These agreement rates assess correspondence with our manual interpretation of the available evidence. They do not establish administrative ground truth. We therefore retain the qualifiers *Potentially* Infrastructure Provider Managed and *Potentially* Owner Managed.

C.3.2 Sensitivity Analysis. Manual validation may not capture systematic errors caused by fanout thresholds, overlapping signals, treatment of unknown cases, or concentration within a major provider. We therefore recompute the March 2026 stock-client hybrid key-exchange rates under alternative classification assumptions. These tests remove fanout-based predicates, vary fanout thresholds, change the treatment of overlapping and unknown cases, and exclude Cloudflare-attributed endpoints.

Under the baseline classification, 316,419 of 416,669 potentially provider-managed endpoints select hybrid key exchange (75.94%), compared with 15,420 of 205,960 potentially owner-managed endpoints (7.49%). This corresponds to a difference of 68.45 percentage points. As shown in Table 7, the difference remains positive under every tested classification variant, ranging from 59.70 to 70.06 percentage points. Removing fanout-dependent predicates moves many domains into the *Unknown* category and reduces the difference to 59.97 percentage points. As a separate concentration check, excluding all Cloudflare-attributed endpoints produces the smallest observed difference, at 50.45 percentage points.

These results show that the aggregate association persists across the tested fanout thresholds, treatments of overlapping evidence, assignments of unknown cases, and the exclusion of Cloudflare.

D Sectoral Distribution of PQ-TLS Adoption

To provide a complete view of PQ-TLS deployment across the dataset, we examine adoption across all SafeDNS domain categories. SafeDNS classifies domains according to the primary type of content or service they host, enabling a broad characterization of sectoral deployment patterns beyond the policy-relevant sectors discussed in the main text.

Figure 5 presents PQ-TLS default adoption rates across all SafeDNS categories observed in the stable panel for the July 2025, November 2025, and March 2026 measurement rounds. Each column corresponds to a category, with the sample size of domains in that category indicated in parentheses. Colors represent the percentage of domains within each category that negotiate PQ-TLS by default.

Note: The domain counts shown in Figure 5 include all domains assigned to each SafeDNS category. In contrast, the sectoral analysis presented in Section 6.3 considers only domains for which TLS management attribution could be determined, excluding those labeled as *Unknown*. Consequently, category totals may appear marginally larger in this figure.

E Protocol-level Security Depth: Labeling Criteria

This appendix documents the labeling criteria and test semantics used in our protocol-level security depth analysis (RQ4). Our goal is to characterize the *configuration surface* exposed by operational TLS deployments rather than to assess exploitability or conduct adversarial testing.

All measurements were conducted using `testssl.sh`, restricted to non-intrusive probes that do not modify server state, attempt credential extraction, or trigger repeated handshake failures. To further minimize operational impact, we enforce a minimum delay of five minutes between successive scans of the same domain.

E.1 Protocol Version Support

A domain is labeled as supporting an obsolete protocol version if it *offers* any of the following during protocol negotiation, regardless of default preference:

- SSLv2
- SSLv3
- TLS 1.0
- TLS 1.1

Support is detected through passive protocol negotiation and version advertisement. Domains that prefer TLS 1.3 but retain support for older versions are still classified as offering obsolete protocols, as such configurations may preserve downgrade paths.

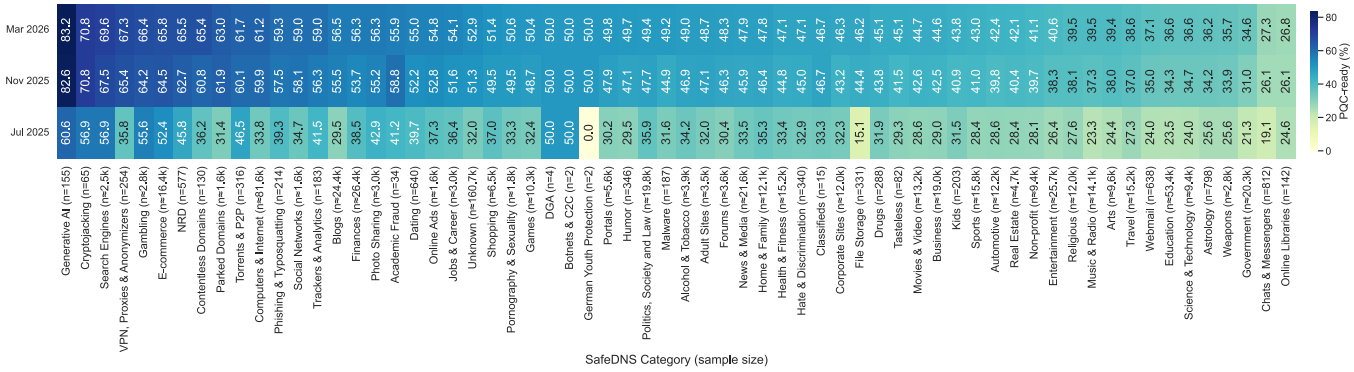

Figure 5: PQ-TLS Adoption Trends Across All SafeDNS Categories

E.2 Cipher Suite Classification

Cipher-related labels are assigned based on advertised cipher suite support according to OpenSSL[7] categorisation. A domain is flagged if it offers any cipher suites falling into the following categories:

- **NULL Ciphers**: Provide no encryption.
- **aNULL Ciphers**: Provide encryption but no authentication, enabling man-in-the-middle attacks.
- **Export Ciphers**: Suites with intentionally weakened 40–56 bit keys, introduced for export control compliance and now trivially broken.
- **Low Ciphers**: Use outdated 64-bit block ciphers.
- **Obsoleted Ciphers**: Suites formally deprecated by standards bodies or removed from modern TLS protocols.
- **3DES/IDEA Ciphers**: Legacy block ciphers with outdated constructions or limited block sizes.
- **Strong-NoFS Ciphers**: Suites that employ strong encryption but lack forward secrecy, exposing past sessions if long-term keys are ever compromised.

Cipher availability is evaluated independently of negotiation preference.

E.3 Known TLS Vulnerability Indicators

We flag domains for *potential exposure indicators* associated with well-known TLS weaknesses. These labels do not imply exploitability, only that the relevant protocol features are present. Indicators include:

- POODLE: SSLv3 with CBC cipher support.
- SWEET32: support for 64-bit block ciphers.
- FREAK: EXPORT-grade RSA cipher support.
- BEAST: TLS 1.0 with CBC ciphers (even if higher versions are also supported).
- Lucky13: support for CBC-mode ciphers under TLS.
- RC4: RC4 cipher support.
- LOGJAM: DH EXPORT cipher support.
- BREACH: HTTP compression enabled, detected via advertised content-encoding headers on a single benign request to the root path.

All vulnerability indicators are evaluated conservatively and without active exploitation attempts. For compression-based indicators (e.g., BREACH), only passive detection of compression support is performed.

E.4 Certificate Properties

Certificate-related issues are identified using the following criteria:

- **Certificate expiration**: the leaf certificate is expired at scan time.
- **Hostname mismatch**: the certificate does not correctly match the supplied domain name, including cases where matching relies solely on CN or wildcard behavior.
- **Chain of trust issues**: including incomplete chains, expired intermediates, self-signed certificates, or self-signed certificate authorities within the chain.

E.5 OCSP Stapling

OCSP stapling is evaluated based on server behavior during the TLS handshake. Domains are classified as not offering OCSP stapling if:

- No stapled OCSP response is provided, or
- The stapled response contains an OCSP error (e.g., `unauthorized`).

E.6 Interpretation

Labels reflect the presence of legacy or potentially risky protocol features, not confirmed vulnerabilities. Many flagged conditions coexist with TLS 1.3 and modern cipher preferences, indicating residual compatibility rather than misconfiguration specific to PQC. As such, results should be interpreted as indicators of protocol hygiene rather than security failures.

[7] https://docs.openssl.org/master/man1/openssl-ciphers/