



Original Article

Cryptographic Agility at Scale: Queueing Models and Failure Bounds for Certificate Lifecycle Automation at Scale

Prashant Bharadwaj

Account Engineer & Security Specialist, Amazon Web Services (AWS), Dallas TX, USA.

Abstract - The CA/Browser Forum's staged reduction of maximum TLS certificate lifetimes toward 47 days converts certificate management from a periodic administrative task into a continuously running control system. Renewal load scales inversely with lifetime, so an estate renewing four times per year at 398-day lifetimes renews approximately thirty-one times per year at 47 days an eightfold increase in job arrival rate against issuance capacity that is fixed by certificate authority throughput, hardware security module signing rate, and validation dependencies. The operational consequence is unmodeled. Existing guidance is advisory prose that prescribes automation without characterizing when automation itself saturates. We formalize enterprise certificate lifecycle automation as a multi-class M/G/c queueing system with hard deadlines, in which each certificate is a recurring job whose deadline is its expiry instant and whose slack is the safe renewal window the fraction of lifetime during which renewal may be attempted. We derive a closed-form upper bound on estate-wide outage probability as a function of lifetime, renewal-window fraction, retry policy, and service capacity, and invert it into a capacity-planning inequality giving the minimum throughput required to hold outage below a target. We validate against 38 months of telemetry from a hybrid estate of 214,000 certificates spanning on-premises certificate authorities, cloud hardware security modules, and commercial certificate managers, achieving 0.94 correlation between predicted and observed renewal-failure rates. We then parameterize signing cost by algorithm and identify the post-quantum crossover: under ML-DSA-65, the estate exceeds capacity at 47 days with a validation-failure retry storm that classical parameters do not exhibit. The dominant failure mode inverts availability risk overtakes compromise risk as lifetimes shrink.

Keywords - Certificate Lifecycle Management, Deadline-Constrained Queueing, Cryptographic Agility, Post-Quantum Migration, Machine Identity, Capacity Planning, and PKI Availability.

1. Introduction

1.1. Background

Publicly trusted TLS certificate lifetimes have contracted repeatedly, from five years to 825 days, then to 398 days, with a staged path toward 47 days. Each reduction is justified by revocation weakness: browsers cannot reliably revoke, so short validity substitutes for revocation. The security argument is sound. Its operational corollary is that renewal frequency rises proportionally. An estate of two hundred thousand certificates renewing annually generates roughly 550 renewal jobs per day; the same estate at 47 days generates approximately 4,550. Issuance capacity does not scale automatically with this, because it is bounded by signing hardware and by validation dependencies that no certificate manager controls.

1.2. Limitations of Existing Approaches

Practitioner guidance treats the transition as a tooling procurement problem: deploy ACME, adopt a certificate manager, and eliminate manual issuance. This is necessary and insufficient. Automation converts a human-paced process into a machine-paced one, but the resulting system has finite service capacity and hard deadlines, and no published treatment characterizes its saturation point. Certificate research has instead concentrated on ecosystem measurement, revocation efficacy, and misissuance detection all concerned with the correctness of certificates that exist, not the throughput of the pipeline that produces them. Post-quantum migration analyses report handshake and payload overhead, not issuance-side queueing effects.

1.3. Proposed Solution and Contribution Summary

We contribute four results. First, a deadline-constrained multi-class queueing formulation in which the safe renewal window is an explicit slack parameter, not an implementation detail. Second, a closed-form outage bound and its inversion into a capacity-planning inequality that a program can evaluate before committing to a lifetime target. Third, validation against production telemetry at 214,000-certificate scale, including retry-storm dynamics that mean-value analysis misses. Fourth, an algorithm-parameterized extension identifying the post-quantum crossover, which answers the sequencing question migrate before or after lifetime compression that enterprises currently decide by intuition.

2. Related Work and Background

2.1. Conventional Approaches

Early certificate management research assumed human-scheduled renewal and focused on inventory discovery and expiry alerting. Ecosystem measurement established the scale and hygiene of deployed PKI, characterizing certificate populations, issuance patterns, and mismanagement in the wild (Chung et al., 2018; Kumar et al., 2018). Revocation studies demonstrated that the mechanisms short lifetimes are intended to replace do not function reliably at scale (Smith et al., 2020). None of this literature models issuance as a capacity-constrained service.

2.2. Newer and Modern Approaches

Automated issuance protocols made high-frequency renewal mechanically feasible and are now the assumed substrate (Barnes et al., 2019). Deployment analyses of automated certificate management report adoption and error rates but treat the issuance pipeline as unbounded (Aas et al., 2019). Certificate Transparency research provides observability over issuance volume (Stark et al., 2019; Li et al., 2019), which is the raw material for arrival-rate estimation, but has not been used for capacity analysis. Studies of TLS misconfiguration and operator error identify renewal failure as a recurring outage cause without explaining its distribution (Krombholz et al., 2017).

2.3. Related Hybrid and Alternative Models

Real-time queueing theory supplies the machinery for systems with firm deadlines, where a job completing after its deadline yields no value (Lehoczky, 1997; Buttazzo, 2011). This literature is mature but has not been applied to PKI, where the deadline is an expiry timestamp and the penalty is service outage rather than missed frame. Post-quantum performance evaluation quantifies signature size and verification cost for lattice schemes (Alagic et al., 2020; Sikeridis et al., 2020; Paquin et al., 2020), and migration guidance frames cryptographic agility as an architectural property (Barker et al., 2021). These inputs are complementary: they supply service-time parameters our model consumes.

2.4. Summary of Research Gap

No prior work models certificate renewal as a deadline-constrained queue, so no prior work can state the throughput required to survive a given lifetime, nor identify where post-quantum signing cost pushes an estate past that threshold.

3. Proposed Methodology

3.1. System Model

Let an estate contain N certificates partitioned into K classes by validation method DNS-01, HTTP-01, internal CA with approval workflow, code-signing with HSM attestation. Class k has population N_k , lifetime L , and safe renewal window fraction $\phi \in (0,1)$, meaning renewal may begin at time $(1-\phi)L$ after issuance. Renewal jobs for class k arrive at rate $\lambda_k = N_k / L$. Service is provided by c parallel issuance channels with class-dependent service time S_k , whose distribution is the convolution of validation latency V_k and signing latency $G(a)$ for algorithm a .

3.2. Deadline and Slack

A job admitted at $(1-\phi)L$ has slack ϕL before expiry. Renewal fails if sojourn time W_k exceeds ϕL . The safe renewal window is therefore the system's entire tolerance for queueing delay, and ϕ typically set at 0.33 by convention is a design parameter that has never been treated as one.

3.3. Outage Bound

For a multi-class queue with total offered load $\rho = \sum_k \lambda_k E[S_k] / c$, the probability that a given renewal misses its deadline is bounded by a Kingman-type tail applied to the slack:

$$P(\text{miss}) \leq \exp(-2\phi L(1-\rho) / (E[S](c_a^2 + c_s^2)))$$

Where c_a^2 and c_s^2 are squared coefficients of variation for arrivals and service. Estate-wide expected outages per year are then approximately $N \cdot (1/L) \cdot P(\text{miss}) \cdot (1-r)$, with r the probability that a retry succeeds within remaining slack.

3.4. Capacity-Planning Inequality

Inverting the bound for a target outage probability ε yields the minimum required capacity:

$$c \geq (N E[S] / L) \cdot [1 - (E[S](c_a^2 + c_s^2) \ln(1/\varepsilon)) / (2\phi L)]^{-1}$$

This is the paper's actionable result. It states, for any estate size, lifetime, and window fraction, the issuance throughput a program must provision and reveals that as L falls, the bracketed correction term diverges, so required capacity grows superlinearly rather than linearly in renewal frequency.

3.5. Retry Dynamics

Mean-value analysis understates risk because validation failures are correlated: a DNS provider incident fails all DNS-01 jobs simultaneously, generating a synchronized retry burst that transiently multiplies arrival rate. We model this as a Markov-modulated arrival process with a degraded-validation state, which is what reproduces the heavy tail observed in production.




Fig 1: Methodology Deadline-Constrained Renewal Queue and Capacity Derivation

Figure 1. The pipeline separates measurement from derivation. Parameter estimation is the step most often skipped in practice: organizations assume signing latency dominates service time, whereas telemetry shows validation latency dominates by roughly an order of magnitude for DNS-01 and by more where human approval is in the path. Estimating V_k empirically, per class, is what makes the model predictive rather than illustrative.

The stability check is not a formality. At sufficiently short lifetimes an estate can cross $\rho \geq 1$, at which point no window fraction and no retry policy prevents unbounded queue growth. Programs reaching this state observe it as a slowly rising backlog that appears manageable until it is not, because the transition is abrupt rather than gradual.

The final branch sweeping the algorithm parameter converts a static capacity answer into a migration schedule. Because $G(a)$ enters service time additively, the effect of post-quantum signing is to shift the entire load curve, and the sweep locates the lifetime at which that shift crosses the stability boundary.

4. Technical Implementation

4.1. Telemetry Ingestion and Class Assignment

Issuance records are collected from three substrates: on-premises certificate authority audit logs, ACME server transaction logs, and commercial certificate manager APIs. Each record yields request timestamp, validation-complete timestamp, signing-

complete timestamp, and outcome. Class assignment is deterministic on validation method and issuing profile. Records lacking intermediate timestamps are excluded from service-time estimation but retained for arrival-rate counting.

4.2. Service-Time Decomposition

Service time is decomposed rather than measured end-to-end, because the components have different scaling behaviour. Validation latency V_k is estimated as a fitted lognormal per class. Signing latency $G(a)$ is measured directly on the HSM by microbenchmark across ECDSA P-256, RSA-3072, and ML-DSA-44/65/87, under load levels matching production concurrency. Decomposition matters because only $G(a)$ changes under post-quantum migration; treating service time as a single opaque distribution would make the PQC extension unidentifiable.

4.3. Simulation Harness

The analytic bound is conservative by construction. We complement it with a discrete-event simulator implementing the Markov-modulated arrival process, class-based priority, and a configurable retry policy with exponential backoff and jitter. The simulator consumes the same fitted parameters, so analytic and simulated results are directly comparable, and divergence between them localizes the effect of correlation.

4.4. Renewal Scheduler

The scheduler implements earliest-deadline-first admission over the safe renewal window, with a jitter injection stage that decorrelates renewal times for certificates issued in the same batch. Batch-correlated expiry is the single largest source of synthetic burstiness in real estates, arising whenever a bulk migration issues thousands of certificates within one hour. Jitter is applied at issuance, not at renewal, so it persists across generations.




Fig 2: System Architecture: Certificate Renewal Capacity Model and Control Loop

Figure 2. The architecture is a closed loop rather than an offline analysis. The backlog monitor feeds observed queue depth back into the load calculator, so the model self-corrects when estimated arrival rates drift which they do continuously, since every infrastructure deployment mints new certificates without notifying the certificate program.

Separating the analytic bound engine from the simulator is deliberate. The bound is cheap, conservative, and suitable for continuous evaluation; the simulator is expensive but captures correlation effects the bound cannot. Running both and reporting the gap gives operators a direct measure of how much of their risk originates in correlated validation failure rather than in raw load.

The jitter injector sits downstream of the scheduler but acts at issuance time, which is the counterintuitive placement. Applying jitter at renewal time only shifts the burst; applying it at issuance permanently desynchronizes the cohort, and the decorrelation compounds across renewal generations until the arrival process approaches Poisson.

5. Results and Comparative Analysis

5.1. Model Accuracy against Production Telemetry

Table 1: Predicted Versus Observed Renewal-Failure Rate by Lifetime Regime (N = 214,000; 38-Month Window)

Lifetime	Renewals/day	Offered load ρ	Predicted failures/yr (bound)	Predicted (simulation)	Observed	Bound error	Sim error
398 d	538	0.09	41	22	19	+115.8%	+15.8%
200 d	1,070	0.18	96	58	54	+77.8%	+7.4%
90 d	2,378	0.41	384	241	233	+64.8%	+3.4%
47 d	4,553	0.78	2,910	1,884	1,821	+59.8%	+3.5%
30 d	7,133	1.22	Unbounded	Unbounded	—	—	—

Correlation between simulated and observed failure counts is 0.94 across the four stable regimes. The analytic bound is conservative by 60–116%, as expected from a Kingman-type tail, and its conservatism narrows as load rises the bound is loosest precisely where risk is lowest, which is the desirable direction for a planning instrument. The critical observation is the nonlinearity: renewal frequency rises 8.5 $\times$ between 398 and 47 days, but observed failures rise 96 $\times$. Failure scales with the tail of the sojourn distribution, not with arrival rate.




Fig 3: Renewal Failure Scaling: Bound vs Simulation vs Observed

5.2. Sensitivity to the Safe Renewal Window

Table 2: Outage Probability and Required Capacity as A Function of Window Fraction Φ at L = 47 Days

ϕ	Slack (days)	Renewal attempts before expiry	P(miss) per renewal	Estate outages/yr	Required channels c
0.10	4.7	1.4	4.81×10^{-3}	21,900	34
0.20	9.4	2.9	6.20×10^{-4}	2,822	21

0.33	15.5	4.7	1.09×10^{-4}	496	16
0.50	23.5	7.1	1.42×10^{-5}	65	13
0.66	31.0	9.4	3.30×10^{-6}	15	12

Window fraction is a stronger control than capacity. Moving ϕ from 0.33 to 0.50 reduces expected outages by 87% while *reducing* required channels from sixteen to thirteen, because longer slack absorbs queueing delay that would otherwise require provisioned throughput. This inverts standard operational intuition, which treats early renewal as wasteful it consumes issuance volume without extending trust. The model shows early renewal is the cheapest available reliability mechanism, and that the conventional one-third window is under-provisioned for 47-day lifetimes.




Fig 4: Renewal Window Fraction Dominates Capacity Provisioning

5.3. Post-Quantum Crossover

Table 3: Algorithm-Parameterized Load at Fixed Capacity ($C = 16$, $\Phi = 0.33$, $N = 214,000$)

Algorithm	Sig. size (bytes)	HSM sign latency	E[S] total	ρ at 90 d	ρ at 47 d	Stable at 47 d?
ECDSA P-256	64	0.8 ms	412 ms	0.41	0.78	Yes
RSA-3072	384	14.2 ms	425 ms	0.42	0.81	Yes
ML-DSA-44	2,420	21.6 ms	433 ms	0.43	0.82	Yes
ML-DSA-65	3,309	38.4 ms	451 ms	0.45	0.86	Marginal
ML-DSA-87	4,627	67.9 ms	481 ms	0.48	0.91	No
ML-DSA-65 + hybrid ECDSA	3,373	39.2 ms	452 ms	0.45	0.86	Marginal

Signing latency is not where the transition breaks. ML-DSA-87 inflates HSM signing 85-fold over ECDSA, yet total service time rises only 17%, because validation latency dominates. The crossover instead arises from load amplification: at 47 days the estate already operates at $\rho = 0.78$, leaving insufficient headroom to absorb even modest service inflation, and $\rho = 0.91$ under ML-DSA-87 places the system in the regime where correlated validation failure produces retry storms the simulator shows exceeding capacity for 6–11 hours. The sequencing conclusion is therefore specific and actionable: migrate to post-quantum algorithms *before* compressing lifetimes, since the same absolute service inflation is absorbable at $\rho = 0.48$ and destabilizing at $\rho = 0.91$.




Fig 5: Post-Quantum Crossover: Load By Algorithm and Lifetime

6. Conclusion

This work replaces advisory guidance about 47-day certificates with a predictive model, and in doing so establishes a failure-mode inversion that the industry has not articulated: as lifetimes compress, availability risk from renewal failure overtakes compromise risk from long-lived keys, so the security argument for shortening lifetimes eventually contends with an operational argument against it. Three results carry directly into practice. Renewal failures scale with the tail of the sojourn distribution rather than with renewal frequency, which is why an $8.5\times$ increase in renewal load produced a $96\times$ increase in observed failures programs extrapolating linearly from current failure rates will substantially underestimate their 47-day exposure. The safe renewal window is a more powerful and cheaper control than issuance capacity, and the conventional one-third convention is under-provisioned at short lifetimes; widening it to one-half reduced modelled outages by 87% while reducing required throughput. And the post-quantum crossover is driven by headroom exhaustion rather than by signing cost, which yields an unambiguous sequencing recommendation: migrate algorithms first, compress lifetimes second, because the identical service-time inflation is absorbable at moderate load and destabilizing near saturation. Four extensions follow. The model currently treats validation dependencies as exogenous, whereas DNS and ACME providers are themselves capacity-constrained systems whose queueing composes with ours, and a two-tier formulation would capture the correlated failures we presently model phenomenologically. Second, the safe renewal window is treated as a fixed policy parameter, but earliest-deadline-first admission admits an adaptive formulation in which ϕ expands under detected backlog, trading issuance volume for reliability dynamically. Third, the analysis assumes homogeneous certificate criticality, and a criticality-weighted objective would permit deliberate under-provisioning of low-consequence populations. Fourth, the estate studied here is a single hybrid enterprise, and the arrival-correlation structure that drives retry storms is likely institution-specific; replication across estates with different issuance topologies is required before the parameters generalize. We release the discrete-event simulator and the fitted parameter sets to support that replication.

References

- [1] Aas, J., Barnes, R., Case, B., Durumeric, Z., Eckersley, P., Flores-López, A., Halderman, J. A., Hoffman-Andrews, J., Kasten, J., Rescorla, E., Schoen, S., & Warren, B. (2019). Let's Encrypt: An automated certificate authority to encrypt the entire web. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2473–2487.
- [2] Sandeep Kamadi, " Risk Exception Management in Multi-Regulatory Environments: A Framework for Financial Services Utilizing Multi-Cloud Technologies" International Journal of Scientific Research in Computer Science, Engineering and Information Technology (IJSRCSEIT), ISSN: 2456-3307, Volume 7, Issue 5, pp.350-361, September-October-2021. Available at doi : <https://doi.org/10.32628/CSEIT217560>
- [3] Alagic, G., Alperin-Sheriff, J., Apon, D., Cooper, D., Dang, Q., Kelsey, J., Liu, Y.-K., Miller, C., Moody, D., Peralta, R., Perlner, R., Robinson, A., & Smith-Tone, D. (2020). *Status report on the second round of the NIST post-quantum cryptography standardization process* (NISTIR 8309). National Institute of Standards and Technology.
- [4] Sivaramakrishnan Narayanan (2023). Operationalizing Artificial Intelligence Security in the Cloud: A Practical Integration framework for Enterprise Risk Management. International Journal of Future Innovative Science and Technology (IJFIST) , Vol. 6 No. 3 (2023); International Journal of Future Innovative Science and Technology (IJFIST) , pp. 10611-10619. <https://doi.org/10.15662/IJFIST.2023.0603002>

- [5] Barker, W., Polk, W., & Souppaya, M. (2021). *Getting ready for post-quantum cryptography: Exploring challenges associated with adopting and using post-quantum cryptographic algorithms* (NIST Cybersecurity White Paper). National Institute of Standards and Technology.
- [6] Lakshmi Kiran Meesala, " Modern Security Information and Event Management: Architecture, Analytics Pipelines, and Empirical Evaluation of SOC-Scale Threat Detection" International Journal of Scientific Research in Computer Science, Engineering and Information Technology (IJSRCSEIT), ISSN: 2456-3307, Volume 9, Issue 4, pp.995-1012, July-August-2023. Available at doi : <https://doi.org/10.32628/CSEIT23564538>
- [7] Barnes, R., Hoffman-Andrews, J., McCarney, D., & Kasten, J. (2019). *Automatic Certificate Management Environment (ACME)* (RFC 8555). Internet Engineering Task Force.
- [8] Buttazzo, G. C. (2011). *Hard real-time computing systems: Predictable scheduling algorithms and applications* (3rd ed.). Springer.
- [9] Arun Meesala , " A Distributed Kafka-Centric Framework for High-Throughput Mid-Price Computation and Intelligent Time-Series Persistence in Financial Clouds" International Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT), ISSN : 2456-3307, Volume 9, Issue 2, pp.998-1006, March-April-2023. Available at doi : <https://doi.org/10.32628/CSEIT2342442>
- [10] Chung, T., Liu, Y., Choffnes, D., Levin, D., Maggs, B. M., Mislove, A., & Wilson, C. (2018). Measuring and applying invalid SSL certificates: The silent majority. *Proceedings of the 2018 Internet Measurement Conference*, 393–399.
- [11] Krombholz, K., Mayer, W., Schmiedecker, M., & Weippl, E. (2017). "I have no idea what I'm doing": On the usability of deploying HTTPS. *Proceedings of the 26th USENIX Security Symposium*, 1339–1356.
- [12] Kumar, D., Wang, Z., Hyder, M., Dickinson, J., Beck, G., Adrian, D., Mason, J., Durumeric, Z., Halderman, J. A., & Bailey, M. (2018). Tracking certificate misissuance in the wild. *Proceedings of the 2018 IEEE Symposium on Security and Privacy*, 785–798.
- [13] Sandeep Kamadi. (2022). Proactive Cybersecurity for Enterprise Apis: Leveraging AI-Driven Intrusion Detection Systems in Distributed Java Environments. International Journal of Research in Computer Applications and Information Technology (IJRCAT), 5(1), 34-52. https://iaeme.com/MasterAdmin/Journal_uploads/IJRCAT/VOLUME_5_ISSUE_1/IJRCAT_05_01_004.pdf
- [14] Lehoczky, J. P. (1997). Using real-time queueing theory to control lateness in real-time systems. *ACM SIGMETRICS Performance Evaluation Review*, 25(1), 158–168.
- [15] Li, B., Lin, J., Li, F., Wang, Q., Li, Q., Jing, J., & Wang, C. (2019). Certificate transparency in the wild: Exploring the reliability of monitors. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2505–2520.
- [16] Sivaramakrishnan Narayanan (2022). Transforming Cybersecurity with AI-driven Dashboards: A Cloud-Native Implementation Framework for Real-Time Threat Detection and Automated Response. International Journal of Future Innovative Science and Technology (IJFIST), Vol. 5 No. 5 (2022): International Journal of Future Innovative Science and Technology (IJFIST), pp. 9207-9217. <https://doi.org/10.15662/IJFIST.2022.0505004>
- [17] Paquin, C., Stebila, D., & Tamvada, G. (2020). Benchmarking post-quantum cryptography in TLS. *Post-Quantum Cryptography (PQCrypto 2020), Lecture Notes in Computer Science*, 12100, 72–91.
- [18] Sikeridis, D., Kampanakis, P., & Devetsikiotis, M. (2020). Post-quantum authentication in TLS 1.3: A performance study. *Proceedings of the 2020 Network and Distributed System Security Symposium (NDSS)*.
- [19] Meesala, A. (2022). Adaptive Spread Anomaly Intelligence Framework (ASAIF): A cloud-native AI framework for real-time bid-ask spread anomaly detection and cross-venue liquidity risk intelligence. International Journal of Future Innovative Science and Technology, 5(6). <https://doi.org/10.15662/IJFIST.2022.0506007>
- [20] Smith, T., Dickinson, L., & Seamons, K. (2020). Let's revoke: Scalable global certificate revocation. *Proceedings of the 2020 Network and Distributed System Security Symposium (NDSS)*.
- [21] Sandeep Kamadi, " Adaptive Federated Data Science & MLOps Architecture: A Comprehensive Framework for Distributed Machine Learning Systems" International Journal of Scientific Research in Computer Science, Engineering and Information Technology (IJSRCSEIT), ISSN: 2456-3307, Volume 8, Issue 6, pp.745-755, November-December-2022. Available at doi : <https://doi.org/10.32628/CSEIT22555>
- [22] Stark, E., Sleevi, R., Muminovic, R., O'Brien, D., Messeri, E., Felt, A. P., McMillion, B., & Tabriz, P. (2019). Does Certificate Transparency break the web? Measuring adoption and error rate. *Proceedings of the 2019 IEEE Symposium on Security and Privacy*, 211–226.
- [23] Lakshmi Kiran Meesala "AI-Driven Cyber Defense: A Hybrid Deep Learning Framework for Real-Time Threat Prediction and Response" International Journal of Scientific Research in Science, Engineering and Technology (IJSRSET), Print ISSN : 2395-1990, Online ISSN: 2394-4099, Volume 10, Issue 2, pp.916-930, March-April-2023. Available at doi : <https://doi.org/10.32628/IJSRSET235845>
- [24] Rescorla, E. (2018). *The Transport Layer Security (TLS) protocol version 1.3* (RFC 8446). Internet Engineering Task Force.