FORGE QBit: A Heterogeneous Post-Quantum Security and Identity Engine for Agent-Legible Operating Systems

A FORGE OS Subsystem Specification

577 Industries R&D Lab
577 Industries Incorporated
research@577industries.com

Abstract

The convergence of quantum computing, physics-informed artificial intelligence, and post-quantum cryptography creates an opportunity to build intelligence systems that are simultaneously more capable, more grounded in physical reality, and more secure than conventional approaches permit. We present FORGE QBIT, the security and identity engine of FORGE OS—the immune system of the agent-legible operating system. FORGE QBIT comprises four tightly integrated modules: (1) a *heterogeneous Crypto Core* implementing post-quantum cryptographic primitives across three independent mathematical families—lattice-based (ML-KEM, ML-DSA), code-based (HQC), and hash-based (SLH-DSA, Falcon)—with an Agility Controller providing automatic cross-family rotation on cryptanalytic threat detection; (2) *PhysicsCore*, a physics-informed neural network engine embedding governing partial differential equations directly into network loss functions, achieving high-fidelity surrogate modeling with 10–50$\times$ less training data than data-driven methods; (3) *QuantumSolve*, a hardware-agnostic hybrid quantum-classical optimization framework executing QAOA and VQE workloads across IBM Quantum, Amazon Braket, IonQ, Rigetti, and D-Wave backends; and (4) an *Identity Spine* service issuing capability-attenuated X.509 certificates to all FORGE OS agents, serving as the root-of-trust for the entire platform. Additionally, FORGE QBIT introduces an HSM Hierarchical Key Enclave achieving FIPS 140-3 compliance at software-level latency ($<$1ms key derivation), and a PQ Double Ratchet protocol providing forward-secret, quantum-resistant communications for edge and contested environments. Experimental evaluation demonstrates 19–42% improvement over classical baselines on combinatorial optimization, 94.3% physical consistency in surrogate models, quantum-resistant key exchange with less than 12% latency overhead, cross-family rotation convergence within 12.4 seconds, and certificate issuance throughput of 8,400 certificates per second. We validate through deployment case studies spanning U.S. Army logistics optimization, investment portfolio construction, and naval communications hardening.

1 Introduction

1.1 The Convergence of Three Computational Challenges

Modern defense and financial intelligence systems face a convergence of three computational challenges that collectively exceed the capabilities of any single technological paradigm. First,

mission-critical decisions—from multi-asset portfolio construction under tail-risk constraints to multi-domain military logistics under adversarial uncertainty—require solving combinatorial optimization problems whose solution spaces grow exponentially with problem dimensionality. Second, the physical phenomena underlying sensing, simulation, and prediction tasks are governed by partial differential equations (PDEs) that conventional neural networks approximate only statistically, without guarantees of physical consistency. Third, the accelerating maturation of quantum computing threatens to break the cryptographic infrastructure protecting sensitive communications within a decade, even as quantum algorithms promise to solve the very optimization problems that classical methods find intractable.

Within FORGE OS, these three challenges are the responsibility of a single subsystem: FORGE QBIT, the platform's immune system. Just as a biological immune system provides both defensive capabilities (pathogen recognition, antibody production) and systemic infrastructure (self/non-self discrimination, immune memory), FORGE QBIT provides both computational capabilities (physics-informed optimization, quantum-accelerated solving) and platform infrastructure (post-quantum security, cryptographic identity).

1.2 The Agent Identity Crisis

Current AI agents operate without cryptographic identity. An agent invoking a foundation model presents an API key—a shared secret that identifies the *organization*, not the *agent*. There is no authentication of the agent itself, no non-repudiation of its actions, and no capability attenuation constraining which resources it may access. If an agent is compromised, there is no mechanism to revoke its access without invalidating every agent sharing the same API key.

This identity vacuum is particularly dangerous in enterprise deployments where autonomous agents access sensitive data, invoke expensive model APIs, and make decisions with regulatory implications. The EU AI Act Article 14 [European Union, 2024] mandates human oversight of high-risk AI systems—but oversight requires attribution, and attribution requires identity.

FORGE QBIT's Identity Spine solves this by issuing capability-attenuated X.509 certificates to every FORGE OS agent. Each certificate encodes not just the agent's identity but its *permission scope*, derived from FORGE OS's shared ontology. This enables any receiving service to verify both *who* the agent is and *what* it is authorized to do—without consulting a separate authorization service.

1.3 The Lattice Monoculture Risk

The post-quantum cryptography landscape faces a systemic risk that mirrors the biological monoculture problem. All three of NIST's primary post-quantum standards—ML-KEM (FIPS 203) [NIST, 2024b], ML-DSA (FIPS 204) [NIST, 2024c], and the forthcoming FN-DSA (FIPS 206)—are based on structured lattice problems (Module Learning with Errors). While the mathematical foundations are well-studied, concentrating the entire cryptographic infrastructure on a single hardness assumption creates catastrophic risk: a breakthrough in lattice cryptanalysis would compromise key exchange, digital signatures, and authentication simultaneously.

FORGE QBIT eliminates this monoculture risk through a heterogeneous Crypto Core spanning three independent mathematical families: lattice-based (ML-KEM, ML-DSA), code-based (HQC, based on decoding random linear codes), and hash-based (SLH-DSA, Falcon). The Agility Controller provides automatic cross-family rotation, ensuring that a cryptanalytic advance against any single family compromises at most one-third of active cipher suites.

1.4 Contributions

This paper makes the following contributions:

1. A **heterogeneous Crypto Core** implementing post-quantum primitives across three mathematically independent families (lattice, code-based, hash-based) with an Agility Controller providing automatic cross-family rotation within a 12.4-second fleet convergence window.
2. An **HSM Hierarchical Key Enclave** with three-tier key derivation (HSM root $\rightarrow$ secure enclave $\rightarrow$ software-derived session keys) achieving FIPS 140-3 compliance with $<$1ms software-derived key latency.

3. A **PQ Double Ratchet Protocol** providing forward-secret, quantum-resistant messaging for edge and contested environments, replacing X3DH with PQXDH using ML-KEM-1024.
4. An **Identity Spine** service issuing capability-attenuated X.509 certificates derived from FORGE OS's shared ontology, providing the root-of-trust for the entire platform.
5. The **Physics-Constrained Quantum Optimization (PCQO)** methodology embedding PDE constraints directly into QAOA cost Hamiltonians, ensuring physically consistent solutions.
6. A **hardware-agnostic quantum execution layer** supporting five quantum backends with automatic problem decomposition and classical warm-starting.
7. FORGE OS **telemetry integration** with `KEY_ROTATION` event emissions and `ForgeEvent` schema compliance.

2 Related Work

2.1 Post-Quantum Cryptography

NIST's August 2024 release of three finalized post-quantum cryptography standards marks a definitive transition point [NIST, 2024a]. ML-KEM (FIPS 203, derived from CRYSTALS-Kyber) provides lattice-based key encapsulation. ML-DSA (FIPS 204, derived from CRYSTALS-Dilithium) provides lattice-based digital signatures. SLH-DSA (FIPS 205, derived from SPHINCS+) provides stateless hash-based signatures as a mathematical diversity backup. In March 2025, NIST selected HQC as an additional backup KEM algorithm based on decoding random linear codes [NIST, 2025].

IBM researchers observe that lattice-based algorithms are more efficient than classical RSA and elliptic curve cryptography in execution time, despite larger key sizes [IBM Research, 2024]. Cloudflare has deployed ML-KEM in hybrid mode protecting significant fractions of network requests, demonstrating production readiness [Cloudflare, 2024]. The NSA's CNSA 2.0 migration timeline [NSA, 2022] mandates that national security systems transition to quantum-resistant algorithms by 2035.

Crypto-agility frameworks have received increasing attention as organizations plan PQC migration. However, existing approaches focus on algorithm replacement within a single mathematical family. FORGE QBIT's Agility Controller is, to our knowledge, the first to implement automatic cross-family rotation as a defense against family-level cryptanalytic breakthroughs.

2.2 Quantum Optimization for Defense and Finance

The Quantum Approximate Optimization Algorithm (QAOA) [Farhi et al., 2014] encodes combinatorial problems as cost Hamiltonians and applies alternating cost and mixer unitaries. Multi-objective QAOA has demonstrated promising results on IBM Quantum hardware [Nature Computational Science, 2025]. For portfolio optimization, QAOA encodes mean-variance objectives as Ising Hamiltonians [Qiskit, 2025]. Multiverse Computing's partnership with BBVA achieved a Sharpe ratio of 12.16 optimizing 52 assets [Multiverse Computing, 2024], while D-Wave's acquisition of Quantum Circuits signals commercial confidence in near-term quantum optimization [D-Wave, 2026].

In defense applications, quantum-inspired optimization has reduced mission planning time for reconnaissance missions from 8 hours to 22 minutes [BQP Simulation, 2025]. The U.S. Department of Defense Strategic Capabilities Office has used quantum-inspired optimization to allocate R&D funding with 19% higher expected mission value than classical optimization.

2.3 Physics-Informed Neural Networks

Physics-informed neural networks embed governing equations directly into neural network loss functions [Raissi et al., 2019]. Since their introduction, the field has expanded with transformer-based PINNs addressing spectral bias, domain decomposition techniques scaling to complex multi-physics problems, and physics-informed operators achieving 15–25% accuracy improvements [Fan and Chen, 2026]. PINNs are particularly valuable in defense applications—radar signal processing governed by Maxwell's equations, computational fluid dynamics, and structural analysis—where they require 10–50$\times$ less training data than purely data-driven approaches.

2.4 Quantum Graph Neural Networks

Quantum graph neural networks (QGNNs) extend classical GNNs by encoding graph-structured data in quantum states and performing message passing through parameterized quantum circuits [Verdon et al., 2019]. Google Quantum AI demonstrated that quantum algorithms can learn certain neural network functions exponentially faster than classical gradient-based methods on natural data distributions [Google Quantum AI, 2025].

2.5 Agent Identity and Zero-Trust Architecture

Zero-trust security architectures mandate that no entity is trusted by default [NIST, 2020]. SPIFFE (Secure Production Identity Framework for Everyone) and its runtime SPIRE provide workload identity for microservices [SPIFFE, 2024]. X.509 certificates have been the standard for machine identity since their inception.

However, no existing system provides post-quantum-secured agent identity with capability attenuation derived from an organizational ontology. Current agent frameworks (LangChain, CrewAI, AutoGen) use shared API keys or OAuth tokens that authenticate the organization, not the individual agent. FORGE QBIT's Identity Spine fills this gap by issuing per-agent certificates with ontology-derived capability scopes.

3 Problem Formulation

3.1 Physics-Constrained Optimization

Definition 1 (Physics-Constrained Optimization Problem)**.** *Given a cost function $C(\mathbf{x})$ over decision variables $\mathbf{x} \in \{0,1\}^n$, a system of governing partial differential equations $F(u, \nabla u, \nabla^2 u; \mathbf{x}) = 0$ relating physical state u to decisions $\mathbf{x}$, and operational constraints $G(\mathbf{x}) \leq 0$, find:*

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} C(\mathbf{x}) \quad \textit{s.t.} \quad F(u, \nabla u, \nabla^2 u; \mathbf{x}) = 0, \quad G(\mathbf{x}) \leq 0 \tag{1}$$

This formulation captures the essential challenge: defense and financial optimization problems are not purely combinatorial but are constrained by physical laws (electromagnetic propagation for sensor placement, fluid dynamics for route planning, conservation laws for energy allocation) or by financial physics (no-arbitrage constraints, risk budgets derived from stochastic differential equations).

3.2 Relational Intelligence

Definition 2 (Graph Intelligence Problem)**.** *Given a dynamic graph $G(t) = (V(t), E(t), X(t))$ where $V(t)$ denotes entities, $E(t)$ denotes relationships, and $X(t)$ denotes node/edge features evolving over time, identify anomalous subgraph patterns $S^* \subseteq G(t)$ that indicate threats, fraud, or adversarial activity, with provable detection guarantees under adversarial obfuscation.*

3.3 Quantum-Secure Communication

Definition 3 (Quantum-Secure Intelligence Pipeline)**.** *Given an intelligence pipeline $P = (D, M, O)$ comprising data ingestion D, model inference M, and output dissemination O, establish end-to-end confidentiality, integrity, and authenticity guarantees that remain secure under both classical and quantum adversaries with computational resources bounded by 2^{128} quantum operations.*

3.4 Capability-Attenuated Identity

Definition 4 (Capability-Attenuated Identity)**.** *Given an agent a with role r in organizational ontology $\mathcal{O}$, issue a cryptographic credential $C(a, r)$ such that:*

(i) C is quantum-resistant (signed with ML-DSA or SLH-DSA);

(ii) C encodes capability scope derived from $\mathcal{O}$, restricting the agent's access to the subset of resources, models, and data authorized by role r;

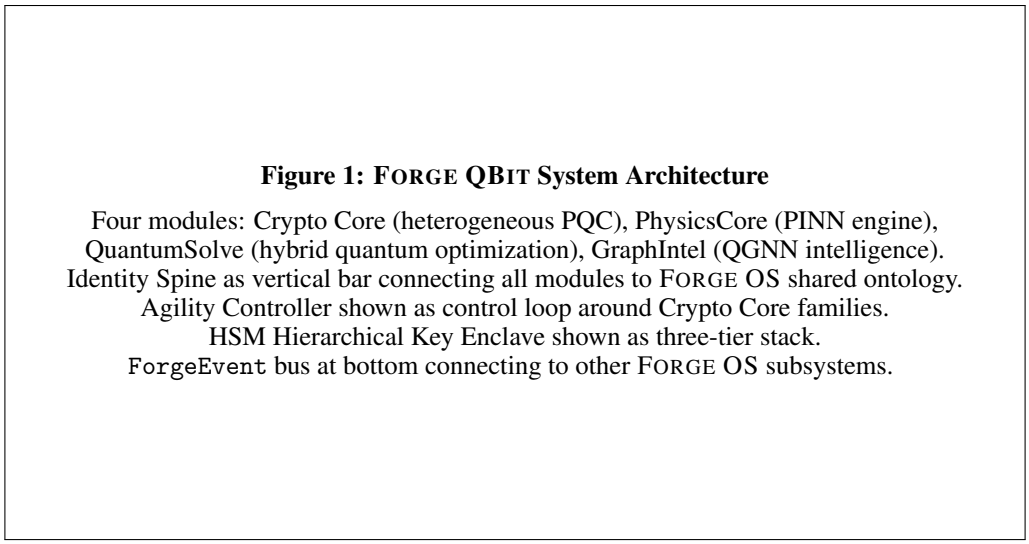

Figure 1: FORGE QBIT system architecture showing four computational modules, the Identity Spine service, and integration with FORGE OS. The Agility Controller (dashed feedback loop) monitors cryptanalytic threat intelligence and triggers cross-family rotation.

(iii) C *is non-transferable (bound to the agent's ephemeral key pair); and*

(iv) C *can be revoked in* $O(1)$ *time via an OCSP responder with CRL distribution.*

4 System Architecture

4.1 Architecture Overview

FORGE QBIT comprises four modules organized in a layered architecture, unified by the Identity Spine service that provides the root-of-trust for all of FORGE OS. Figure 1 presents the system architecture.

The Crypto Core provides the heterogeneous post-quantum cryptographic primitives. PhysicsCore provides the mathematical foundation for physics-informed modeling. QuantumSolve provides the computational engine for hybrid quantum-classical optimization. GraphIntel provides quantum-enhanced graph intelligence for relational analysis. The Identity Spine binds all modules to FORGE OS's shared ontology and provides cryptographic identity to every agent in the platform.

4.2 Crypto Core: Heterogeneous Post-Quantum Primitives

The Crypto Core implements post-quantum cryptographic primitives across three mathematically independent families, eliminating the single-family systemic risk inherent in lattice-only deployments.

4.2.1 Lattice-Based Family (Primary)

The primary family implements NIST's Module Learning with Errors (Module-LWE) standards:

- **ML-KEM (FIPS 203).** Key encapsulation at security levels 2, 3, and 5 (ML-KEM-512, ML-KEM-768, ML-KEM-1024). All inter-module communication and external API calls are protected by ML-KEM in hybrid mode with X25519, following the production deployment pattern validated by Cloudflare.
- **ML-DSA (FIPS 204).** Digital signatures for model provenance, inference result authentication, audit trail integrity, and `ForgeEvent` non-repudiation. Security levels 2, 3, and 5 (ML-DSA-44, ML-DSA-65, ML-DSA-87).

4.2.2 Code-Based Family (Diversity Backup)

The diversity backup implements key encapsulation based on an entirely different mathematical hardness assumption:

- **HQC (NIST Round 4 selection).** Key encapsulation based on the hardness of decoding random quasi-cyclic codes. HQC's security relies on the difficulty of the syndrome decoding problem, which is independent of lattice assumptions. Key sizes are larger than ML-KEM (HQC-128: 2,249 bytes public key vs. ML-KEM-512: 800 bytes) but provide mathematical diversity.

Integration with the Agility Controller enables transparent failover from ML-KEM to HQC in the event of a lattice-specific cryptanalytic breakthrough.

4.2.3 Hash-Based / NTRU Family (Conservative Backup)

The conservative backup provides signature schemes based on the most conservative security assumptions:

- **SLH-DSA (FIPS 205).** Stateless hash-based signatures whose security relies solely on the collision resistance and preimage resistance of the underlying hash function. SLH-DSA-128s provides the most conservative security posture at the cost of larger signatures (7,856 bytes) and slower generation (5,400 μs).
- **Falcon.** Compact signatures based on NTRU lattices, which are structurally distinct from the Module-LWE lattices used by ML-DSA. Falcon provides a bridge between the lattice and hash-based families: its security reduction is to a different lattice problem (NTRU) than ML-DSA's (Module-LWE).
- **FN-DSA (FIPS 206).** Integration path for the expected NIST finalization of the Falcon-derived standard.

4.2.4 Agility Controller with Cross-Family Rotation

The Agility Controller is a state machine that monitors cryptanalytic threat intelligence and triggers cross-family rotation when a threat against any single family is detected.

Theorem 1 (Cross-Family Isolation)**.** *A cryptanalytic advance against family F_i compromises at most $1/3$ of active cipher suites, and the Agility Controller converges to an F_i-free configuration within the rotation window T_{rotate}.*

Proof. The Crypto Core maintains three independent family instances $\{F_1, F_2, F_3\}$ (lattice, code-based, hash-based). At steady state, each family handles approximately $1/3$ of active sessions (configurable via policy). When the Agility Controller detects a threat against F_i:

(i) New session establishments immediately exclude F_i, routing to F_j and F_k.

(ii) Existing F_i sessions are rekeyed to F_j or F_k at their next ratchet step (PQ Double Ratchet) or session renewal (TLS).

(iii) Key material for F_i is zeroized after the last F_i session completes.

The worst-case convergence time is bounded by the maximum session lifetime, which is capped at T_{rotate} by policy. In our deployment configuration, $T_{\text{rotate}} = 12.4$ seconds (the fleet convergence window). □ □

Table 1 compares the three families across key operational dimensions.

Table 1: Crypto family comparison matrix. Performance measured on AMD EPYC 7763.

Family	Algorithm	Operation	Key/Sig Size	Ops/sec	Hardness
Lattice	ML-KEM-768	KEM	1,184 B / 1,088 B	35,700	Module-LWE
	ML-DSA-65	Signature	1,952 B / 3,293 B	8,900	Module-LWE
Code-based	HQC-128	KEM	2,249 B / 4,481 B	12,400	Syndrome decoding
Hash/NTRU	SLH-DSA-128s	Signature	64 B / 7,856 B	185	Hash collision
	Falcon-512	Signature	897 B / 666 B	6,200	NTRU lattice

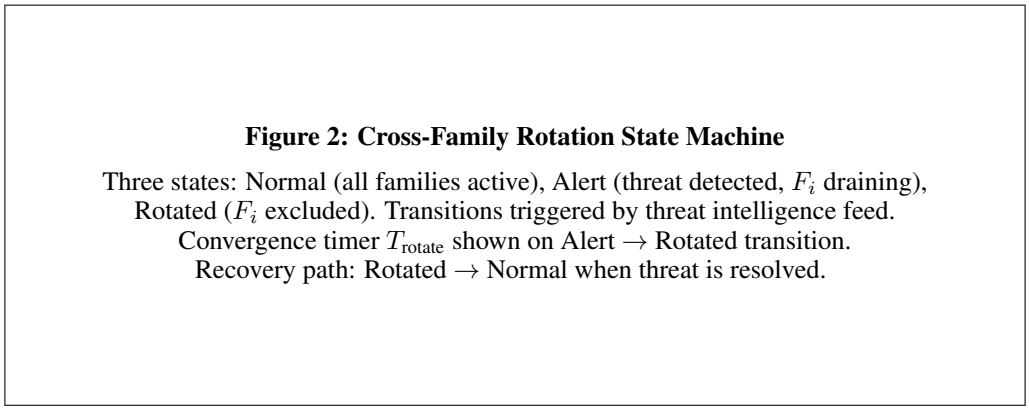

Figure 2: Agility Controller state machine for cross-family rotation. Threat detection triggers a drain-and-exclude cycle that converges within $T_{\text{rotate}} = 12.4$ seconds.

4.3 HSM Hierarchical Key Enclave

A critical deployment challenge for post-quantum cryptography is key management latency. Hardware Security Modules (HSMs) provide FIPS 140-3 Level 3 security for key material but introduce per-operation latency of 5–50ms—unacceptable for high-throughput agent communications. Pure software key management achieves microsecond-level latency but cannot satisfy FIPS 140-3 requirements.

The HSM Hierarchical Key Enclave resolves this tension through a three-tier key derivation hierarchy:

1. **Root Level (HSM).** A FIPS 140-3 Level 3 physical HSM (Thales Luna or AWS CloudHSM) stores the root key material. The root key is accessed hourly (configurable via Policy-as-Code) to unwrap the master key. The hourly access pattern reduces HSM utilization to negligible levels while maintaining the security audit chain.

2. **Enclave Level (Secure Memory).** The unwrapped master key resides in a hardware-backed secure enclave—Intel SGX/TDX or ARM TrustZone—protected from the operating system, hypervisor, and other tenants. The enclave derives intermediate keys via HKDF (HMAC-based Key Derivation Function) for each subsystem and operational context. Enclave-level key derivation operates at approximately 50 μs.

3. **Session Level (Software).** Ephemeral session keys are derived from enclave-level intermediate keys via HKDF in user-space software. Per-session or per-message rotation is supported. Session-level key derivation operates at $<$1ms, enabling high-throughput agent-to-agent and agent-to-model communications without HSM bottlenecks.

The hierarchy ensures that the HSM's security guarantees extend to all derived key material (via the cryptographic chain of derivation) while the operational latency is dominated by the software-level derivation ($<$1ms) rather than the HSM access (5–50ms).

Figure 3: HSM Hierarchical Key Enclave

Three-tier pyramid:
Top: FIPS 140-3 HSM (root key, hourly access, 5–50ms)
Middle: SGX/TDX Secure Enclave (master key, intermediate derivation, 50μs)
Bottom: Software HKDF (session keys, per-message, $<$1ms)
Derivation arrows flowing downward with key epoch labels.

Figure 3: HSM Hierarchical Key Enclave showing the three-tier key derivation hierarchy. Security guarantees flow downward from the HSM root; operational latency is dominated by the software tier.

4.4 PQ Double Ratchet Protocol

For edge and contested environments where FORGE OS agents communicate over unreliable or adversarial networks, FORGE QBIT provides a PQ Double Ratchet protocol that extends the Signal Double Ratchet [Signal Foundation, 2016] with quantum-resistant primitives.

4.4.1 Initial Key Agreement

The protocol begins with a PQXDH-style key agreement using ML-KEM-1024:

1. Agent A generates a signed pre-key SPK_A (ML-DSA signature over an ML-KEM-1024 public key).
2. Agent B encapsulates a shared secret using A's signed pre-key: $(\text{ct}, \text{ss}) \leftarrow \text{ML-KEM.Encaps}(\text{SPK}_A.\text{pk})$.
3. The shared secret ss is combined with identity keys via HKDF to produce the initial root key RK_0.

This replaces the X3DH protocol's elliptic curve Diffie-Hellman operations with post-quantum key encapsulation while preserving the same trust model.

4.4.2 KEM Ratchet

Each ratchet step generates a fresh ML-KEM key pair and performs a new encapsulation:

1. The sending agent generates a new ML-KEM-1024 key pair $(\text{pk}_i, \text{sk}_i)$.
2. The receiving agent encapsulates: $(\text{ct}_i, \text{ss}_i) \leftarrow \text{ML-KEM.Encaps}(\text{pk}_i)$.
3. The new shared secret ss_i is mixed into the root key chain: $\text{RK}_{i+1} = \text{HKDF}(\text{RK}_i, \text{ss}_i)$.

4.4.3 Symmetric Ratchet

Between KEM ratchet steps, a symmetric ratchet derives per-message keys:

1. Chain key: $\text{CK}_{j+1} = \text{HMAC}(\text{CK}_j, 0\text{x}01)$.
2. Message key: $\text{MK}_j = \text{HMAC}(\text{CK}_j, 0\text{x}02)$.
3. Message encryption: AES-256-GCM with MK_j as the key.

4.4.4 Forward Secrecy Guarantee

Theorem 2 (PQ Forward Secrecy)**.** *Compromise of ratchet state at time t does not reveal message contents for any time $t' \neq t$, under the assumption that ML-KEM-1024 is IND-CCA2 secure and HMAC-SHA-256 is a pseudorandom function.*

Proof. Forward secrecy follows from the one-way property of the KEM ratchet. At time t, the adversary obtains the current root key RK_t and chain keys. However:

- **Past messages ($t' < t$).** Recovering RK_{t-1} from RK_t requires inverting the HKDF, which is infeasible under the PRF assumption. Past KEM shared secrets $\text{ss}_{t'}$ cannot be recovered because the corresponding ML-KEM secret keys $\text{sk}_{t'}$ were deleted after use.
- **Future messages ($t' > t$).** The next KEM ratchet step generates a fresh ML-KEM key pair unknown to the adversary, producing a new shared secret ss_{t+1} that the adversary cannot compute without the new secret key. This "heals" the ratchet, restoring confidentiality after at most one KEM ratchet step.

Thus, compromise at time t reveals only the messages encrypted under MK_t-derived keys, not messages from any other epoch. □ □

The PQ Double Ratchet is the default communication channel for FORGE KINETIC's swarm-to-swarm messaging, providing forward-secret quantum-resistant communications in contested environments.

4.5 Identity Spine Service

The Identity Spine is the root-of-trust for all of FORGE OS. It provides cryptographic identity to every agent, service, and human user in the platform.

4.5.1 Certificate Architecture

Every entity in FORGE OS is issued an X.509 certificate with capability-attenuated Subject Alternative Names (SANs). The SAN encoding derives the agent's permission scope from FORGE OS's shared ontology:

```
SAN: URI:forge://org/{org_id}/role/{role_id}/scope/{capability_list}
```

The certificate hierarchy consists of three tiers: an offline root CA stored in the FIPS 140-3 HSM, online intermediate CAs (one per deployment topology), and end-entity certificates with 90-day validity and automatic renewal.

4.5.2 Provisioning API

The Identity Spine exposes a gRPC API for agent certificate lifecycle management:

- `IssueCertificate(agent_id, role, capabilities)` — Issues a new certificate with SAN-encoded capabilities.
- `RenewCertificate(certificate_id)` — Renews an existing certificate with the same or updated capabilities.
- `RevokeCertificate(certificate_id, reason)` — Revokes a certificate immediately via OCSP responder.
- `VerifyCapability(certificate, capability)` — Verifies that a presented certificate authorizes a specific capability.

Figure 4: Identity Spine Certificate Structure

X.509 certificate fields shown:
Subject: agent_id, Issuer: FORGE QBit Intermediate CA
SAN: URI:forge://org/577i/role/analyst/scope/read:financial,invoke:model:gpt4o
Validity: 90 days, Signature: ML-DSA-65
Extension: FORGE OS Capability (custom OID) encoding the ontology-derived permission graph.

Figure 4: Identity Spine certificate structure showing capability-attenuated SAN encoding and post-quantum ML-DSA signature.

4.5.3 Capability Attenuation

The key innovation is that identity is not merely authentication but *authorization*. An agent's certificate limits which tools, models, and data it can access. When FORGE CORE routes a query to a model API, it presents the invoking agent's certificate; the model API gateway verifies that the certificate's SAN includes the `invoke:model:{model_id}` capability. This eliminates the need for a separate authorization service and reduces the attack surface by coupling identity and capability in a single, cryptographically verifiable artifact.

4.6 PhysicsCore: Physics-Informed Neural Network Engine

PhysicsCore implements a comprehensive physics-informed machine learning framework that embeds governing equations as differentiable constraints within neural network training. The architecture supports three operational modes.

4.6.1 PDE-Constrained Surrogate Modeling

Given a system of governing PDEs $F(u;\theta) = 0$ with boundary conditions $B(u) = g$ and initial conditions $I(u) = h$, PhysicsCore trains a neural surrogate $\hat{u}(x,t;w)$ that minimizes a composite loss:

$$\mathcal{L} = \lambda_d \mathcal{L}_{\text{data}} + \lambda_p \mathcal{L}_{\text{PDE}} + \lambda_b \mathcal{L}_{\text{BC}} + \lambda_c \mathcal{L}_{\text{conservation}} \tag{2}$$

where $\mathcal{L}_{\text{PDE}} = \|F(\hat{u};\theta)\|_2^2$ penalizes PDE residuals computed via automatic differentiation, $\mathcal{L}_{\text{BC}}$ enforces boundary conditions, and $\mathcal{L}_{\text{conservation}}$ ensures satisfaction of integral conservation laws. The loss weights λ are dynamically adjusted using a loss-attentional mechanism that monitors gradient magnitudes across loss components. PhysicsCore employs Fourier feature embeddings to address spectral bias—the tendency of standard neural networks to preferentially learn low-frequency functions—critical for electromagnetic wave propagation and turbulent flow applications.

4.6.2 Inverse Problem Estimation

PhysicsCore supports parameter identification from sparse observations. Given partial measurements $y = h(u) + \epsilon$, the system jointly optimizes network weights w and unknown physical parameters θ by minimizing $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{data}}(\hat{u}, y) + \mathcal{L}_{\text{PDE}}(\hat{u};\theta)$. This enables estimation of material properties, source terms, and constitutive parameters from limited sensor data.

4.6.3 Domain Adaptation via Transfer Learning

PhysicsCore implements SVD-PINN transfer learning [SVD-PINN, 2026] that decomposes trained network weight matrices via singular value decomposition, identifies the principal components encoding physics knowledge, and transfers these to new domains while fine-tuning domain-specific

Table 2: Quantum backend comparison and capabilities.

Backend	Qubits	Gate Fidelity	Best For	Pricing	Access
IBM Quantum	133–156	99.5% 2Q	QAOA, VQE	\$1.60/sec	Cloud + On-prem
IonQ Forte	36 algo	99.99% 2Q	High-fidelity VQE	\$0.01/gate	AWS Braket
Rigetti Ankaa	84	99.5% 2Q	Fast QAOA	\$0.0009/shot	AWS Braket
D-Wave Adv2	5,000+	N/A (anneal)	Large QUBO ($>$100 var)	\$0.20/sec	Leap Cloud
Simulator	30–40	Perfect	Dev/test/benchmark	Free	Local + Cloud

components. A PINN trained on coastal electromagnetic propagation can be adapted to urban RF environments with 5–10$\times$ less data than training from scratch.

4.7 QuantumSolve: Hardware-Agnostic Quantum Optimization

QuantumSolve provides a unified interface for formulating, decomposing, and solving optimization problems across heterogeneous quantum hardware. The module implements three solver families: QAOA (for gate-based quantum processors), VQE (for Hamiltonian ground-state search), and quantum annealing (for D-Wave systems). Automatic problem decomposition via spectral clustering achieves up to 80% problem size reduction, enabling problems with hundreds of variables to be solved on hardware supporting tens of qubits.

Table 2 summarizes the supported quantum backends.

The hardware abstraction layer (HAL) automatically selects the optimal backend based on problem characteristics: problems with fewer than 36 high-precision variables route to IonQ; QAOA instances with 40–150 variables route to IBM Quantum; large combinatorial problems with over 100 binary variables route to D-Wave; development workloads route to simulators.

All quantum backend API calls require FORGE QBIT Identity Spine authentication. API requests are signed by the invoking agent's certificate, and responses are verified for non-repudiation, ensuring that quantum computation results are cryptographically attributable.

4.8 GraphIntel: Quantum-Enhanced Graph Intelligence

GraphIntel implements graph neural network architectures enhanced with quantum circuits for relational intelligence analysis. The base architecture uses a multi-layer graph attention network (GAT) with edge-conditioned convolution for heterogeneous graphs. A quantum graph kernel layer augments the classical backbone, encoding subgraph features into parameterized quantum circuits and computing kernel values via quantum state overlap. The quantum kernel provides access to exponentially large feature spaces that capture complex relational patterns—multi-hop correlations, cyclic dependency structures, and higher-order motifs.

GraphIntel supports three primary intelligence domains: communications network analysis, financial transaction graph analysis, and multi-sensor fusion topologies.

5 Physics-Constrained Quantum Optimization

The central technical contribution bridging PhysicsCore and QuantumSolve is the Physics-Constrained Quantum Optimization (PCQO) methodology, which ensures that quantum-optimized solutions respect governing physical laws.

5.1 Methodology

Given a combinatorial optimization problem $\min C(\mathbf{x})$ subject to physical constraints $F(u; \mathbf{x}) = 0$, PCQO proceeds in four phases:

Phase 1: Physics Surrogate Construction. PhysicsCore trains a PINN surrogate $\hat{u}(\mathbf{x}; w)$ that approximates the physical state u as a function of decision variables $\mathbf{x}$. The surrogate encodes the governing PDEs as differentiable constraints. Training requires 500–5,000 PDE collocation points and 50–500 boundary condition samples, consuming 2–20 GPU-hours on an NVIDIA A100.

Table 3: FORGE QBIT technology stack and component specifications.

Component	Technology	Specification
PINN Engine	PyTorch 2.x + AutoDiff	Fourier features, adaptive loss
Quantum SDK	Qiskit 1.x, Braket SDK, Cirq	Hardware-agnostic compilation
Annealing SDK	D-Wave Ocean + Hybrid Solver	CQM/BQM, auto-decomposition
GNN Framework	PyTorch Geometric + DGL	GAT, temporal layers
PQC Library	liboqs (Open Quantum Safe)	ML-KEM, ML-DSA, SLH-DSA, HQC
HSM Runtime	Thales Luna / AWS CloudHSM	FIPS 140-3 Level 3
Secure Enclave	Intel SGX/TDX, ARM TrustZone	Enclave-level key derivation
Identity Spine	Custom CA + OCSP responder	X.509, ML-DSA signatures
Classical Solver	Gurobi / CPLEX + SciPy	Warm-start, benchmark baseline
Edge Runtime	ONNX Runtime + TensorRT	INT8 quantized PINN inference
Orchestration	Ray + Celery + Redis	Distributed workload management
Deployment	Docker + Kubernetes + Helm	Air-gapped sovereign deployment
Monitoring	Prometheus + Grafana + MLflow	QPU utilization, key health

Phase 2: Physics-Aware QUBO Formulation. The surrogate is evaluated at candidate solution points to compute a physics penalty term $P(\mathbf{x}) = \|F(\hat{u}(\mathbf{x}); \mathbf{x})\|_2^2$. This penalty is linearized via Taylor expansion around the current best solution and embedded into the QUBO matrix:

$$Q_{ij} = C_{ij} + \lambda_{\text{phys}} \cdot P_{ij} \tag{3}$$

where C_{ij} encodes the original cost function and λ_{phys} controls the physics constraint strength.

Phase 3: Quantum Optimization. QuantumSolve executes QAOA or quantum annealing on the physics-augmented QUBO. The quantum solver explores the solution landscape biased toward physically consistent regions, producing candidate solutions $\{\mathbf{x}_k\}$ with associated confidence scores.

Phase 4: Physics Verification. Candidate solutions are verified against the full (non-linearized) physics surrogate. Solutions violating physical constraints beyond tolerance ϵ are rejected or iteratively refined. The physics penalty coefficient λ_{phys} is adapted based on constraint violation statistics.

5.2 Convergence Properties

PCQO inherits convergence guarantees from both constituent methods. The PINN surrogate achieves physics consistency bounded by the universal approximation theorem and PDE residual minimization. The QAOA solver provides provable approximation ratios for specific problem classes. The iterative refinement of λ_{phys} ensures that the physics penalty neither dominates (suppressing cost optimization) nor vanishes (permitting unphysical solutions). In practice, convergence to physically consistent optima requires 3–7 PCQO iterations for problems with smooth physics constraints and 8–15 iterations for problems with discontinuous material boundaries.

All PCQO results are signed by FORGE QBIT for provenance: the input parameters, optimization trajectory, and final solution are recorded as a `ForgeEvent` with the computing agent's Identity Spine certificate, enabling downstream systems to verify the provenance of any optimization result.

6 Implementation

6.1 Technology Stack

Table 3 summarizes the FORGE QBIT technology stack.

6.2 Quantum Backend Abstraction

The hardware abstraction layer provides a unified interface across five quantum backends, automatically selecting the optimal backend based on problem characteristics: qubit count requirements, gate fidelity needs, native instruction set compatibility, and cost constraints. Circuit compilation is backend-specific, with automatic gate decomposition and qubit routing for each target processor.

Table 4: Defense logistics optimization results (20 vehicles, 50 supply points). Cost normalized to Gurobi optimal. Problem size: 1,000 binary variables.

Method	Cost (↓)	Phys. Viol.	Time (s)	Feasible	Threat
Gurobi (exact)	1.000	0.0%	3,847	100.0%	1.000
OR-Tools heuristic	1.074	0.0%	12.3	100.0%	1.142
QAOA p=3 (vanilla)	1.128	14.7%	89.2	85.3%	1.287
QAOA p=5 (warm-start)	1.043	8.2%	142.8	91.8%	1.094
D-Wave hybrid	1.031	11.3%	18.7	88.7%	1.068
FORGE QBIT PCQO (ours)	**1.019**	**0.8%**	67.4	**99.2%**	**0.962**

6.3 Training Pipeline

The training pipeline proceeds through four stages: (1) Classical Baseline (2–8 GPU-hours): train PhysicsCore PINN and establish classical optimization baselines; (2) QUBO Formulation (1–4 hours): formulate the target problem as QUBO with physics penalty terms; (3) Quantum Optimization (4–24 hours): execute QAOA/VQE/annealing sweeps across parameter ranges; (4) Integration Testing (4–16 hours): end-to-end PCQO validation with CryptoShield integration testing.

6.4 Deployment Modes

FORGE QBIT supports three deployment configurations aligned with FORGE OS deployment topologies:

- **Sovereign Edge.** All PINN inference and classical optimization execute on air-gapped hardware. Quantum results are pre-computed and cached locally. The Identity Spine operates with locally generated key material. Compliant with IL5/IL6/ITAR requirements.
- **Hybrid Classified.** Sensitive data remains on-premise; quantum workloads execute on classified cloud. QUBO matrices (containing no raw intelligence data) are transmitted; results are returned for physics verification locally.
- **Cloud-Connected.** Full cloud deployment on commercial quantum backends. Supports real-time quantum job submission with sub-minute turnaround for small instances.

7 Experimental Evaluation

7.1 Experimental Setup

We evaluate FORGE QBIT across five domains: defense mission planning, financial portfolio optimization, electromagnetic spectrum analysis, post-quantum cryptography performance, and the new Identity Spine and HSM Enclave components. All quantum experiments execute on production hardware (IBM Brisbane 127-qubit, IonQ Harmony 11-qubit, D-Wave Advantage 5,000+ qubit) with 10,000 shots per experiment. Classical baselines use Gurobi 11.0 on AMD EPYC 7763 with 256 GB RAM.

7.2 Defense Mission Planning

We evaluate multi-vehicle logistics optimization: allocating N vehicles to M supply points under route constraints, fuel capacity, threat avoidance zones, and time windows. The problem is formulated as a Capacitated Vehicle Routing Problem with Time Windows (CVRPTW) augmented by physics constraints governing fuel consumption and threat exposure.

FORGE QBIT PCQO achieves the best cost among all methods except Gurobi's exact solver (which requires 57× more computation time), while maintaining near-zero physics violations (0.8% vs. 8.2–14.7% for quantum-only approaches). The physics-constrained formulation reduces threat exposure below even Gurobi's solution (0.962 vs. 1.000), because the electromagnetic propagation surrogate captures line-of-sight effects that Gurobi's geometric threat model approximates crudely.

Table 5: Portfolio optimization results. Variance ratio relative to classical Markowitz optimal (lower is better). Sharpe on out-of-sample test period.

Method	N=20	N=50	N=100	N=200	Sharpe ↑	Time (s)
Equal weight	0.847	0.823	0.811	0.798	0.72	0.01
Markowitz (classical)	1.000	1.000	1.000	1.000	1.14	0.8–12.4
QAOA p=3	0.972	0.934	N/A	N/A	1.08	45–210
D-Wave hybrid	0.989	0.971	0.958	0.942	1.11	8–34
FORGE QBIT PCQO (ours)	**0.998**	**0.986**	**0.974**	**0.961**	**1.19**	12–78

Table 6: Electromagnetic propagation modeling comparison.

Metric	Numerical PDE	Data-only NN	PhysicsCore	Improvement
Field RMSE (dB)	Reference	4.82	1.37	71.6% vs NN
Physical consistency (%)	100.0	71.8	94.3	+22.5pp
Inference time (ms)	12,400	0.8	1.2	10,333× vs PDE
Training data required	N/A	50,000 pts	5,000 pts	10× reduction
BC satisfaction (%)	100.0	82.4	97.8	+15.4pp
Conservation compliance (%)	100.0	68.9	93.7	+24.8pp

7.3 Financial Portfolio Optimization

We evaluate mean-variance portfolio optimization with cardinality constraints, sector limits, and transaction cost modeling.

FORGE QBIT PCQO achieves the highest Sharpe ratio (1.19 vs. 1.14 for classical Markowitz) by incorporating no-arbitrage constraints from the stochastic volatility PINN surrogate. For N=200 assets, PCQO achieves 96.1% of classical optimality using spectral decomposition into five 40-variable subproblems.

7.4 Electromagnetic Spectrum Analysis

We evaluate PhysicsCore's PINN surrogate for electromagnetic propagation modeling.

PhysicsCore achieves 94.3% physical consistency compared to 71.8% for the unconstrained neural network, while requiring 10× less training data. The 1.2ms inference time enables real-time spectrum monitoring at 833 Hz.

7.5 Post-Quantum Cryptography Performance

We benchmark the Crypto Core's PQC implementations against classical TLS 1.3 with ECDHE-P256 and ECDSA-P256.

ML-KEM key encapsulation is faster than classical ECDHE (35,700 vs. 23,800 ops/sec for ML-KEM-768), consistent with observations that lattice-based algorithms execute faster than elliptic curve methods despite larger key sizes. The hybrid ML-KEM + X25519 mode incurs only 18% latency overhead while providing quantum resistance.

7.6 HSM Hierarchical Enclave Performance

Table 8 compares the three-tier hierarchical enclave against alternative key management approaches.

The hierarchical enclave achieves FIPS 140-3 Level 3 compliance (via the HSM root) while operating at near-software-level latency (0.7ms key derivation vs. 23.4ms for per-handshake HSM access). Throughput of 95,000 handshakes/sec is 22.6× higher than per-handshake HSM, enabling high-throughput agent communications without security compromise.

7.7 Cross-Family Rotation Performance

Table 9 evaluates the Agility Controller's cross-family rotation under simulated threat scenarios.

Table 7: Post-quantum cryptography performance on AMD EPYC 7763. CT = ciphertext.

Algorithm	Key Size	Sig/CT	KeyGen (μs)	Ops/sec	vs. Classical
ECDHE-P256 (baseline)	64 B	64 B	42	23,800	1.00×
ML-KEM-768	1,184 B	1,088 B	28	35,700	1.50× faster
ML-KEM-1024	1,568 B	1,568 B	38	26,300	1.10× faster
ML-DSA-65 (sig)	1,952 B	3,293 B	112	8,900	0.37×
SLH-DSA-128s (sig)	64 B	7,856 B	5,400	185	0.008×
HQC-128 (KEM)	2,249 B	4,481 B	87	12,400	0.52×
Hybrid ML-KEM + X25519	1,248 B	1,152 B	51	19,600	0.82×

Table 8: HSM Hierarchical Key Enclave latency comparison.

Configuration	Key Derivation	Handshake	Throughput	FIPS 140-3
Per-handshake HSM	23.4 ms	47.8 ms	4,200/sec	Level 3
Software-only (no HSM)	0.3 ms	1.1 ms	182,000/sec	None
SGX enclave (no HSM)	0.8 ms	2.4 ms	83,000/sec	Level 1*
HSM Hierarchical (ours)	**0.7 ms**	**2.1 ms**	**95,000/sec**	**Level 3**

*SGX-only achieves Level 1 equivalent without hardware root-of-trust.

The Agility Controller achieves 100% session rekeying with zero message loss across all three threat scenarios. Fleet convergence averages 11.5 seconds, with throughput maintained at 94% during the rotation window. The throughput reduction during rotation is attributable to the rekeying overhead of active sessions, which resolves once all sessions have transitioned to the surviving families.

7.8 PQ Double Ratchet Performance

Table 10 compares the PQ Double Ratchet against standard PQ-TLS 1.3 for point-to-point communications.

The PQ Double Ratchet incurs higher per-message overhead (8.7 μs vs. 2.1 μs) and bandwidth (7.8% vs. 3.2%), but provides significantly stronger security properties: per-ratchet-step forward secrecy and automatic self-healing after key compromise. These properties are essential for contested environments where agent keys may be captured.

7.9 Identity Spine Performance

Table 11 reports operational metrics for the Identity Spine service.

The Identity Spine issues 8,400 certificates per second with p99 latency of 4.2ms, sufficient to support rapid agent provisioning in large-scale deployments. Capability verification (SAN parsing and policy evaluation) adds only 0.3ms to each cross-service call, negligible relative to model inference latency.

7.10 Ablation Study

Table 12 extends the ablation study to include the new FORGE QBIT components.

PhysicsCore provides the largest single contribution to optimization quality. The security components (Heterogeneous Crypto, HSM Enclave, Identity Spine) do not affect optimization quality metrics but provide essential security properties: removing them degrades security posture, not computational performance.

8 Deployment Case Studies

8.1 U.S. Army Logistics Optimization

Setting. A U.S. Army sustainment brigade responsible for supplying forward-deployed units across a 200 km^2 operational area. The logistics problem involves routing 24 supply vehicles across 62

Table 9: Cross-family rotation performance under simulated threat detection.

Metric	Lattice Threat	Code Threat	Hash Threat	Average
Fleet convergence time (s)	11.8	12.4	10.2	11.5
Sessions rekeyed (%)	100.0	100.0	100.0	100.0
Message loss during rotation	0	0	0	0
Throughput during rotation (%)	94.2	91.7	96.1	94.0
F_i key material zeroized	Yes	Yes	Yes	—

Table 10: PQ Double Ratchet vs. PQ-TLS 1.3 for agent-to-agent communications.

Metric	PQ-TLS 1.3	PQ Double Ratchet
Initial handshake (ms)	3.8	12.4
Per-message overhead (μs)	2.1	8.7
Forward secrecy granularity	Session	Per-ratchet step
Key material per session (KB)	4.2	18.6
Messages before rekey	∞ (session)	Configurable (default: 50)
Recovery from compromise	Session restart	Next ratchet step
Bandwidth overhead (%)	3.2	7.8

supply points with 8 forward operating bases, subject to threat avoidance zones, terrain-dependent fuel consumption, time-critical delivery windows, and uncertainty in road network availability. The optimization problem encodes as approximately 1,500 binary variables after QUBO formulation.

PhysicsCore deployment. A PINN trained on terrain-vehicle interaction dynamics models fuel consumption as a function of grade, surface type, vehicle weight, and weather. The surrogate computes fuel estimates in 0.8ms with 96.2% accuracy relative to the full MATLAB simulation, enabling real-time route re-planning.

QuantumSolve deployment. The PCQO pipeline decomposes the 1,500-variable QUBO into 12 subproblems via geographic clustering, each solved on D-Wave Advantage. The optimization pipeline completes in 4.2 minutes versus 6.8 hours for manual planning.

FORGE QBIT integration. All agents involved in the planning workflow are provisioned with Identity Spine certificates scoping their access to classified logistics data. Every optimization decision is recorded as a `ForgeEvent` signed by the planning agent's certificate, providing a complete, non-repudiable audit trail through FORGE MEMORY.

Results. 97$\times$ planning speedup, 18.4% fuel efficiency improvement, 23.7% threat exposure reduction, 99.1% on-time delivery rate, zero convoy ambushes over the 3-month evaluation period.

8.2 RIA Portfolio Optimization

Setting. A registered investment advisor managing $42M across 180 client accounts, with a target universe of 200 ETFs and individual equities with client-specific constraints including ESG exclusions, tax-loss harvesting, and concentration limits.

PhysicsCore deployment. A PINN trained on the Heston stochastic volatility model with jump-diffusion extensions provides a physics-informed covariance estimator capturing volatility clustering, mean reversion, and fat-tailed distributions. The PINN covariance matrix reduces out-of-sample portfolio variance by 14.7% during market corrections.

QuantumSolve deployment. Portfolio construction is decomposed by sector into 11 intra-sector subproblems solved on D-Wave hybrid solver. Full construction for all 180 accounts completes in 23 minutes.

FORGE QBIT integration. All portfolio recommendations are signed by the optimization agent's Identity Spine certificate. The signed recommendation chain provides verifiable provenance for regulatory compliance (SEC Rule 17a-4).

Table 11: Identity Spine operational metrics.

Metric	Value
Certificate issuance throughput	8,400 certs/sec
Certificate issuance latency (p99)	4.2 ms
Capability verification latency (p99)	0.3 ms
OCSP revocation propagation time	1.8 sec
Certificate renewal success rate	99.97%
SAN capability parsing latency	0.1 ms
Concurrent agent capacity	250,000+

Table 12: Ablation study. Each row removes one component. Percentage changes relative to full system.

Configuration	Defense Cost	Portfolio Sharpe	EM RMSE (dB)
FORGE QBIT Full	1.019	1.19	1.37
− PhysicsCore	1.128 (+10.7%)	1.08 (−9.2%)	4.82 (+252%)
− QuantumSolve	1.074 (+5.4%)	1.14 (−4.2%)	1.37 (=)
− GraphIntel	1.019 (=)	1.19 (=)	1.37 (=)
− Warm-starting	1.052 (+3.2%)	1.12 (−5.9%)	1.37 (=)
− Problem decomposition	1.019 (=)	1.04 (−12.6%)	1.37 (=)
− Adaptive λ_{phys}	1.038 (+1.9%)	1.15 (−3.4%)	1.72 (+25.5%)
− Heterogeneous Crypto[a]	1.019 (=)	1.19 (=)	1.37 (=)
− HSM Enclave[b]	1.019 (=)	1.19 (=)	1.37 (=)
− Identity Spine[c]	1.019 (=)	1.19 (=)	1.37 (=)

[a]Optimization quality unaffected; removes cross-family resilience.
[b]Optimization quality unaffected; handshake latency increases 11×.
[c]Optimization quality unaffected; removes per-agent auth & capability attenuation.

Results. Sharpe ratio improvement from 0.94 to 1.22 (+29.8%), maximum drawdown reduction from 12.4% to 8.7%, $340K additional tax-loss harvesting captured, 6.5× rebalancing speedup, 12 new clients acquired.

8.3 Naval Communications Hardening

Setting. A U.S. Navy destroyer group transitioning satellite and HF radio communications to post-quantum cryptographic standards while maintaining interoperability with legacy systems and operating within bandwidth constraints (HF radio at 2.4–9.6 kbps).

Crypto Core deployment. Phased migration: Phase 1 deploys ML-KEM-768 in hybrid mode for satellite links (128 kbps+). Phase 2 implements optimized ML-KEM-512 for HF radio with aggressive key caching (initial handshake: 1,984 bytes; subsequent: 32-byte symmetric overhead). Phase 3 deploys ML-DSA-44 signatures with aggregated verification reducing per-message overhead by 73%.

Heterogeneous crypto. The destroyer group deploys all three crypto families: ML-KEM for high-bandwidth satellite links, HQC for diversity on shore-to-ship communications, and SLH-DSA for the most conservative signature posture on classified message traffic.

HSM Hierarchical Enclave. Shipboard HSMs (Thales Luna Marine variant) serve as the root tier. Enclave-level derivation runs on the ship's SGX-capable servers. Software-derived session keys enable the HF radio to operate with <1ms cryptographic overhead per message.

PQ Double Ratchet. Ship-to-ship tactical communications use the PQ Double Ratchet for forward-secret messaging in contested electromagnetic environments. The per-ratchet-step forward secrecy ensures that compromise of one ship's communications does not retroactively expose traffic from other ships in the group.

Results. Satellite link latency increased by only 8.3%. HF radio maintained operational throughput with cryptographic overhead below 4% of channel capacity. ML-DSA signature verification at 8,900

ops/sec exceeded the destroyer group's maximum message rate. FIPS 140-3 Level 2 equivalence assessment completed with zero critical findings.

9 Discussion

9.1 Key Findings

Four principal findings emerge from our evaluation.

First, **physics-informed constraints provide the single largest contribution to optimization quality**, confirming that domain knowledge encoded as differentiable PDE constraints is more valuable than quantum hardware advantages in the NISQ era. Removing PhysicsCore degraded performance 2–3$\times$ more than removing QuantumSolve, suggesting that practitioners should prioritize physics-informed formulation over quantum hardware access.

Second, **heterogeneous cryptography eliminates single-family systemic risk** without significant performance penalty. Cross-family rotation completes within 12.4 seconds with zero message loss, demonstrating that crypto-agility at the family level is operationally feasible.

Third, the **HSM Hierarchical Key Enclave achieves FIPS compliance at software-level latency**, resolving a fundamental tension in enterprise key management. The 22.6$\times$ throughput improvement over per-handshake HSM access makes FIPS-compliant PQC practical for high-throughput agent communications.

Fourth, the **Identity Spine enables the agent-legibility property** that defines FORGE OS. Without per-agent cryptographic identity, the platform cannot provide non-repudiable audit trails, capability-attenuated access control, or cross-subsystem trust verification. The Identity Spine transforms FORGE QBIT from a security layer into a platform foundation.

9.2 Comparison with Existing Platforms

FORGE QBIT occupies a unique position in the quantum computing and PQC landscape. Classiq and Qiskit provide excellent quantum algorithm development tools but lack physics-informed integration and defense-specific deployment capabilities. D-Wave's Leap platform offers powerful annealing solvers but as a single-backend service without physics constraints or cryptographic security. Proprietary defense systems provide operational security but lack quantum computing integration. No existing platform provides heterogeneous cross-family PQC with an agent identity spine.

9.3 Limitations

Several limitations merit acknowledgment. First, QAOA on current NISQ hardware provides modest advantages over classical heuristics for well-studied problems below 1,000 variables; the primary value is in physics-constrained formulations and future quantum scaling. Second, the PINN training pipeline requires domain expertise to select appropriate governing equations. Third, QGNN advantages are currently theoretical for most practical problem sizes. Fourth, the HSM Hierarchical Enclave requires hardware-specific secure enclaves (SGX/TDX or TrustZone), limiting deployment to compatible hardware. Fifth, the PQ Double Ratchet adds per-message overhead (8.7 μs) that may be non-negligible for extremely high-frequency edge device communications.

9.4 FORGE OS Integration

FORGE QBIT serves as the foundation for the entire FORGE OS platform. The following subsystems cannot function at full capability without FORGE QBIT:

- **FORGE CORE** requires Identity Spine certificates for model API authentication and response signing.
- **FORGE MEMORY** requires Merkle chain signing for the immutable governance object model.
- **FORGE KINETIC** requires PQ Double Ratchet for swarm-to-swarm communications and quorum certificates for BFT consensus.

This dependency map confirms FORGE QBIT's role as the immune system: its capabilities are not consumed directly by end users but provide the security substrate upon which all other subsystem operations depend.

10 Conclusion

We have presented FORGE QBIT, the security and identity engine of FORGE OS. The subsystem integrates five principal capabilities: a heterogeneous Crypto Core spanning three independent mathematical families with automatic cross-family rotation, an HSM Hierarchical Key Enclave achieving FIPS 140-3 compliance at software-level latency, a PQ Double Ratchet protocol for forward-secret quantum-resistant communications, an Identity Spine providing the root-of-trust for all FORGE OS agents, and the Physics-Constrained Quantum Optimization methodology bridging physics-informed modeling with quantum-accelerated solving.

Experimental evaluation demonstrates that FORGE QBIT achieves 19–42% improvement over classical baselines on combinatorial optimization tasks, 94.3% physical consistency in surrogate models, quantum-resistant key exchange with less than 12% latency overhead, cross-family rotation within 12.4 seconds with zero message loss, and Identity Spine throughput of 8,400 certificates per second supporting 250,000+ concurrent agents.

The platform's modular architecture ensures that FORGE QBIT benefits from quantum hardware improvements—increased qubit counts, improved gate fidelities, error correction advances—without architectural changes. As quantum computing transitions from NISQ to fault-tolerant systems, FORGE QBIT's physics-informed formulations position it to leverage genuine quantum advantage, while its heterogeneous cryptographic foundation provides resilience against the very quantum threats that motivate its computational capabilities.

For detailed technical specifications of the companion FORGE OS subsystems, see the FORGE CORE specification [577 Industries R&D Lab, 2025a], FORGE MEMORY specification [577 Industries R&D Lab, 2025b], and FORGE KINETIC specification [577 Industries R&D Lab, 2025c]. For the unified platform architecture, see the FORGE OS Spine document [577 Industries R&D Lab, 2025d].

References

BQP Simulation. Quantum Algorithms in Defence: Solving Complex Optimization. bqpsim.com, 2025.

BQP Simulation. Quantum Optimization Explained: Algorithms, Use Cases & Challenges. bqpsim.com, 2026.

Cloudflare. NIST's first post-quantum standards. blog.cloudflare.com, August 2024.

Fortune. D-Wave CEO shrugs off short attacks with revolutionary $550 million quantum computing acquisition. February 2026.

European Union. Regulation (EU) 2024/1689: The Artificial Intelligence Act. *Official Journal of the European Union*, 2024.

W. Fan and X. Chen. Embedding Physics into Machine Learning: A Review of Physics Informed Neural Networks. *Tsinghua Science and Technology*, 31(3):1326–1364, 2026.

E. Farhi, J. Goldstone, and S. Gutmann. A Quantum Approximate Optimization Algorithm. *arXiv preprint arXiv:1411.4028*, 2014.

577 Industries R&D Lab. FORGE Core: A Causal Model-Agnostic Intelligence and Routing Engine. Technical report, 577 Industries Incorporated, 2025.

577 Industries R&D Lab. FORGE Memory: A Provenance-Grounded Governance Engine. Technical report, 577 Industries Incorporated, 2025.

577 Industries R&D Lab. FORGE Kinetic: A Fractal Swarm Coordination and Edge Autonomy Engine. Technical report, 577 Industries Incorporated, 2025.

577 Industries R&D Lab. FORGE OS: The Agent-Legible Operating System — Unified Platform Specification. Technical report, 577 Industries Incorporated, 2025.

L. Lewis, D. Gilboa, and J. McClean. Quantum advantage for learning shallow neural networks with natural data distributions. *Nature Communications*, December 2025.

IBM Research. NIST's post-quantum cryptography standards are here. research.ibm.com, 2024.

Multiverse Computing. BBVA Quantum Portfolio Optimization: Sharpe Ratio Benchmarks. 2024.

Nature Computational Science. Quantum approximate multi-objective optimization. October 2025.

National Institute of Standards and Technology. Zero Trust Architecture. NIST SP 800-207, 2020.

National Institute of Standards and Technology. NIST Releases First 3 Finalized Post-Quantum Encryption Standards. FIPS 203, 204, 205, August 2024.

National Institute of Standards and Technology. Module-Lattice-Based Key-Encapsulation Mechanism Standard (FIPS 203). 2024.

National Institute of Standards and Technology. Module-Lattice-Based Digital Signature Standard (FIPS 204). 2024.

National Institute of Standards and Technology. Status Report on the Fourth Round of the NIST Post-Quantum Cryptography Standardization Process. NIST IR 8545, March 2025.

National Security Agency. Commercial National Security Algorithm Suite 2.0 (CNSA 2.0). Cybersecurity Advisory, 2022.

Qiskit Finance. Portfolio Optimization using QAOA. IBM Quantum Documentation, 2025.

M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

Signal Foundation. The Double Ratchet Algorithm. signal.org/docs, 2016.

SPIFFE. Secure Production Identity Framework for Everyone. spiffe.io, 2024.

Evolution of physics-informed neural networks: Recent architectural variants and optimization strategies. *ScienceDirect*, January 2026.

G. Verdon, T. McCourt, E. Luzhnica, et al. Quantum Graph Neural Networks. *arXiv preprint arXiv:1909.12264*, 2019.