

**Evolution of Trust in Federated Systems:  
Architecting Secure Identity and Authorization for  
Enterprise Exchange in an Agentic Era**

*A Technical Framework for  
Banking, Financial Services, and Cross-Enterprise Collaboration*

February 2026

Copyright 2026 NEO PKI

# Table of Contents

Table of Contents .....	2
1. The Crisis of Static Trust in Dynamic Systems .....	4
1.1 Why Traditional Trust Models Are Breaking .....	4
1.2 The Financial Services Imperative.....	4
2. The Layered Trust Architecture .....	4
2.1 Layer 1: Static Trust Anchors .....	5
2.2 Layer 2: Federated Workload Identity.....	5
2.3 Layer 3: Policy and Authorization .....	5
2.4 Layer 4: Agent Trust Semantics (Emerging) .....	6
3. SPIFFE as the Foundation for Workload Trust .....	7
3.1 Why SPIFFE Is Exceptionally Well-Aligned.....	7
3.2 Where SPIFFE Needs Augmentation.....	7
4. Bridging Identity and Authorization: OIDC and FAPI 2.0.....	8
4.1 Why Plain OAuth Is Not Enough .....	8
4.2 FAPI 2.0: OAuth Without Foot-Guns.....	8
4.3 The Composite Architecture.....	8
4.4 Practical Boundaries .....	9
5. ASCX9 Framework - X9 Intermediate CAs: Governance .....	9
5.1 The Mistake to Avoid .....	9
5.2 The Cryptographic Functions X9 ICAs Should Play .....	9
5.3 The Clean Separation.....	10
5.4 When X9 ICAs Are Needed .....	10
6. Building for the Agentic Future .....	11
6.1 Browsers as Agent Runtimes.....	11
6.2 The Role of Public CAs .....	11
6.3 Cross-Enterprise Federation at Scale .....	11
7. Recommendations for Financial Services and Enterprise Leaders .....	12
7.1 Immediate Actions .....	12
7.2 Medium-Term Strategic Moves .....	12
7.3 Long-Term Vision.....	12
8. Conclusion .....	13

## Executive Summary

The emergence of autonomous agents and cloud-native architectures has fundamentally challenged the traditional trust models that have governed financial services and enterprise exchanges for decades. Static, human-curated trust stores, the bedrock of browser-based commerce and traditional PKI, are proving inadequate for a world where machines make millisecond decisions, workloads are ephemeral, and trust relationships must be both cryptographically verifiable and dynamically revocable.

This whitepaper presents a comprehensive framework for federated trust that bridges the established practices of traditional financial exchanges with the emerging requirements of agentic systems. Drawing from real-world experience in launching the ASC X9 mTLS federation, we explore how workload identity, cryptographic binding, policy-driven authorization, and federation governance can be architected to enable secure, scalable, and auditable inter-enterprise exchange.

The core argument is straightforward but consequential: static trust stores remain necessary as bootstrap anchors, but the primary trust fabric must shift to dynamic, workload-centric identity systems like SPIFFE, hardened with authorization frameworks like FAPI 2.0, and governed by federation models that maintain legal accountability. This layered architecture—combining SPIFFE for identity, OIDC for interoperability, FAPI 2.0 for authority, and policy engines for context-aware decisions—represents the most viable path forward for high-stakes, regulated environments.

For executives and technical leaders in banking, payments, healthcare, and other regulated industries, this paper provides both strategic vision and tactical guidance for modernizing trust infrastructure while maintaining the assurance and auditability that regulators and stakeholders demand...add Board Members as well.

# 1. The Crisis of Static Trust in Dynamic Systems

## 1.1 Why Traditional Trust Models Are Breaking

For the past three decades, the internet's trust infrastructure has rested on a remarkably simple foundation: a small set of certificate authorities whose public keys are embedded in browsers and operating systems. This model, static trust stores maintained by human curators, worked because it matched the architecture of the web: relatively stable servers, human operators making decisions, and trust relationships that changed infrequently.

The shift to cloud-native, microservices-based architectures, and now autonomous agents, has broken every assumption underlying this model:

- **Ephemeral workloads:** Services spin up and down in seconds, not years. Identity must be as dynamic as the infrastructure.
- **Machine-to-machine decisions:** Agents make thousands of trust decisions per second without human intervention. Trust evaluation must be automated and contextual.
- **Workload-centric trust:** Identity must be bound to what is running, not to the host it runs on or the person who deployed it.
- **Bounded trust domains:** Global, universal trust is dangerous when agents operate autonomously. Trust relationships must be explicit, scoped, and revocable.

The result is clear: static trust stores remain necessary as bootstrap anchors, but they cannot serve as the primary trust fabric for systems where identity, authorization, and risk assessment must happen in real-time, at scale, across organizational boundaries.

## 1.2 The Financial Services Imperative

This challenge is particularly acute in financial services, where the stakes are highest. Payment networks, securities exchanges, and interbank communication systems require:

- **Legal accountability:** Every transaction must be attributable to a verifiable legal entity.
- **Non-repudiation:** Cryptographic proof of who authorized what action, when.
- **Auditability:** Comprehensive logs that satisfy regulators and forensic investigators.
- **Controlled risk exposure:** The ability to revoke trust instantly when fraud or compromise is detected.

Traditional PKI federation models, cross-certification hierarchies, policy OIDs, and name constraints, were designed for these requirements. But they assumed stability, human oversight, and infrequent changes. They cannot support the velocity and autonomy that modern financial systems demand.

## 2. The Layered Trust Architecture

The solution is not to abandon traditional trust mechanisms, but to recast them within a layered architecture where each layer addresses a specific trust question. This approach, drawing from both established practice and emerging standards, provides a clear path forward.

## 2.1 Layer 1: Static Trust Anchors

**Purpose:** Bootstrapping and interoperability

Browser roots, OS trust stores, public certificate authorities, and hardware security modules (TPMs, HSMs) continue to serve a critical function: they provide the initial trust anchors that enable systems to establish cryptographic relationships. These anchors change infrequently, are broadly recognized, and offer a stable foundation.

**However:** These are bootstrap mechanisms, not the runtime trust fabric. Once initial trust is established, systems must transition to more dynamic, contextual trust evaluation.

## 2.2 Layer 2: Federated Workload Identity

**Purpose:** Machine trust

This is where the paradigm shift occurs. Instead of binding identity to hosts (DNS names, IP addresses) or users (login credentials), workload identity systems bind cryptographic identity to what is running: **the specific service, application, or agent.**

**Technologies that enable this:**

- **SPIFFE (Secure Production Identity Framework for Everyone):** Provides cryptographic workload identity with clear semantics. SPIFFE IDs uniquely identify workloads, independent of network location or deployment mechanism.
- **SPIRE (SPIFFE Runtime Environment):** Automated issuance, rotation, and verification of workload credentials with minimal human intervention.
- **Kubernetes Service Accounts with OIDC:** Native workload identity for containerized environments.
- **Cloud Provider Workload Identity:** AWS IRSA, GCP Workload Identity Federation, Azure Workload Identity.
- **mTLS with short-lived certificates:** Cryptographic binding of identity to every connection, with credentials that expire in minutes or hours, not years.

**Critical properties:**

- Credentials are short-lived (measured in minutes or hours)
- Automatic rotation without human touch
- Cryptographically verifiable provenance
- **Federation as a first-class concept**

SPIFFE answers the question: *What is this workload, and who vouches for it?*

## 2.3 Layer 3: Policy and Authorization

**Purpose:** Decision making

Cryptography proves identity. Policy decides authority. In agent systems, knowing who a workload is does not automatically grant it permission to act. Authorization must be contextual, attribute-based, and continuously evaluated.

### Technologies that scale:

- **OPA (Open Policy Agent):** Policy-as-code with a declarative language (Rego) for expressing authorization rules.
- **Cedar (AWS):** Amazon's open-source authorization framework with formal verification.
- **Attribute-based access control (ABAC):** Decisions based on attributes of the requester, resource, and context (time, location, risk score, transaction value).
- **Intent-aware authorization:** Emerging systems that evaluate not just identity but declared purpose and expected behavior.

Policy engines answer the question: *Should I trust this workload for this task, right now, given what I know about risk and context?*

## 2.4 Layer 4: Agent Trust Semantics (Emerging)

**Purpose:** Safe autonomy

As agents gain greater autonomy, a new layer of trust semantics is emerging to address questions unique to agentic systems:

- **Delegation:** When can an agent delegate authority to another agent, and under what constraints?
- **Capability constraints:** What actions is this agent authorized to perform, and with what limits?
- **Verifiable intent:** Can the agent cryptographically prove what it intends to do before it acts?
- **Auditability:** Comprehensive, tamper-evident logs of agent actions and decisions.

This layer is still taking shape, but its contours are clear: browsers may evolve into agent runtimes, and when they do, they will need mechanisms to delegate user authority safely, revoke it granularly, and maintain accountability.

## 3. SPIFFE as the Foundation for Workload Trust

### 3.1 Why SPIFFE Is Exceptionally Well-Aligned

SPIFFE has emerged as the de facto standard for workload identity in cloud-native systems because it solves precisely the problems that traditional PKI does not:

- **Clear identity semantics:** SPIFFE IDs uniquely identify workloads in a structured, verifiable format. They are not DNS names, not user principals, but workload identifiers.
- **Automated credential lifecycle:** SPIFFE servers issue X.509 SVIDs (SPIFFE Verifiable Identity Documents) and JWT SVIDs automatically, rotating them without human intervention.
- **Federation as a first-class concept:** SPIFFE trust bundles enable explicit federation between trust domains, with clear boundaries and revocation semantics.
- **Separation of identity and authorization:** SPIFFE intentionally does not encode authorization policies in credentials. This clean separation allows policy engines to make context-aware decisions.

**SPIFFE's sweet spot:** It answers "what is this workload, and who vouches for it?" That is foundational. But it is not sufficient by itself.

### 3.2 Where SPIFFE Needs Augmentation

SPIFFE is an identity substrate, not a complete trust framework. To build production systems, particularly in regulated industries, several gaps must be addressed:

- **No native authorization semantics:** SPIFFE proves identity but does not define what that identity is allowed to do. Authorization must be layered on top.
- **No economic or commercial trust model:** SPIFFE does not address questions of legal liability, insurance, or financial responsibility that matter in banking and payments.
- **No consumer/browser trust integration:** SPIFFE is designed for backend systems, not for integration with browser trust stores or consumer-facing applications.
- **Federation governance is out-of-band:** SPIFFE trust bundles enable federation, but the governance—who decides which trust domains can join, under what terms—is not specified by the standard.
- **Does not express intent, risk, or business context:** SPIFFE identities are static assertions. They do not convey the purpose of a connection, the risk level of an operation, or business-specific constraints.

These are not criticisms of SPIFFE—they are by design. SPIFFE is focused, composable, and interoperable. The gaps are filled by other layers of the architecture.

## 4. Bridging Identity and Authorization: OIDC and FAPI 2.0

### 4.1 Why Plain OAuth Is Not Enough

OAuth 2.0 and OpenID Connect (OIDC) are ubiquitous standards for delegated authorization and federated authentication. But vanilla OAuth assumes:

- A human in the loop making authorization decisions
- Bearer tokens that can be used by anyone who possesses them
- Browser-based redirects for user consent
- Weak guarantees about client identity

Agent systems break every one of these assumptions. Agents operate autonomously, handle high-value transactions, and require strong non-repudiation. Bearer tokens are unacceptable when a stolen token can authorize a million-dollar transfer.

### 4.2 FAPI 2.0: OAuth Without Foot-Guns

FAPI (Financial-grade API) 2.0 is a security profile of OAuth 2.0 and OIDC that eliminates ambiguity and enforces best practices. Originally developed for financial services, it has been generalized for any high-security application. FAPI 2.0 is exceptionally prescriptive about how tokens are issued, bound, and used.

#### Key requirements:

- **Strong client authentication:** FAPI 2.0 requires `private_key_jwt` or mTLS for client authentication. No shared secrets, no anonymous clients. This aligns directly with SPIFFE-issued X.509 SVIDs and workload-bound keys.
- **Cryptographic binding everywhere:** Access tokens must be sender-constrained using DPoP (Demonstrating Proof-of-Possession) or mTLS binding. This prevents token replay, confused deputy attacks, and delegation abuse.
- **Explicit authorization semantics:** FAPI pushes OAuth toward capability-based authorization with fine-grained scopes, explicit consent artifacts, audience restriction, and proof-of-possession tokens.

**Why this matters:** FAPI 2.0 is already accepted in financial services, healthcare, and government APIs. For agents acting on behalf of users or handling regulated data, FAPI 2.0 provides the assurance that regulators demand.

### 4.3 The Composite Architecture

The right model for agent systems combines SPIFFE, OIDC, and FAPI 2.0 in a layered approach:

- SPIFFE establishes **who**: agent identity, workload provenance, cryptographic trust
- OIDC provides **interoperability**: federation across organizations, standard claims model, token exchange
- FAPI 2.0 hardens **authority**: how tokens are issued, how delegation works, how abuse is prevented

- Policy engines decide **should**: context, intent, risk, business rules

This gives you: strong identity + safe delegation + enforceable authority.

## 4.4 Practical Boundaries

FAPI 2.0 is not a silver bullet. It still assumes centralized authorization servers, does not solve policy reasoning, and is heavyweight for purely internal service meshes.

### Recommended pattern:

- **Inside a cluster:** SPIFFE + mTLS + OPA policy
- **Cross-boundary, high-value APIs:** OIDC + FAPI 2.0
- **User-delegated agents:** FAPI 2.0 becomes critical

If SPIFFE answers "who are you?", FAPI 2.0 answers "how may you act, and how do we know you didn't cheat?"

## 5. ASCX9 Framework - X9 Intermediate CAs: Governance

### 5.1 The Mistake to Avoid

There is a temptation to use X9 Intermediate Certificate Authorities (ICAs) as direct workload credential issuers—essentially rebuilding traditional PKI federation with new names. This would be a mistake. It reintroduces all the problems of classic cross-certification: large blast radius, policy OID complexity, path validation ambiguity, and browser-style constraints leaking into private trust.

**The right role for X9 ICAs is different:** they should cryptographically vouch for federation membership and compliance, not for individual workloads.

### 5.2 The Cryptographic Functions X9 ICAs Should Play

Think of X9 ICAs as trust-assurance anchors, not identity factories. Their cryptographic functions should be:

- **Trust domain attestation:** An X9 ICA issues a certificate asserting that a trust domain (e.g., a SPIRE server) is a recognized X9 federation member and meets compliance requirements. This certificate is not presented in workload mTLS. Instead, it is used to sign SPIFFE trust bundle distribution, gate federation onboarding, and support audit and dispute resolution.
- **Cryptographic binding of legal entity to trust domain:** X9 ICAs can issue certificates that bind a legal entity to a trust domain name and root public key, creating non-repudiable linkage for accountability and clear revocation semantics. SPIFFE intentionally does not solve this layer.
- **Federation admission and revocation control:** Instead of trusting a bundle because someone says so, trust bundles are signed by an X9-recognized ICA. SPIFFE federation still works without it, but governance becomes cryptographically enforceable with it.

- **Optional policy signaling:** X9 ICAs can embed high-level policy identifiers or compliance assertions, but these are consumed by federation controllers and auditors, not by workloads. This avoids the policy-in-cert failure mode of classic PKI.

### 5.3 The Clean Separation

The architecture looks like this:

[X9 ICA] → (certifies) → [Trust Authority / Trust Domain] → (issues) → [SPIFFE SVIDs] → [Workloads / Agents]

**Key principle:** No workload ever chains to the X9 ICA. No SPIFFE SVID needs to include X9 semantics. The X9 ICA never appears in runtime path validation.

**Why this improves assurance:**

- Smaller blast radius—compromise of a workload credential does not compromise federation governance
- Explicit federation boundaries with clear revocation semantics
- Cleaner audits as federation membership and workload activity are separately traceable
- You can rotate workload crypto without touching federation assurance

### 5.4 When X9 ICAs Are Needed

Be honest: if all federation members already have bilateral legal agreements, bundle exchange is tightly controlled, and audit requirements are light, an X9 ICA adds governance clarity but not cryptographic necessity. SPIFFE federation alone is technically sufficient.

**X9 ICAs become valuable when:**

- **Scale increases (dozens of members, not just a handful)**
- **Liability matters (financial transactions, regulated data)**
- **Regulators show up (audits, compliance reviews, dispute resolution)**
- **You need neutral third-party attestation (insurance, legal frameworks)**

**Bottom line:** X9 ICAs have a cryptographic role, but not as traditional issuing CAs. Their real value is in certifying who is allowed to run a trust authority, binding legal identity to trust domains, and anchoring federation governance with cryptographic evidence. Used this way, **SPIFFE handles identity, X9 handles assurance, and PKI stops being overloaded.**

## 6. Building for the Agentic Future

### 6.1 Browsers as Agent Runtimes

If, or more likely, when browsers evolve into agent runtimes, the trust model will need to support:

- **Delegated authority:** Users granting agents permission to act on their behalf, with granular and revocable constraints.
- **High-stakes operations:** Agents managing money, assets, and sensitive personal data.
- **Regulatory compliance:** Accountability to financial regulators, data protection authorities, and consumer protection agencies.

Browsers will not throw away their trust stores. Instead, they will use those roots to bootstrap federated, workload-scoped trust. The layered architecture described in this paper is designed for precisely that transition.

### 6.2 The Role of Public CAs

Public certificate authorities will not disappear. But their role may shift from issuing endpoint certificates to providing attestation and audit services. They may become trust brokers rather than identity issuers, vouching for the compliance and security posture of trust domains rather than signing individual workload certificates.

### 6.3 Cross-Enterprise Federation at Scale

The future of enterprise exchange is not monolithic trust hierarchies but explicit, selective, policy-driven federation. Organizations will:

- Run their own trust domains (SPIFFE, SPIRE, or equivalent)
- Federate selectively with partners based on legal agreements and risk assessment
- Use X9 ICAs (or equivalent) to anchor federation membership
- Rely on FAPI 2.0 for high-value, cross-boundary interactions
- Enforce policy dynamically based on context, not just identity

This is not a distant future. It is happening now in payment networks, interbank messaging, and healthcare data exchange. The question is not whether to adopt this model, but how quickly and how well.

## 7. Recommendations for Financial Services and Enterprise Leaders

### 7.1 Immediate Actions

- **Deploy SPIFFE/SPIRE for internal workload identity:** Start with internal microservices and cloud-native applications. Prove the technology works at scale before federating externally.
- **Adopt mTLS with short-lived certificates:** Replace long-lived certificates and bearer tokens with cryptographically bound, automatically rotated credentials.
- **Implement policy-as-code:** Deploy OPA or Cedar to separate authorization logic from application code. Test context-aware, attribute-based access control.
- **Pilot FAPI 2.0 for external APIs:** Start with partner integrations that handle sensitive data or high-value transactions. Measure the reduction in security incidents.

### 7.2 Medium-Term Strategic Moves

- **Establish federation governance:** Define clear policies for who can join your trust domain, under what terms, and with what audit requirements. Consider X9 ICA attestation for legal accountability.
- **Federate with key partners:** Start with bilateral SPIFFE federation for API integrations. Use trust bundles to establish explicit, revocable trust relationships.
- **Build intent-aware authorization:** Move beyond 'who' to 'what for.' Require agents to declare their purpose and validate it against policy before granting access.
- **Engage with standards bodies:** Participate in ASC X9, OpenID Foundation, and SPIFFE/SPIRE working groups. Shape the standards to meet your industry's needs.

### 7.3 Long-Term Vision

- **Prepare for agentic systems:** Design your trust infrastructure to support autonomous agents, not just human-operated systems. Test delegation, revocation, and auditability at scale.
- **Advocate for browser-agent trust models:** Work with browser vendors and standards organizations to define how browsers will manage trust for user-delegated agents.
- **Build multi-layered defense:** Do not rely on any single trust mechanism. Layer SPIFFE, OIDC, FAPI 2.0, policy engines, and X9 governance to create defense in depth.
- **Invest in observability and forensics:** Comprehensive logging, tamper-evident audit trails, and real-time anomaly detection are essential for agent systems. You cannot audit what you cannot see.

## 8. Conclusion

The shift from static, human-curated trust to dynamic, machine-mediated trust is not a distant possibility, it is the present reality in cloud-native architectures and the immediate future in agentic systems. For financial services, healthcare, and other regulated industries, this transition must be managed carefully, preserving the accountability and auditability that stakeholders demand while embracing the velocity and autonomy that modern systems require.

The layered trust architecture presented in this paper offers a clear path forward:

- **Static trust anchors** provide bootstrapping and broad interoperability.
- **SPIFFE** establishes cryptographic workload identity with federation as a first-class concept.
- **OIDC and FAPI 2.0** enable secure, interoperable authorization across organizational boundaries.
- **Policy engines** make context-aware, risk-informed decisions in real-time.
- **X9 ICAs and governance frameworks** provide legal accountability and audit-ready attestation.

This is not a replacement for existing trust infrastructure—it is an evolution. Public CAs, browser trust stores, and traditional PKI continue to play important roles, but they are recast as bootstrap mechanisms rather than the primary runtime trust fabric.

The organizations that succeed in this transition will be those that recognize the paradigm shift early, invest in the right technologies, and build governance frameworks that enable federation at scale.

**The ASC X9 mTLS federation is a pioneering example of this approach in action, demonstrating that financial institutions can collaborate securely without centralized control or monolithic trust hierarchies.**

**The future of trust is federated, workload-centric, policy-driven, and agentic.** The question is not whether to embrace this model, but how quickly your organization can adopt it while maintaining the assurance and accountability that your stakeholders and regulators expect.

- *Static trust stores anchor legitimacy*
- *SPIFFE enables velocity*
- *Federation and policy make agent autonomy safe*