primitive

AGENTIC FINANCIAL SERVICES PROTOCOL (AFSP)

AFSP Specification

Version 0.1 — Public Review Draft

Status: Open for public comment through November 16, 2026
Published September 28, 2026 · agenticfinanceprotocol.org · https://github.com/primitive-os/AFSP

> **Important Notice:** This document is a draft specification released for public comment. It does not constitute a final standard. Primitive invites review and comment from financial institutions, AI platform operators, data aggregators, card networks, and regulatory bodies. Nothing in this document creates legal obligations on any party. The technical specifications herein are subject to revision based on implementation experience and working group feedback.

License: Fair, Reasonable and Non-Discriminatory (FRAND) terms · See Section 17.5 · Feedback: afsp@primitive.com

Contents

Executive Summary 2
1. Introduction 5
2. AFSP Standard Components — Complete Architecture Overview 9
3. Security Considerations 11
4. Extension Process and Registry 15
5. Consumer Layer Operator Architecture and Discovery Routing 16
6. The AFSP Attestation Package 22
7. AFSP-01: Consumer Authorization and Biometric Binding 24
8. AFSP-02: Financial Profile Attestation 28
9. AFSP-03: Reverse Inquiry Protocol and Fit-Based Marketplace 32
10. AFSP-04: Clearance Return Protocol and Session Token 35
11. AFSP-05: Behavioral Session Intelligence Interface 37
12. AFSP-06: Non-Participant Disclosure Standard 40
13. Return Trip API 42
14. Certification and Participation Levels 42
15. Regulatory Considerations 43
16. AFSP v1.0 Roadmap 44
17. Governance, Versioning, and Feedback 45
18. Items for Further Consideration 49
Appendix A — Non-Normative: End-to-End Worked Example 51
Appendix B — References 57

Executive Summary

The Agentic Financial Services Protocol (AFSP) is an open industry protocol for one moment: when an AI agent arrives at a regulated financial institution on behalf of a person. Before the agent reaches the institution's systems, AFSP lets the institution verify who the agent is, who it is acting for, and that the person is present and in control. Version 0.1 covers the first use case, agent-assisted opening of a new deposit account, and is open for public comment until November 16, 2026.

Key points

AI agents are becoming a banking channel, and banks have no standard way to say yes safely. AFSP gives an institution verified answers about the agent, the person behind it and whether that person is present, before the agent reaches the institution's systems.

Standards grow channels, not just protect them. Card payments, online banking and open banking each grew sharply once a shared standard made them safe to join; accounts using the FDX API standard roughly quadrupled in about three years.[1] AFSP is designed to do the same for AI agents in banking.

It fills a gap existing standards do not cover. OAuth, FAPI, FDX and FIDO assume a relationship already exists; AFSP governs the moment an agent asks to start one.

The bank stays in charge. Every identity signal traces to a regulated institution that already knows the consumer, and every KYC and account decision stays with the receiving institution.

The consumer stays in control, and existing relationships get first consideration. The consumer approves each consequential step on their own device, through a confirmation generated outside the AI model. Requests go first to the institutions the consumer already banks with; others only if the consumer chooses to look further, and none learns who the consumer is until chosen.

Built with the industry and open for comment now. AFSP emerged from hundreds of conversations across banks, data aggregators, technology and institutional platforms, and strategic consultancies. Founding Endorsers will be announced ahead of the November 16, 2026 comment deadline, working groups begin in November, first proofs of concept are targeted for Q1 2027, and governance passes to an independent AFSP Foundation targeted for early 2027.

Why the industry needs it

AI agents are already searching, comparing and buying on people's behalf. Consumers will soon ask them to find a better savings account or move money, and those agents will arrive at banks ready to act. Today there is no agreed industry standard for deciding whether to let one in, so the only safe answer for a bank is no. Opening a new agentic banking channel requires the industry to act now.

If the industry does not agree on an answer, one will arrive anyway: through bilateral deals between every bank and every AI platform, through rules set by the largest platforms, or through regulation after the fact.

Existing standards do not cover this moment. OAuth, FAPI, FDX and FIDO assume a relationship already exists. A new generation of agentic commerce protocols — covering agent-to-agent messaging, agent payments, and agent identity — is emerging to govern what an agent does once it is already acting on an open account. AFSP is complementary to that work, not competing with it: it occupies the pre-credential space, the moment before any of those protocols apply, when an agent is asking a bank to begin a relationship for someone the bank has never met.

Every trusted banking channel has followed the same arc. The technology arrives first, and usage stays small while each institution carries the risk alone. Then a shared standard makes the channel safe to join,

[1] Financial Data Exchange (FDX) press releases, January 24, 2022 (28 million consumer accounts) and April 25, 2025 (about 114 million customer connections, which FDX treats as comparable), financialdataexchange.org.

and usage multiplies. Card payments scaled once common network rules let any card work at any participating merchant. Online and mobile banking went mainstream once shared security standards gave institutions and customers confidence. Open banking is the most recent proof: according to FDX, consumer accounts using its API standard grew from 28 million in early 2022 to about 114 million in 2025, and UK open banking, built on a single API standard, has passed 100 billion API calls and one billion payments since its 2018 launch.[2] In each case the standard did more than reduce risk. It grew the channel for everyone in it: more customers for institutions, more usage for platforms and more choice for consumers.

The agentic banking channel is at that inflection point now and we have a unique opportunity to define the new network standard together in a way that keeps consumer relationships sovereign, banks as the trust anchor, and regulatory guidance as the north star.

What AFSP does

AFSP answers the five questions an institution must answer before admitting an agent session: who is this agent; is it acting for a real person who authorized it; has that person's identity been verified; is their financial profile established; and are they present and in control right now. The answers arrive as a cryptographically signed package, verified in real time, before any open banking API is called.

AFSP decides nothing about the person. Every KYC, underwriting and account-opening decision stays with the institution, under its own regulatory accountability. The principle throughout is simple: the agent presents, the human decides, the bank executes.

The worked example in Appendix A shows a consumer going from "find me the best high-yield savings account" to an open account in under three minutes from authorizing the session. She confirms each consequential step, and the bank runs its own KYC.

What makes AFSP different

Identity anchored in regulated institutions. Every signal traces back to a bank or credit union that has already verified the consumer, not to AFSP or any single company.

Authorization on the consumer's own device. An agent can never act beyond what the person explicitly authorized, confirmed with a biometric check on their own device.

A confirmation the AI model cannot rewrite. Before any consequential action, the consumer sees a plain-language summary generated from the request itself, outside the AI model, so a manipulated agent cannot disguise what is being approved.

Existing relationships first. By default, a request goes only to institutions the consumer already banks with, giving existing relationships first consideration before an account moves. Other institutions are considered only if the consumer chooses to look further, and they compete without seeing who the consumer is.

Examination-ready by design. Every session produces a tamper-evident record of who authorized what, when and on whose behalf, making examination easier, not harder.

What it offers each participant

Consumers: an AI agent can act for them with a bank, starting with opening an account, with proof that it is working for them and that they are still in control.

[2] FDX press releases of January 24, 2022 and April 25, 2025; Open Banking Limited, "Open Banking surpasses one billion payments and 100 billion API calls," press release, July 30, 2026 (CMA9 banks, since the January 2018 launch).

Financial institutions: a verified, pre-screened session before their own processes begin, with no change to how they perform KYC or make decisions, and first consideration for their existing customers.

AI platforms and agent operators: one certification and one integration pattern to reach participating institutions, instead of a separate arrangement with each.

Regulators and examiners: a tamper-evident record of who authorized what, when, on whose behalf, and with what signals present. AFSP creates no new data rights and no new decision-making authority. AFSP is built to align with regulatory guidance and examiner expectations, not to work around them.

What is open, and who does what

AFSP is open governance, not open source. The specification is public and free to read, reproduce and comment on. Any organization with a genuine interest in the agentic banking channel can join the working groups that will decide what it becomes.

The AFSP Foundation, an independent non-profit targeted for formation in early 2027 (with final timing set by the independent entity itself), will govern the standard. It ratifies and maintains the specification, runs the working groups, and administers certification and the participation registry. Primitive is the founding author and acts as interim administrator until the Foundation is formed; after that, it is one voice in the Foundation's governance.

Primitive developed the underlying technology, holds the patent filings and licenses the claims essential to implementing AFSP on FRAND terms, with license terms published before v1.0. It will also offer commercial platform services, including the network's initial marketplace, under separate agreements that sit alongside the standard, not inside it. Institutions can use AFSP in their own branded channels without routing to any marketplace, consumers remain free to compare, and certifying additional marketplace operators is under consideration (Section 18).

How v0.1 is structured

This document is AFSP v0.1, the first version of the standard. It specifies six of AFSP's twelve core components (AFSP-01 to AFSP-06) in full, enough to deliver a complete, end-to-end agentic account opening, and sets the stage for the full agentic network: the other six (AFSP-07 to AFSP-12) supply the shared network services, from the agent registry to the regulatory audit log, that turn one use case into a complete agentic banking channel. The Technical Working Groups will develop them to reflect industry needs and practices. Section 2 summarizes all twelve and maps the rest of the document.

What v0.1 does not yet include. The protocol binding (transport, endpoints and message formats), a reference implementation and conformance tests follow in later versions, starting with the working groups in November 2026 and the first proofs of concept against the network specifications in Q1 2027. v0.1 is a draft for scrutiny, not a finished standard. Section 18 lists the open items, including how a bank verifies a cloud-hosted agent, and we want the industry's help solving them.

How to get involved

Comment by November 16, 2026: https://github.com/primitive-os/AFSP or afsp@primitive.com

Endorse or join a working group: afsp@primitive.com. No implementation commitment required.

1. Introduction

1.1 Purpose and Scope

The Agentic Financial Services Protocol (AFSP) defines the pre-credential trust framework governing how AI agents interact with regulated financial institutions on behalf of consumers. AFSP addresses a specific, bounded problem:

> **AFSP governs one question only:** Should this AI agent session proceed to the open banking API layer?

AFSP does not govern what happens after that decision. OAuth authorization, consumer identity verification, Know-Your-Customer (KYC) processes, credit underwriting, and account opening decisions remain entirely within the receiving financial institution's existing processes and regulatory accountability. AFSP is designed to provide the governed front door that the agentic banking channel does not yet have. Version 0.1 addresses the first use case: agent-assisted opening of new deposit accounts.

The protocol answers five questions that every regulated financial institution must answer before opening its open banking rails to an AI agent session:

SIGNAL	QUESTION	WHAT IT VERIFIES
S1	Who is this agent?	The AI agent is a certified, registered, unmodified software build from an accountable operator
S2	On whose behalf — and did they authorize it?	A real human, on their enrolled device, explicitly authorized this specific session
S3	Has their identity been verified?	A trusted network institution has previously KYC'd this consumer (cryptographic attestation only, no raw data)
S4	Are they financially real?	The consumer's behavioral financial signals are consistent with an established consumer profile
S5	Are they present right now?	This session is currently human operated within normal behavioral parameters at this moment

1.2 Relationship to Existing Standards

AFSP is designed to complement — not replace — existing financial services standards. Governance will transition from Primitive to a non-profit foundation as described in Section 17. AFSP builds on established frameworks as follows:

OAuth 2.0 / FAPI 2.0: AFSP operates before the OAuth authorization flow begins and is designed to complement FAPI 2.0. How AFSP artifacts bind to the subsequent authorization request (for example, through Rich Authorization Requests, RFC 9396) will be specified by the Technical Working Group.

Section 1033 (Dodd-Frank): AFSP Signal 4 financial identity attestation is derived from consumer-permissioned financial data accessed under Section 1033 consumer data rights. AFSP does not create new data rights — it operates within the rights framework Section 1033 establishes.

FDX API Standard: AFSP governs the pre-credential session trust layer. FDX governs the data sharing layer within an established relationship. They are adjacent and complementary — AFSP answers whether a session should proceed; FDX governs what data flows within that session.

ISO 20022: Where AFSP data accompanies ISO 20022 payment messages, such as initial account funding, it MAY be carried in the Supplementary Data element. AFSP audit records are signed, tamper-evident, and exportable in a documented format for examination.

FFIEC Authentication Guidance: AFSP pre-access attestation is designed to support a receiving institution's layered, risk-based approach to authentication and access under the FFIEC's 2021 guidance. Each institution remains responsible for its own risk assessment and control decisions.

1.3 Design Philosophy and Regulatory Foundation

AFSP is designed to embrace regulatory frameworks, not work around them. Three foundational principles inform every architectural decision in this specification and explain why the protocol is structured as it is. These principles are stated here, at the outset, because they are the lens through which every component should be read. Section 15 addresses specific regulatory frameworks and compliance considerations in detail.

Principle 1 — Consumer Sovereignty. The consumer is the sovereign authority in every AFSP-governed session. The consumer defines the scope of their agent's authority through explicit biometric authorization. The consumer selects which institution receives their financial identity. The consumer makes every consequential decision. No AFSP component enables an agent to act beyond what the consumer has explicitly authorized. Constraining the agent to the authority the consumer explicitly granted, and making that authority verifiable by the receiving institution, is itself a consumer protection.

Principle 2 — The Regulated Institution as Trust Anchor. Every AFSP signal is ultimately tethered to a regulated institution. Signal 2 (biometric consent) is bound to a root key accepted by an institution that has already KYC'd the consumer. Signal 3 (federated referral) is an attestation from another regulated network institution. Signal 4 (financial identity) is derived from consumer data held at regulated institutions under Section 1033. The AFSP platform delivers its pre-credential trust signals to a regulated institution that makes every consequential decision independently under its own regulatory accountability. The bank is not a passive endpoint — it is the trust anchor to which every signal in the protocol is ultimately tethered. AFSP makes the bank's regulatory accountability more legible, not less.

Principle 3 — Protocol as Regulatory Reinforcement. AFSP does not create regulatory arbitrage. It creates a documented, cryptographically auditable chain of custody for consumer intent and agent authorization that makes regulatory examination of agentic banking sessions easier, not harder. The AFSP audit record (AFSP-04, AFSP-11) captures every signal present at session initiation, every authorization event, and every execution action in a tamper-evident format compatible with BSA/AML lookback requirements. Every agent session produces a verifiable, examiner-accessible record of who authorized what, when, on whose behalf, and with what trust signals present.

1.4 What AFSP Is Not

To prevent mischaracterization, the following clarifications apply to all AFSP implementations:

> **AFSP is NOT intended to be a model.** AFSP is designed as an access-governance protocol. It delivers attested trust signals that a receiving institution may consider in deciding whether a session proceeds to the institution's own KYC and account opening processes. AFSP does not make credit, eligibility, KYC, or account opening decisions, and it does not prescribe how an institution weighs AFSP signals; each institution decides whether and how to use them under its own policies.

> **AFSP is NOT a fraud score.** Signal 5 is a session integrity signal that estimates whether the session is being operated by a human in a normal manner. It does not score the consumer's fraud risk or creditworthiness.

> **AFSP does NOT replace KYC.** The receiving financial institution retains full KYC accountability and decision authority. AFSP delivers pre-credential trust signals that may inform the institution's KYC process. The institution makes every KYC, underwriting, and account opening decision independently. AFSP is not present in those decisions.

> **The agent does NOT choose for the consumer.** It presents offers ranked by the published FitScore formula using the consumer's own preference weights, discloses the basis of the ranking, and requires explicit consumer confirmation for every consequential action. The agent presents. The human decides. The bank executes.

1.5 Terminology

AI Agent: A software system that acts autonomously on behalf of a consumer within defined authorization parameters. In the context of this specification, an AI agent is the entity that initiates a financial product inquiry session on behalf of a consumer.

AFSP Platform: The trust infrastructure operator that assembles PCAA attestation packages, operates the agent certification registry, manages the participation registry, and routes structured financial product inquiries. In v0.1, this role is performed by Primitive, which provides platform services under the Platform Access Agreement (Section 17.2).

Consumer Layer Operator (CLO): The entity operating the consumer-facing interface through which the consumer authorizes an AI agent session. The consumer layer operator and the AFSP platform may be the same entity or distinct entities. This specification is not limited to any specific entity performing either role.

Receiving Institution: A financial institution that receives AFSP attestation packages and independently evaluates them as part of its pre-KYC access governance.

PCAA Package: Pre-Credential Agent Attestation package — the assembled collection of five trust signals delivered to the receiving institution before any open banking API endpoint is invoked.

Return Trip API: The mechanism by which a receiving institution returns a signed confirmation of account opening outcome to the AFSP platform, enabling closed-loop conversion analytics and network intelligence.

Agent Operator: The entity that develops or operates the AI agent that initiates an AFSP session. The agent operator and the CLO may be the same entity (Section 5.2).

Issuing Institution: The regulated institution that has previously verified the consumer's identity and that binds the consumer's enrolled device key to its KYC record at enrollment (Section 7.2).

Signal Provider: A certified signal or endpoint operator contributing to one or more of Signals 1–5, including a Signal 4 financial profile attestation provider (S4 provider) and a Signal 5 session score provider (S5 provider).

AFSP Foundation: The Agentic Financial Services Protocol Foundation, the independent non-profit body intended to govern the specification and administer certification (Section 17.1). Before its formation, Primitive acts as interim protocol administrator.

Network Participation Agreement (NPA): The agreement with the AFSP Foundation that governs an institution's certified participation (Section 17.2).

Platform Access Agreement (PAA): The commercial agreement with Primitive that governs access to its platform services (Section 17.2).

Intent Verification Checkpoint (IVC): The mandatory plain-language confirmation of a specific Authorization Tier 2 or Authorization Tier 3 action before it is signed (Section 5.8).

Human-in-the-Loop (HITL): Routing of a session to explicit human reauthorization, step-up verification, or review (Sections 5.8 and 11.4).

Session Continuity Token (SCT): The token that authenticates agent re-entry sessions for accounts opened through AFSP (AFSP-12, Section 5.7; v1.0 roadmap).

Non-Decline Principle: The AFSP behavioral governance principle stating that when a Signal 5 session integrity anomaly is detected — including during post-opening re-entry sessions — the session SHALL NOT be automatically declined. Instead, the anomalous session is routed to human reauthorization, step-up verification, or enhanced review. This reflects the consumer sovereignty principle: a system that may incorrectly flag a legitimate consumer session must give the consumer the opportunity to reconfirm, not unilaterally block access. Automatic session termination without human reauthorization opportunity is not AFSP-compliant behavior for Signal 5 anomaly handling. Suspension of a session pending out-of-band reauthorization through the receiving institution (Section 11.4, Geographic Risk Tier 5) is consistent with this principle.

SHALL / MUST: Absolute requirement. Non-conforming implementations are not AFSP-compliant.

SHOULD: Recommended practice. Deviation is permitted with documented rationale.

MAY: Optional capability. Conforming implementations may or may not implement.

1.6 Roles and Responsibilities

The following table summarizes the roles in the AFSP network and who performs them in v0.1.

ROLE	WHO	RESPONSIBILITIES
AFSP specification	Open protocol (this document)	Defines the trust signals, attestation package, session flows, and requirements that participants implement.
AFSP Foundation	Independent non-profit (Section 17.1)	Will govern the specification, certify participants under the Network Participation Agreement, and govern the Participation Registry (Section 17.2).
Primitive	Founding author and platform operator	Serves as interim protocol administrator until the Foundation is formed; operates the AFSP platform and the network's initial marketplace under the Platform Access Agreement (Section 17.2); holds the patent applications described in Section 17.5.
Receiving institutions	Banks and credit unions	Make every decision on whether a session proceeds, on KYC, and on account opening; declare the products they offer through AFSP (Section 12.2).

ROLE	WHO	RESPONSIBILITIES
Consumer Layer Operators (CLOs)	AI platforms, bank-embedded assistants, widgets, and platform companies (Section 5.3)	Present sessions to the consumer, obtain the consumer's authorization, and route discovery as specified in Section 5.5.
Signal providers	Certified signal and endpoint operators contributing to Signals 1–5	Produce signal attestations and integrity protections in the defined formats, as needed to fulfill each signal's intended purpose.
Consumer	The sovereign authority (Section 1.3, Principle 1)	Authorizes each session, decides where to look for products and which institution receives their identity, and makes every consequential decision.

2. AFSP Standard Components — Complete Architecture Overview

AFSP comprises twelve standard components organized across two publication phases: eleven core components (AFSP-01 through AFSP-11), plus AFSP-12 (Post-Opening Session Continuity). The six v0.1 components are specified in enough detail for prototyping and implementation feedback. Several dependencies, including the agent provenance attestation format, the Discovery Declaration companion specification, and the interim EKU guidance, will be completed through the working group before v1.0. Version 0.1 defines the AFSP architecture and information model; the protocol binding — transport, endpoints, message serialization, signature format, and error handling — will be specified through the working group (Section 18). Six v1.0 roadmap components are architecturally defined and IP-protected: five (AFSP-07 through AFSP-11) covered by the original provisional filings, plus AFSP-12 (Post-Opening Session Continuity) identified as a further extension with a separate provisional filing. The AFSP Standard components are described in and covered by provisional patent applications filed by CiDR Technologies, Inc. (D/B/A Primitive).

Together the twelve components answer the five pre-credential trust questions from Section 1.1 and govern the full lifecycle of an AFSP-governed agent session from initial certification through ongoing transaction monitoring and regulatory audit.

AFSP v0.1 — Six Components for Publication with Founding Endorsers

COMPONENT	NAME	SIGNAL(S)	WHAT IT GOVERNS	SECTION
AFSP-01	Consumer Authorization and Biometric Binding	S1, S2	Hardware-bound biometric consent ceremony; authorization token with scope, action classes, and temporal validity; step-up reconfirmation for out-of-envelope actions	Section 7
AFSP-02	Financial Profile Attestation	S4	Six-stage privacy-preserving attestation pipeline from Section 1033-permissioned data; differential privacy (Laplace epsilon=1.0; k-anonymity k=5); ECDSA-P256 signed output; three-layer replay prevention	Section 8
AFSP-03	Fit-Based Reverse Inquiry Marketplace	S3, S4	Reverse inquiry payload schema; FitScore formula with default weights; promotional weighting structural exclusion; tiered participation states (Default, Expanded, Non-Participant); non-participant disclosure	Section 9

COMPONENT	NAME	SIGNAL(S)	WHAT IT GOVERNS	SECTION
AFSP-04	Clearance Return Protocol and Session Token	All signals	Post-selection handshake; institution clearance evaluation; session-scoped authorization token minting; platform KYC absence enforcement; tamper-evident audit record generation	Section 10
AFSP-05	Session-Initiation Behavioral Score	S5 (initiation)	Vendor-agnostic abstraction layer for behavioral session intelligence at session initiation; weighted ensemble normalization; ECDSA-P256 attestation; six signal categories. Note: moment-level continuous detection is AFSP-10 (v1.0)	Section 11
AFSP-06	Non-Participant Disclosure Standard	Network governance	Participation registry; factual disclosure when named institution is not a network participant; three consumer routing options; stated-preference data logging; non-participant analytics	Section 12

AFSP v1.0 Roadmap — Six Components (IP Filed, Working Group Development)

COMPONENT	NAME	SIGNAL(S)	WHAT IT GOVERNS	SECTION
AFSP-07	Agent Provenance and Certification Registry	S1	Full certification authority infrastructure; three-tier credential hierarchy (Discovery/Transactional/Exception); cryptographic scope enforcement via Extended Key Usage; real-time revocation propagation	Section 16 (roadmap)
AFSP-08	Federated Credential Verification	S1, S3	Institution verifies agent credentials via federated registry without bilateral relationships or KYC data transmission; directly answers S1 and S3	Section 16 (roadmap)
AFSP-09	Pre-Execution Revalidation Protocol	S5 (continuous)	Every agent-initiated execution intercepted before execution; synchronous multi-factor revalidation; no bypass path; session-airlock enforcement; reinforces downstream payment protocols via tamper-detection baseline	Section 16 (roadmap)
AFSP-10	Behavioral Envelope Engine (Moment-Level)	S5 (moment-level)	Per-credential ML behavioral baseline at every transaction execution; non-decline principle; cross-credential federated dataset; correct IP home for moment-level continuous detection (complements AFSP-05 session-initiation score)	Section 16 (roadmap)
AFSP-11	Tamper-Evident Regulatory Audit Log	Compliance	Cryptographically chained append-only log across credential lifecycle; 7-year retention (exceeds the BSA's general five-year period); extends AFSP-04 from single clearance event to full lifecycle chain	Section 16 (roadmap)
AFSP-12	Post-Opening Session Continuity	S2, S5	Three-leg triangulation re-authentication for agent-mediated servicing of AFSP-opened accounts: biometric continuity (same TEE root key), behavioral continuity (Signal 5 baseline comparison), relationship confirmation (Return Trip API cross-reference). Session Continuity Token; Re-Entry PCAA Package; SERVICE action class extension. See §5.7.	Section 5.7 (roadmap)

Section 3 (Security Considerations) consolidates the threat model and adversary analysis that informs all subsequent technical specifications. Section 4 (Extension Process) defines how the protocol evolves. Section 5 (Consumer Layer Operator Architecture and Discovery Routing) defines the discovery layer, entry point types, and consumer disclosure requirements. The detailed technical specifications for AFSP-01 through AFSP-06 follow in Sections 6 through 12. Section 13 (Return Trip API) and Section 14 (Certification and Participation) apply across all components. Section 15 (Regulatory Considerations), Section 16 (AFSP v1.0 Roadmap), Section 17 (Governance), and Section 18 (Items for Further Consideration) complete the document. Appendix A provides a non-normative end-to-end worked example, and Appendix B lists references.

3. Security Considerations

This section consolidates the security analysis for AFSP v0.1. It describes the adversary model, the attack surface, the mitigations built into the protocol, and the residual risks that implementing organizations should address through their own security programs. This section is normative for the purpose of security evaluation; the specific mitigations described are required unless otherwise stated.

3.1 Adversary Model

AFSP is designed to defend against the following adversary classes:

A1 — Fraudulent agent operator: An operator who deploys an AI agent that is not enrolled in the AFSP agent certification registry, or who deploys a modified version of a certified agent. Goal: access open banking APIs as if operating a legitimate agent. *Additionally, the three-tier Authorization Scope Envelope (Section 7.4) provides a second independent layer of protection: a fraudulent agent cannot obtain a valid INITIATE or EXECUTE payload, because only the consumer's hardware-held tier keys can produce one, each higher-tier signature requires consumer user verification, and verifiers reject out-of-tier signatures (VAL-08a).*

A2 — Synthetic identity attacker: An attacker who presents fabricated consumer identity and financial profile data to qualify for financial products. Goal: pass pre-credential attestation without a genuine consumer relationship.

A3 — Session hijacker: An attacker who intercepts or replays a valid PCAA package to gain access to a financial institution's open banking API for an unauthorized purpose. Goal: reuse a legitimate attestation for a different session or institution. Additionally, the credential constraint layer of the Authorization Scope Envelope (Section 7.4) limits the blast radius of a successful replay: a replayed package is constrained to the tier of the credential embedded in it and cannot be escalated to a higher tier.

A4 — Coercion attacker: An attacker who forces or deceives a legitimate consumer into authorizing an agent session they do not intend. Goal: pass biometric authorization (Signal 2) with a non-consenting consumer. Additionally, the Intent Verification Checkpoint (Section 5.8) requires the consumer to confirm the specific action in plain language before any INITIATE or EXECUTE payload is signed. This helps a consumer who is being deceived notice a mismatch; it does not stop a coercer who dictates the confirmation.

A5 — Provider compromise: An attacker who compromises a certified Signal 4 or Signal 5 provider's signing key. Goal: issue fraudulent attestations that pass signature verification.

A6 — Platform compromise: An attacker who compromises the AFSP platform operator's signing key. Goal: issue fraudulent PCAA packages that pass receiving institution validation.

A7 — Prompt injection via financial data: An attacker who embeds malicious instructions in financial data fields accessible via Section 1033 data pull — including merchant names, transaction descriptions, and memo fields — with the goal of manipulating AI agent behavior during Signal 4 assembly. Unlike A1 through A6, this adversary does not attack the AFSP protocol directly; they attack the data pipeline that feeds into it, attempting to cause the agent to take unauthorized actions by corrupting the financial data from which behavioral signals are derived. Mitigation: AFSP-02 Stage 2 normalization (Section 8.2) strips all free-text fields before signal derivation. Instructions embedded in financial data therefore cannot influence the S4 attestation. Stage 2 does not protect the agent itself: an agent with READ access may retrieve raw transaction text, so CLOs SHOULD treat all retrieved financial data as untrusted input. Additionally, the Intent Verification Checkpoint at Authorization Tier 2 and Authorization Tier 3 (Section 7.4) provides a second independent layer against prompt injection reaching the action-execution layer: the IVC confirmation is generated from structured fields, not model output, so a redirected action shows the fraudulent payee before the payload is signed.

A8 — Fraudulent or deficient receiving institution: An institution that obtains AFSP certification despite inadequate KYC/BSA-AML controls, or that is itself fraudulent, and uses legitimately-received PCAA packages to collect consumer data or applications under false pretenses. Goal: exploit certified-network trust to harvest consumer data or applications without genuine institutional legitimacy.

A9 — Physical device compromise: An attacker who has gained physical possession of a consumer's enrolled device (theft, loss, or unauthorized access) and operates it directly, potentially with location spoofing to defeat geographic plausibility checks. Distinct from A3 (no interception or replay involved — the attacker operates the device live) and A4 (no legitimate consumer is present to coerce). Goal: use the enrolled device's cryptographic identity and biometric access to initiate or execute unauthorized actions.

A10 — Prompt injection via CLO tool chain: An attacker who injects a malicious instruction into the CLO tool chain that redirects an authorized action to a fraudulent target while remaining within the authorized scope envelope. Signal 5 detects no anomaly because the action clears the Authorization Tier 2 scope check and the human is genuinely present. Distinct from A7, the instruction arrives through the CLO's own tools rather than through financial data. Goal: have a legitimate, present consumer authorize an action directed to a fraudulent target. Mitigation: Intent Verification Checkpoint (IVC), Section 5.8. The CLO presents a plain-language summary generated directly from the structured Servicing Intent Package — not from model output — and obtains explicit consumer confirmation before transmission. A redirected action shows the fraudulent payee; the consumer catches the mismatch before the payload is signed. Under IVC Requirement 4 (Tool Chain Disclosure), any tool chain call not in the certified build's manifest triggers automatic Authorization Tier 3 escalation. This assumes the IVC rendering module is uncompromised, which Signal 1 build certification is intended to ensure.

THREAT	ATTACK VECTOR	AFSP MITIGATION
Fraudulent agent (A1)	Unregistered agent presents itself as a certified build	VAL-05: agent provenance attestation signature verified against AFSP agent registry. VAL-06: agent build hash verified against registry record. Unregistered or modified agents fail both checks. Additionally, the Authorization Scope Envelope (Section 7.4) provides a second independent layer: a fraudulent agent cannot obtain a valid Authorization Tier 2 (INITIATE) or Authorization Tier 3 (EXECUTE) payload, because only the consumer's hardware-held tier keys can produce one, each higher-tier signature requires consumer user verification, and verifiers reject out-of-tier signatures (VAL-08a).

THREAT	ATTACK VECTOR	AFSP MITIGATION
Synthetic identity (A2)	Fabricated financial profile submitted as Signal 4	S2 and S4 together substantially raise the cost of a synthetic-identity attack: the attacker must control an enrolled device bound to a KYC record at a participating institution and sustain a multi-institution financial history consistent with the behavioral consistency markers in Section 8.4. Deliberately aged synthetic identities may still pass, so S2 and S4 supplement, rather than replace, the receiving institution's own synthetic-identity controls.
Package replay (A3)	Valid PCAA package captured and replayed to same or different institution	VAL-03: temporal validity window (max 3600 seconds). VAL-04: unique package_id consumed-package registry. VAL-02: receiving_institution_id binding — package invalid if presented to a different institution. Additionally, the Authorization Scope Envelope (Section 7.4) limits replay blast radius: a replayed package is constrained to the Authorization Tier of the captured credential and cannot be escalated.
Biometric coercion (A4)	Consumer forced to authenticate under duress	S5 may flag sessions whose behavior departs from normal patterns, including some automated or coerced sessions, and flagged sessions route to human review under the Non-Decline Principle. The IVC (Section 5.8) shows the consumer the exact action before any Authorization Tier 2 (INITIATE) or Authorization Tier 3 (EXECUTE) payload is signed. Neither control reliably detects a consumer operating their own device under duress; coercion remains a residual risk (Section 3.2).
Provider key compromise (A5)	Attacker compromises Signal 4 or Signal 5 provider signing key	Each provider's signing key is independently verified (VAL-09, VAL-10). Key rotation is required upon compromise. AFSP platform maintains a revocation list for compromised provider keys.
Platform key compromise (A6)	Attacker compromises AFSP platform signing key	VAL-01: all packages verified against platform public key. Key rotation procedure required. Receiving institutions subscribe to platform key update notifications. Prior packages with revoked keys are invalidated.
Prompt injection via financial data (A7)	Malicious instructions embedded in transaction descriptions or merchant names attempt to manipulate agent behavior	AFSP-02 Stage 2 normalization removes all free-text fields (merchant names, transaction descriptions) before signal derivation. Raw text never reaches the attestation layer. Additionally, the IVC confirmation at Authorization Tier 2 (INITIATE) and Authorization Tier 3 (EXECUTE) (Section 5.8, Section 7.4) reads from the structured Servicing Intent Package rather than model output — a redirected action shows the fraudulent payee at the IVC confirmation step.
Fraudulent or deficient institution (A8)	Institution obtains AFSP certification despite inadequate KYC/BSA-AML controls, or is itself fraudulent	Section 17.2 KYC Rigor Condition: certification excludes institutions subject to an active public enforcement action related to BSA/AML, CIP, CDD, or identity verification, supported by an annual officer attestation. Annual reverification by the AFSP Foundation.
Physical device compromise (A9)	Attacker gains physical possession of enrolled device and operates it directly, with location spoofing	Section 11.4 Environmental Consistency Layer: ambient device ecosystem cross-check independent of GPS. Geographic Risk Tier 5 suspends the session pending out-of-band reauthorization through the receiving institution.
Prompt injection via CLO tool chain (A10)	Malicious instruction injected through a CLO tool, plugin, or API call redirects an authorized action to a fraudulent target within the authorized scope	IVC (Section 5.8): the confirmation is generated from the structured Servicing Intent Package rather than model output, so a redirected action shows the fraudulent payee before the payload is signed. IVC Requirement 4: a tool chain call not in the certified manifest triggers automatic Authorization Tier 3 escalation. Signal 1 build certification covers the IVC rendering module.

3.2 Out of Scope

The following threats are explicitly out of scope for AFSP v0.1 and should be addressed by implementing organizations through complementary controls:

Device compromise: An attacker who has fully compromised the consumer's enrolled device before the PCAA session begins. AFSP assumes the TEE on the enrolled device is uncompromised. Device compromise is addressed through device management and TEE attestation infrastructure outside AFSP's scope.

Receiving institution KYC failure: AFSP delivers a pre-credential trust package. It does not govern the receiving institution's KYC process. Fraudulent applications that pass AFSP pre-qualification but fail to satisfy KYC are handled by the institution's existing fraud and KYC controls.

Post-account-opening fraud: AFSP governs the pre-credential session. Account takeover, authorized push payment fraud, and other post-opening fraud categories are outside AFSP's scope and are governed by existing payment network and bank fraud controls.

Agent behavior after session establishment: AFSP governs the session gate. What the agent does within an authorized session is governed by the consumer's authorization scope (AFSP-01 action classes) and the receiving institution's own controls. AFSP v1.0's AFSP-09 and AFSP-10 components extend governance into the session but are not in scope for v0.1.

Coercion of a consumer operating their own device: AFSP controls may surface some coerced sessions but cannot reliably detect a consumer acting under duress. Receiving institutions should apply their existing scam and fraud controls.

3.3 Signal Availability and Partial Package Behavior

AFSP-governed sessions may encounter signal unavailability due to provider outages, network failures, or enrollment gaps. The following rules govern partial package behavior:

S1 (Agent Provenance) — MANDATORY: A PCAA package without a valid S1 attestation SHALL be rejected. An unverified agent is the foundational risk AFSP addresses. No partial session is permitted without S1.

S2 (Biometric Consent) — MANDATORY: A PCAA package without a valid S2 authorization token SHALL be rejected. Consumer authorization is non-negotiable for any agent-initiated session.

S3 (Federated Referral) — OPTIONAL: S3 may be absent when the consumer is not a returning network participant or when no federated referral attestation is available. Absent S3 is noted in the signals_present field. The receiving institution may apply a lower composite trust confidence score but SHALL NOT automatically reject the package.

S4 (Financial Identity) — MANDATORY: A PCAA package without a valid S4 attestation SHALL be rejected. Financial identity attestation is required for the fit-based marketplace to function and for the fraud protections built into S2+S4 to apply.

S5 (Session Integrity) — CONDITIONAL: If no certified S5 provider is available (e.g., provider outage), the package MAY be delivered with S5 absent, noted in the signals_present field. The receiving institution's configured risk-tiered decisioning framework governs whether to proceed, request step-up verification, or decline. AFSP platform operators SHOULD maintain at least two certified S5 providers to minimize single-provider outage risk.

> **Summary:** S1 and S2 are unconditionally mandatory — no session proceeds without them. S4 is mandatory for marketplace sessions. S3 and S5 are conditional — their absence is logged, disclosed to the receiving institution, and handled by the institution's risk framework rather than automatic rejection.

3.4 Cryptographic Assumptions

AFSP v0.1 security depends on the following cryptographic assumptions:

ECDSA-P256: All AFSP signatures use ECDSA with the P-256 curve. This is assumed to provide at least 128-bit security. Implementing organizations should monitor NIST guidance on elliptic curve security and plan for algorithm agility in future versions.

SHA-256: Package binding uses SHA-256 hashing. Assumed collision-resistant for the purposes of inquiry event binding and build hash verification.

TEE integrity: AFSP Signal 2 assumes the hardware TEE on the consumer's enrolled device is uncompromised and that TEE attestation from the device is authentic. Implementing organizations should verify TEE attestation chains against manufacturer root certificates.

Differential privacy: Signal 4 privacy protections depend on correct implementation of Laplace-mechanism differential privacy with epsilon=1.0 and k-anonymity with k=5. Implementing organizations should independently validate their differential privacy implementations against established test vectors. The privacy model for individual attestations is an item for further consideration (Section 18).

4. Extension Process and Registry

AFSP is designed to be extended through a defined process that prevents forking while enabling the protocol to evolve as the agentic banking channel matures. This section defines how new Signal types, action classes, schema fields, and protocol components may be proposed and adopted.

4.1 Extension Categories

Type 1 — Minor extension: Addition of optional fields to existing schemas, addition of values to existing enumerated types (e.g., new product categories in the reverse inquiry payload, new feature_set values), or clarification of existing normative language without changing behavior. Minor extensions require working group review and are incorporated into the next point release (e.g., v0.1 → v0.2).

Type 2 — Major extension: Addition of a new Signal type, a new AFSP component number, a new action class, or a change to mandatory behavior. Major extensions require working group consensus, a 60-day public comment period, and are incorporated into the next minor or major version (e.g., v0.1 → v1.0 or v1.0 → v2.0).

Type 3 — Profile extension: A named implementation profile (e.g., "AFSP-EU" for PSD3 markets, "AFSP-SMB" for small business agents) that constrains or extends the base specification for a specific deployment context. Profile extensions are published as companion documents and do not modify the base specification.

4.2 Proposal Process

Any organization may submit an extension proposal through the AFSP GitHub repository (https://github.com/primitive-os/AFSP) by opening an issue tagged "extension-proposal." Proposals SHALL include: (a) the problem statement and use case; (b) the proposed technical specification; (c) backward compatibility analysis; (d) security considerations specific to the extension; and (e) the proposing organization's identity and contact.

The AFSP Technical Working Group reviews extension proposals on a rolling basis. Founding endorsers have standing participation in the working group. Other organizations may participate as observers or contributors as defined in the working group charter.

4.3 Registered Values

The following values are registered under AFSP v0.1 and may be extended through the Type 1 process:

Action classes: READ, DISCOVER, PRESENT, INITIATE, EXECUTE

Product categories: SAVINGS, CHECKING, CD, MMA

Signal types (S5): KEYSTROKE_DYNAMICS, TOUCH_PATTERNS, DEVICE_MOTION, MOUSE_MOVEMENT, NAVIGATION_TIMING, DEVICE_HARDWARE

Clearance decisions: PROCEED_TO_KYC, REQUEST_ADDITIONAL_VERIFICATION, DECLINE_SESSION

Return trip outcomes: OPENED, DECLINED, PENDING, ABANDONED

5. Consumer Layer Operator Architecture and Discovery Routing

> **AFSP is AI assistant platform-agnostic.** The protocol governs the trust signals that accompany an agent session — not the platform, company, or product that operates the agent. Any AI assistant from any operator that achieves AFSP agent certification may initiate AFSP-governed sessions. No AI platform receives preferential treatment in the agent registry, the reverse inquiry marketplace, or the FitScore ranking algorithm.

5.1 Overview

AFSP governs two distinct phases of an agentic financial services session. The first phase — discovery — is how an AI agent learns that AFSP-governed product discovery is available and routes the consumer's financial intent to the AFSP marketplace. The second phase — authorization and attestation — is what the rest of this specification governs in detail.

Discovery is architecturally necessary. Without it, AFSP has no answer for how an agent finds the marketplace. A consumer telling their AI assistant "find me the best savings account" produces no AFSP session unless the agent knows that AFSP-governed discovery exists and where to reach it.

The discovery phase operates on public information. AI agents crawl and query publicly available financial institution content in the same way search engines index public web content. No consumer authentication, no data sharing, and no regulatory authorization is required for the discovery phase. Regulatory obligations — consumer consent, biometric authorization, data access rights — attach at the moment the consumer initiates an AFSP session in Phase 2, not at the moment of discovery.

> **Phase 1 — Discovery (this section):** Agent crawls public content, finds AFSP Discovery Declaration, routes consumer intent to AFSP marketplace. No authentication. No data sharing. No regulatory event. Phase 2 — Authorization and Attestation (Sections 6–12): Consumer authorizes session biometrically. PCAA package assembled. Reverse inquiry submitted. Clearance handshake. Account opening. All regulatory obligations attach here.

5.2 Consumer Layer Operator (CLO) Role

The Consumer Layer Operator (CLO) is the entity operating the interface through which the consumer authorizes an AFSP session and through which the agent acts. The CLO is responsible for: (a) implementing the consumer-facing biometric authorization ceremony defined in AFSP-01; (b) managing the hardware TEE enrollment that binds the consumer's biometric to their AFSP identity; (c) presenting PCAA-governed sessions to the consumer in a manner consistent with the consumer disclosure

requirements in Section 5.9; and (d) implementing AFSP Discovery Declaration publication or registration as required for its entry point type.

The CLO and the agent operator may be the same entity or distinct entities. A bank that builds its own AI assistant and deploys it within its own app is simultaneously the CLO and the agent operator. An AI platform company whose agent initiates a session on behalf of a consumer is the agent operator; the CLO may be the AI platform, the consumer's primary bank, or a certified third-party enrollment provider. AFSP does not require consumers to use any specific CLO or AI assistant platform.

5.3 Entry Point Types

AFSP recognizes four entry point types through which consumers may reach the AFSP marketplace. Each has distinct CLO responsibilities and discovery implementation requirements. The exact technical implementation of AFSP Discovery Declarations for each entry point type — including metadata formats, well-known URL schemas, and AI platform tool registration protocols — will be specified by the AFSP Technical Working Group and published as a companion specification. v0.1 establishes the architectural framework and certification requirements.

ENTRY POINT TYPE	EXAMPLES	CLO ROLE	CERTIFICATION PATH
Standalone AI Assistant	General-purpose AI assistants and AI platform products	AI platform is the CLO. Implements biometric authorization within the platform's interface. Registers with Primitive as a certified AFSP Consumer Layer Operator. Implements AFSP discovery routing in its financial services tool stack.	AI Platform CLO Agreement with Primitive. Tool registration in Primitive's CLO registry.
Bank-Embedded AI Assistant	Bank mobile app AI assistant, bank website conversational agent	Bank is simultaneously CLO and institution. Implements AFSP-01 biometric authorization in its existing mobile app or web interface. Consumer enrollment is the natural extension of existing bank onboarding.	Bank achieves Level 2/3 certification. CLO capability included in standard certification scope.
Embedded Widget or SDK	Third-party financial assistant widget embedded in non-financial app, personal finance aggregator	Widget operator is the CLO. Implements AFSP-01 via Primitive's certified CLO SDK. Consumer authorizes through the widget's interface. CLO SDK handles TEE enrollment and biometric binding.	CLO SDK certification. Widget operator signs CLO Agreement.
Platform Company (Infrastructure Layer)	Core banking platform providers, digital banking platform providers, account opening platform providers	Platform company acts as CLO on behalf of all client institutions. Single CLO certification covers all institutions on the platform. Implements AFSP Discovery Declarations across all client institution digital properties as part of platform rollout.	Platform CLO certification. Single agreement covers all client institutions. Discovery Declaration publication is a platform certification requirement.

5.4 AFSP Discovery Declaration — Architectural Framework

The AFSP Discovery Declaration is a standardized machine-readable signal published by participating institutions, platform companies, or CLOs on their public-facing digital properties. It tells any AFSP-capable AI agent that AFSP-governed product discovery is available and where to route the consumer's financial intent.

The Discovery Declaration operates on public information — it is a routing signal, not a consumer disclosure or data sharing authorization. No consumer authentication is required to read a Discovery Declaration. The Declaration simply signals availability; the consumer's authorization decision happens in Phase 2.

> **Note on regulatory review:** The exact format, required fields, and implementation requirements for AFSP Discovery Declarations — including website metadata tags, well-known URL endpoints, and AI platform tool registration protocols — are subject to review by the AFSP Technical Working Group and counsel before finalization. The v0.1 specification establishes the architectural requirement and the certification obligation. The companion specification governing technical implementation will be published prior to v1.0.

The Discovery Declaration framework SHALL support at minimum the following discovery pathways:

Website metadata: A machine-readable declaration published on the institution's or platform's public-facing website, discoverable by AI agents crawling public financial institution content. The specific format (HTML metadata, well-known URL endpoint, or equivalent) will be defined by the Technical Working Group.

AI platform tool registration: Primitive registers as a verified financial services tool provider with major AI platforms (including but not limited to major AI platform operators). When any registered AI agent encounters financial product discovery intent, the platform's tool routing directs it to Primitive's AFSP marketplace. AI platform tool registration is a CLO certification requirement for AI platform entry point types.

Platform-level declaration: A certified platform company (core banking platform providers, digital banking platform providers, or equivalent) publishes a single AFSP Discovery Declaration covering all client institutions on its platform. This declaration is updated automatically as institutions join or leave the platform, eliminating the need for individual institution-level declaration management.

5.5 Discovery Routing and the AFSP Marketplace

Regardless of entry point type, AFSP discovery across institutions routes to the AFSP marketplace, which Primitive operates as the network's initial marketplace. Certification of additional marketplace operators is an item for further consideration (Section 18).

Certification requirement: Any entity seeking AFSP CLO certification SHALL implement the AFSP Discovery Declaration appropriate for its entry point type and SHALL ensure that AFSP-capable agents operating through its interface route financial product discovery across institutions to the AFSP marketplace, except in institution-direct sessions as provided below.

Institution-direct sessions: The routing requirement above applies to discovery across institutions. A Bank-Embedded AI Assistant, or a Platform Company CLO operating within a client institution's branded experience, MAY limit its sessions to that institution's products and proceed directly to the clearance handshake (Section 10.2). Such a CLO SHALL disclose at the outset that it presents that institution's products only. If the consumer asks for other options, the CLO SHALL tell the consumer how to seek them through AFSP discovery.

A Standalone AI Assistant or an Embedded Widget or SDK CLO MAY proceed directly to a single institution only when the consumer has named that institution. If the consumer asks for broader options at any point, the CLO SHALL offer scope expansion under Section 9.2.

An institution's choice of participation mode SHALL NOT restrict a consumer's ability to seek offers from other institutions.

5.6 CLO Certification Requirements Summary

Full CLO certification requirements will be published in the AFSP CLO Certification Specification (companion document, v1.0). The following requirements apply to all CLO types in v0.1:

(a) Implement AFSP-01 consumer authorization and biometric binding as specified.

(b) Publish or implement an AFSP Discovery Declaration appropriate for the CLO's entry point type.

(c) Route financial product discovery across institutions to the AFSP marketplace, except in institution-direct sessions (Section 5.5).

(d) Execute a Consumer Layer Operator Agreement with Primitive covering data handling obligations, TEE implementation requirements, liability allocation, and any consumer disclosure obligations beyond those required by Section 5.9.

(e) For AI platform CLOs: register with Primitive as a certified AFSP tool provider and implement AFSP financial product discovery routing in the platform's financial services tool stack.

(f) For platform company CLOs: implement Discovery Declaration publication as a standard platform feature covering all client institutions and update the declaration as institutions join or leave the platform.

(g) Meet the consumer disclosure requirements in Section 5.9.

5.7 Post-Opening Session Continuity (AFSP-12 — v1.0 Roadmap)

AFSP v0.1 governs the pre-credential session preceding a new account opening. AFSP-12 (Post-Opening Session Continuity) extends AFSP governance to subsequent agent-mediated interactions with an account opened through an AFSP session. This section establishes the architectural concept. Full technical specification will be published as AFSP-12 in the v1.0 roadmap following the working group process defined in Section 4.

The Post-Opening Trust Anchor. At the conclusion of an AFSP-governed account opening session, three data points exist: (a) a verified biometric root key bound to the consumer’s enrolled device and cryptographically accepted by the receiving institution; (b) a confirmed account relationship evidenced by the Return Trip API confirmation; and (c) a behavioral baseline from Signal 4 and Signal 5 established during the opening session. Together these constitute an AFSP post-opening trust anchor for subsequent authenticated interactions through an AI agent.

The Three-Leg Triangulation. Re-entry agent sessions for AFSP-opened accounts are authenticated through three independently verifiable legs: Leg 1 (Biometric Continuity) — the re-entry session presents a Signal 2 token signed by the same TEE root key as the original opening; Leg 2 (Behavioral Continuity) — the re-entry Signal 5 score is evaluated against the opening baseline, with significant deviation triggering the non-decline principle; Leg 3 (Relationship Confirmation) — the AFSP platform’s continuity registry cross-references the re-entry session against the Return Trip API confirmation from the original opening. A re-entry session satisfying all three legs is authenticated through the AFSP post-opening trust anchor without requiring the consumer to re-establish identity through the institution’s legacy authentication stack.

Scope of Post-Opening Agent Authority. Re-entry sessions inherit action class constraints from the opening authorization token. Actions outside the consumer’s defined authority envelope require a fresh

biometric step-up authorization event. The agent presents. The human decides. The bank executes. CLOs implementing bank-embedded or SDK entry point types SHOULD implement AFSP-12 to support agent-mediated account servicing. Platform company CLOs SHOULD implement AFSP-12 as a standard platform feature. Full certification requirements will be defined in the v1.0 working group.

5.8 Intent Verification Checkpoint (IVC)

The Servicing Intent Package is the structured, machine-generated payload assembled by the CLO for any Authorization Tier 2 (INITIATE) or Authorization Tier 3 (EXECUTE) action — the specific action type, payee or recipient identifier, amount and currency, timing, and receiving institution, plus, for account opening, the product terms required by Section 5.9 (CD-5). It is distinct from the PCAA package (Section 6): the PCAA package establishes pre-credential trust before account opening; the Servicing Intent Package describes one specific action once a session is authorized. *(See Section 7.4 for the complete Authorization Tier definitions and cryptographic enforcement mechanism.)*

The Intent Verification Checkpoint (IVC) is a mandatory control applied to all Authorization Tier 2 (INITIATE) and Authorization Tier 3 (EXECUTE) actions before transmission (for action class definitions see Section 7.4) of a Servicing Intent Package to a receiving institution, regardless of whether Signal 5 has detected any session anomaly. It addresses the prompt injection attack vector described in Section 3.1 (A10): a malicious instruction that redirects a legitimate authorized action to a fraudulent target while remaining within the consumer's authorized scope envelope, where no behavioral anomaly exists for Signal 5 to detect.

IVC Requirement 1 — Plain-Language Confirmation.

Before transmitting any Authorization Tier 2 (INITIATE) or Authorization Tier 3 (EXECUTE) Servicing Intent Package, the certified CLO MUST present to the consumer a plain-language summary including at minimum: action type, payee or recipient name and account identifier, amount and currency, timing, and receiving institution name. The consumer MUST provide explicit confirmation before transmission.

IVC Requirement 2 — Confirmation Integrity.

The summary MUST be generated directly from the structured fields of the Servicing Intent Package — not from CLO language model output. This prevents second-order prompt injection where the model manipulates both the action and the confirmation description. The IVC rendering module MUST be included in the agent's build hash scope for Signal 1 certification.

IVC Requirement 3 — Audit Trail Logging.

The confirmation event MUST be logged in the session audit trail, including: exact summary text presented, presentation timestamp, confirmation timestamp, and authentication method used. This log is part of the tamper-evident session record.

IVC Requirement 4 — Tool Chain Disclosure.

The CLO certified build MUST disclose all external API calls and plugins in a manifest. Any tool chain call not in the certified manifest triggers automatic Authorization Tier 3 escalation. Undisclosed calls cannot be resolved within Authorization Tier 2 regardless of scope.

IVC and HITL Relationship.

HITL is triggered by Signal 5 anomalies detected at session initiation. IVC is a mandatory pre-transmission step for every Authorization Tier 2 or Authorization Tier 3 action regardless of session-initiation integrity status — action content fidelity. Both MUST be satisfied for an Authorization Tier 2 or Authorization Tier 3 action to proceed.

5.9 Consumer Disclosure Requirements

Certified CLOs SHALL provide the following disclosures to the consumer in plain language, in the language and accessibility mode of the CLO's interface, at the points indicated. These requirements set a minimum. They do not replace disclosures that receiving institutions must provide under applicable law, and a Consumer Layer Operator Agreement may add to them but may not reduce them.

CD-1 — Before Authorization.

Before the consumer authorizes a session under AFSP-01, the CLO SHALL disclose the identity of the CLO and of the agent operator; that the session uses AFSP and the AFSP platform; which financial data the Signal 4 provider will access, from which accounts, and for what purpose; that raw data is processed transiently and not retained (Section 8.2); and that declining does not affect the consumer's existing accounts.

CD-2 — Inquiry Scope.

Before a reverse inquiry is submitted, the CLO SHALL disclose which institutions will receive the anonymous inquiry under the active participation state (Section 9.2), and that the consumer's identity is not shared with any institution until the consumer selects an offer and that institution returns PROCEED_TO_KYC.

CD-3 — Ranking and Compensation.

When presenting offers, the CLO SHALL disclose that offers are ranked by the published FitScore formula (Section 9.5) using the consumer's preference weights, or the default weights if the consumer has not set them; that no institution can pay for ranking position; that the AFSP platform receives fees from participating institutions (Section 17.2); and that offers come only from participating institutions, so products from other institutions are not shown.

CD-4 — Offer Terms.

Each presented offer SHALL identify the offering institution's legal name and deposit insurer (FDIC or NCUA) and SHALL state the terms required by applicable deposit advertising rules (Appendix B), including, where applicable: the annual percentage yield and the date it is accurate as of; whether the rate may change; minimum opening deposit and balance requirements; fees that could reduce earnings, with any waiver conditions; and, for certificates of deposit, the term and early withdrawal penalty.

CD-5 — Account Opening Confirmation.

For an INITIATE action that opens an account, the IVC summary (Section 5.8) SHALL include, in addition to the elements in IVC Requirement 1, the product name, annual percentage yield, recurring fees, minimum balance, and deposit insurer. The receiving institution remains responsible for delivering its account agreement and legally required account-opening disclosures through institution-controlled channels before the account is opened.

CD-6 — Outcomes.

If the receiving institution returns DECLINE_SESSION or REQUEST_ADDITIONAL_VERIFICATION, the CLO SHALL tell the consumer promptly and plainly, identify the receiving institution as the party that made the decision, and provide the institution's contact channel. The CLO SHALL NOT describe the outcome as a decision of AFSP or the AFSP platform. Any notice that applicable law requires of the institution remains the institution's responsibility.

CD-7 — Data Use.

The CLO SHALL disclose that account-opening outcomes are reported to the AFSP platform through the Return Trip API (Section 13) and used for fee settlement, FitScore calibration, and aggregate network analytics that contain no individual consumer data and SHALL provide a link to the AFSP platform's privacy notice.

CD-8 — Revocation and Records.

The CLO SHALL tell the consumer how to revoke AFSP enrollment and any data-access consent under the CLO's implementation, and how to request a record of the authorizations the consumer has given, which the CLO SHALL provide from the session audit trail (IVC Requirement 3).

6. The AFSP Attestation Package

6.1 Package Structure

An AFSP attestation package (PCAA package) is a signed JSON object assembled by the AFSP platform and delivered to a receiving institution's registered API endpoint before any open banking API is invoked. The package contains five independently sourced trust signals.

```
{
  "afsp_version": "0.1",
  "package_id": "pkg_[256-bit unique identifier]",
  "session_id": "sess_[session correlation identifier]",
  "assembly_timestamp": "[ISO 8601 UTC timestamp]",
  "receiving_institution_id": "[institution identifier from participation registry]",
  "institution_nonce": "[institution-bound session nonce, added at binding]",
  "nonce_bound_institution": "[hash of the receiving institution's registered public key]",
  "nonce_expiry": "[ISO 8601 UTC timestamp, max 1800 seconds after assembly_timestamp]",
  "validity_window": {
    "not_before": "[ISO 8601 UTC timestamp]",
    "not_after": "[ISO 8601 UTC timestamp, max 3600 seconds after not_before]"
  },
  "signals": {
    "S1": { [AFSP-01 Agent Provenance Attestation] },
    "S2": { [AFSP-01 Consumer Authorization Token] },
    "S3": { [AFSP-03 Federated Referral Attestation — optional] },
    "S4": { [AFSP-02 Financial Profile Attestation] },
    "S5": { [AFSP-05 Behavioral Session Intelligence Score] }
  },
  "signals_present": ["S1", "S2", "S4", "S5"],
  "platform_signature": {
    "algorithm": "ECDSA-P256",
    "signed_by": "platform.afsp.org",
    "signature": "[ECDSA-P256 signature over canonical package JSON]",
    "certificate_chain": "[X.509 certificate chain]"
  }
}
```

Package binding and replay prevention (institution-specific nonce binding): Every PCAA package must contain an institution-specific session nonce that cryptographically binds the package to a single receiving institution and a single session. The nonce is generated by the AFSP platform using the receiving institution's registered public key from the Foundation's Participation Registry. A package with a nonce bound to Institution A is cryptographically unusable at Institution B — nonce verification fails regardless of the package's signature validity and timestamp freshness. Nonce generation and binding is performed within the AFSP platform's TEE.

Signal presence: Signal S3 (Federated Referral) is optional in AFSP v0.1. When absent, the package is valid but carries a lower composite trust confidence. Signal S3 availability is dependent on network adoption and will become more prevalent as the certified institution network grows.

Anti-replay and nonce expiry: Each package contains a unique 256-bit package_id and an institution-bound nonce. The nonce expires at the earlier of: (a) successful processing by the bound receiving institution; (b) thirty (30) minutes from package assembly timestamp; or (c) explicit revocation by the AFSP platform. A package's effective validity ends at the earlier of its validity_window not_after time and its nonce_expiry. The institution nonce is added when the package is bound to the selected institution (Section 10.2); a package held by the AFSP platform before selection carries no institution nonce. Receiving institutions SHALL maintain a consumed-package registry and SHALL implement nonce validation as the mandatory first step in PCAA package processing, rejecting and logging any package that fails nonce validation or presents a previously-processed package_id, regardless of other package validity indicators. Processing packages without nonce validation is a breach of NPA certification requirements.

6.2 Package Validation Requirements

Receiving institutions that elect to implement AFSP SHALL perform the following validation sequence upon receipt of a PCAA package:

REQ	SIGNAL	REQUIREMENT	RATIONALE
VAL-01	Package	Verify ECDSA-P256 platform signature against published AFSP platform public key	Confirms package was assembled by the certified AFSP platform operator and has not been tampered with in transit
VAL-02	Package	Confirm receiving_institution_id matches the institution's own registered identifier	Prevents package redirection attacks where a valid package is presented to an unintended institution
VAL-03	Package	Confirm current time falls within validity_window (not_before, not_after)	Enforces temporal validity — stale packages are rejected
VAL-04	Package	Verify package_id has not been previously consumed	Prevents replay attacks using previously valid packages
VAL-04a	Package	Verify institution_nonce cryptographically matches nonce_bound_institution for this receiving institution and current time is before nonce_expiry	Prevents package redirection and replay via the institution-specific nonce (Section 6.1) — independent of the VAL-02/VAL-04 checks
VAL-05	S1	Verify agent provenance attestation signature against AFSP agent registry public key	Confirms the AI agent is currently certified and unrevoked
VAL-06	S1	Confirm agent build hash matches registry record for the agent_id	Ensures the presenting agent is the certified build, not a modified version
VAL-07	S2	Verify consumer authorization token ECDSA signature against enrolled device public key	Confirms biometric binding event occurred on the consumer's enrolled device
VAL-08	S2	Confirm authorization scope permits the requested action class	Enforces consumer-defined authorization boundaries

REQ	SIGNAL	REQUIREMENT	RATIONALE
VAL-08a	S2	Confirm the EKU of the signing key's certificate permits the payload's action class	Enforces Authorization Tier separation (Section 7.4)
VAL-09	S4	Verify Signal 4 attestation ECDSA signature against certified provider public key	Confirms financial profile attestation was generated by a certified AFSP Signal 4 provider
VAL-10	S5	Verify Signal 5 composite score ECDSA signature against certified provider public key	Confirms behavioral session score was generated by a certified AFSP Signal 5 provider

> **Failure handling:** Any validation failure in VAL-01 through VAL-04, or VAL-04a, (package-level) SHALL result in immediate session termination. Failures of VAL-05 through VAL-09, including VAL-08a, which verify the mandatory signals S1, S2, and S4, SHALL result in rejection of the package, consistent with Section 3.3. A VAL-10 failure (Signal 5) SHALL be logged and reported in the institution's configured response and MAY allow conditional session continuation at the institution's discretion under its risk-tiered decisioning framework, consistent with the conditional treatment of Signal 5 in Section 3.3.

6.3 Transport Security (Interim)

Until the protocol binding is specified (Section 18), all AFSP messages SHALL be exchanged over TLS 1.2 or later, and implementations SHOULD use TLS 1.3. Connections between the AFSP platform and receiving institutions SHALL be mutually authenticated. The message-level signatures defined in this specification apply in addition to, not instead of, transport security.

7. AFSP-01: Consumer Authorization and Biometric Binding

7.1 Overview

AFSP-01 governs the consumer-side authorization ceremony that initiates an AFSP-governed agent session. It answers Signal S1 (Who is this agent?) and Signal S2 (On whose behalf — and did they authorize it?).

The core architectural principle: the consumer's biometric credential never leaves their enrolled device. All cryptographic operations involving the biometric are performed within a hardware-secured trusted execution environment (TEE) on the consumer's personal computing device. No biometric data is transmitted to the AFSP platform, the receiving institution, or any other party.

7.2 Biometric Root Key Architecture

Upon consumer enrollment, the following key generation sequence SHALL occur within the device TEE:

1. The TEE generates an asymmetric root key pair (ECDSA-P256 or platform-equivalent hardware key).

2. The TEE generates a hardware attestation certifying that the key pair was created within the TEE. This attestation is made available to relying parties to verify the hardware-bound nature of the root key.

3. The root private key SHALL remain within the TEE at all times. Export of the root private key in any form is prohibited.

4. The root public key is exported to the consumer's issuing institution and bound to the consumer's existing KYC record at that institution.

5. The consumer's enrollment is confirmed by the issuing institution upon successful binding of the root public key to the KYC record. The enrolled device and the confirmed enrollment record together constitute the consumer's AFSP biometric root identity. No AFSP session may be initiated from an unenrolled device.

6. All subsequent authorization tokens issued in AFSP sessions are signed using the TEE-held root private key, making them cryptographically enforceable rather than merely asserted.

Enrollment portability: An issuing institution SHALL NOT condition a consumer's AFSP enrollment on exclusive use of its own channels, or restrict use of that enrollment with other certified CLOs.

7.3 Authorization Token Structure

An AFSP consumer authorization token is generated within the TEE upon biometric authentication and encodes the following fields:

```
{
  "token_type": "AFSP_CONSUMER_AUTH",
  "version": "0.1",
  "consumer_device_id": "[enrolled device identifier]",
  "biometric_event": {
    "type": "FIDO2_BIOMETRIC",
    "timestamp": "[ISO 8601 UTC timestamp of biometric event]",
    "tee_attestation": "[hardware attestation from device TEE]"
  },
  "authorized_agent_id": "[agent identifier from AFSP agent registry]",
  "receiving_institution_id": "[institution identifier — bound element]",
  "scope": {
    "action_classes": ["READ", "DISCOVER"],
    "product_categories": ["SAVINGS", "CHECKING", "CD", "MMA"],
    "geographic_constraints": "[optional — state/region codes]",
    "consumer_preference_weights": {
      "w_rate": 0.35, "w_fee": 0.25,
      "w_balance": 0.20, "w_features": 0.20
    }
  },
  "temporal_validity": {
    "not_before": "[ISO 8601 UTC timestamp]",
    "not_after": "[ISO 8601 UTC timestamp, max 3600 seconds]"
  },
  "scope_expansion": {
    "default_state": "EXISTING_RELATIONSHIPS_ONLY",
    "expansion_requires": "BIOMETRIC_RECONFIRMATION",
    "expansion_token_ttl_seconds": 3600
  },
  "signature": "[ECDSA-P256 signature over canonical token JSON using TEE root private key]"
}
```

7.4 Action Class Definitions

The following action class taxonomy governs what an AI agent may do within an authorized session. Each higher class implies all lower classes plus additional consent:

ACTION CLASS	TIER	PERMITTED ACTIVITIES	STEP-UP REQUIRED	ADVERSARY MITIGATED	EXAMPLE
READ	1	Access and display account information, balances, and transaction history within consumer-authorized scope. No financial action. No consumer identity transmitted.	None	A1 (credential constraint)	Agent retrieves balance summary for consumer review.

ACTION CLASS	TIER	PERMITTED ACTIVITIES	STEP-UP REQUIRED	ADVERSARY MITIGATED	EXAMPLE
DISCOVER	1	Submit anonymous reverse inquiry to AFSP marketplace. Receive and display ranked product offers. No identity disclosure at this stage.	None	A1 (credential constraint)	Agent submits anonymous profile to marketplace and presents ranked savings account offers.
PRESENT	1	Present ranked marketplace results with fit rationale and comparative product information to the consumer for their independent evaluation and selection. The CLO presents offers ranked by the published FitScore formula; the consumer makes the selection.	None	A1, A7 (credential constraint + IVC scope)	Agent presents: “Based on your profile, these accounts ranked by fit” with comparative rate and fee information.
INITIATE	2	Begin account opening process at selected institution. Transmit PCAA package to receiving institution. Consumer identity is disclosed to the selected institution only after it returns PROCEED_TO_KYC (Section 10.2). Subject to mandatory Intent Verification Checkpoint (IVC) before payload is signed.	Biometric reconfirmation + IVC confirmation	A1, A3, A4, A7	Consumer selects product; agent presents IVC confirmation; consumer confirms; agent transmits PCAA package.
EXECUTE	3	Complete account opening. Authorize initial funding transfers within consumer-confirmed parameters. Only available after receiving institution returns PROCEED_TO_KYC clearance. Subject to mandatory IVC confirmation including specific funding parameters before payload is signed.	Biometric reconfirmation + IVC confirmation (including funding parameters) + institution PROCEED_TO_KYC	A1, A3, A4, A7 (all envelope layers)	Agent completes application after institution clearance; consumer confirms activation and funding parameters through IVC.

Authorization Tiers are calibrated to the risk level of the action classes they govern — higher tiers involve higher-consequence actions, require stronger consumer confirmation controls, and apply a narrower Authorization Scope Envelope. The tiers define what controls apply to each action class; they are not risk ratings AFSP assigns to sessions or consumers.

Action Class to Authorization Tier Mapping.

TIER	ACTION CLASSES	RISK LEVEL	CONSUMER STATE	IDENTITY DISCLOSED?
Authorization Tier 1 — Discovery and Advisory	READ, DISCOVER, PRESENT	Lowest — information only	Consumer has authorized session scope. No financial action yet.	No — anonymous profile only
Authorization Tier 2 — Transactional Initiation	INITIATE	Medium — account relationship begins	Consumer has biometrically reconfirmed and completed IVC for this specific action.	Yes — to the selected institution only, after PROCEED_TO_KYC
Authorization Tier 3 — Execution and Authorization	EXECUTE	Highest — full financial commitment	Consumer has biometrically reconfirmed, completed IVC with funding parameters, and institution has returned PROCEED_TO_KYC clearance.	Yes — full disclosure to institution

Cryptographic Enforcement — How Tier Separation Is Enforced.

Tier separation is enforced by three mechanisms working together.

Key separation. At enrollment, the device's secure hardware generates a distinct key pair for each Authorization Tier, bound to the consumer's enrolled root identity (Section 7.2). The private keys never leave that hardware. Authorization Tier 2 and Authorization Tier 3 keys require consumer user verification for every use; the Authorization Tier 1 key requires it at session authorization.

Certificate policy. Each tier key is certified with a tier-specific Extended Key Usage (EKU) value: the Authorization Tier 1 (Discovery) EKU permits READ, DISCOVER, and PRESENT payloads; the Authorization Tier 2 (INITIATE) EKU permits INITIATE payloads; and the Authorization Tier 3 (EXECUTE) EKU permits EXECUTE payloads. Because the EKU is part of the signed certificate, it cannot be altered without invalidating the certificate.

Verifier enforcement. The AFSP platform and receiving institutions SHALL reject any payload whose action class is not permitted by the EKU of the signing key's certificate (VAL-08a).

An Authorization Tier 1 signature over an INITIATE payload therefore fails validation, as does an Authorization Tier 1 or Authorization Tier 2 signature over an EXECUTE payload. A compromised or fraudulent agent cannot produce a valid higher-tier payload unless the consumer completes the step-up and IVC for that action (see adversary class A1, Section 3.1). These guarantees depend on every verifier performing the EKU check and on the integrity of the device's secure hardware (Section 3.4).

The full cryptographic specification for tier key generation, the EKU value registry, and the certificate chain structure is contained in AFSP-07 (Agent Provenance and Certification Registry, v1.0 Roadmap). AFSP v0.1 implementations SHALL implement Authorization Tier 1 credential enforcement as specified in AFSP-01. Full three-tier credential hierarchy enforcement is a MUST requirement for AFSP-07 certification. Until AFSP-07 is ratified, implementations SHALL use the interim EKU OID guidance published by the AFSP Technical Working Group at https://github.com/primitive-os/AFSP. Implementers SHOULD NOT register custom EKU OIDs for AFSP tier key designation without reference to the interim guidance — non-standard OIDs will produce implementations incompatible with other certified CLOs.

Authorization Scope Envelope.

The Authorization Scope Envelope is the complete set of constraints that bind what an agent may do within a session at a given tier. Every AFSP-governed agent action must satisfy all three envelope layers simultaneously before any payload is signed and transmitted. The three layers nest inside each other — the outermost is immovable regardless of instruction, the innermost is the most specific.

Layer 1 — Credential Constraint (outermost).
The EKU designation on each tier key's certificate limits which payload types verifiers will accept from that key. The constraint rests on hardware-held tier keys and verifier-enforced certificate policy: an agent or CLO cannot extract a tier key or alter its EKU, and verifiers reject out-of-tier signatures (VAL-08a). Addresses adversary classes A1 (a fraudulent agent cannot construct a valid higher-tier payload) and A3 (a replayed package is limited to the tier of the captured credential).

Layer 2 — Consumer-Defined Scope (middle).

The parameters the consumer authorized at session initiation: which institutions, which product categories, which action classes are active in this session, and which amount parameters apply, encoded in the AFSP consumer authorization token (Section 7.3). The consumer-defined scope can be narrower than what the credential key permits — a consumer may authorize Authorization Tier 2 capability but restrict it to a specific institution or product category — but cannot be broader than what the credential permits. Addresses adversary class A4 (a coercion attacker cannot expand scope beyond what the consumer authorized at session initiation).

Layer 3 — Intent Verification Checkpoint (IVC) Confirmation (innermost).

For Authorization Tier 2 (INITIATE) and Authorization Tier 3 (EXECUTE) actions: the specific action the consumer confirmed at the Intent Verification Checkpoint (Section 5.8) before the payload was signed. The IVC confirmation is the innermost envelope — it pins not just the class of action but the specific instance: the exact payee, amount, account identifier, and timing. Addresses adversary classes A4 (it helps a consumer who is being deceived notice a mismatch, though it does not stop a coercer who dictates the confirmation) and A7 (prompt injection that manipulates the agent's context is visible to the consumer at the IVC confirmation step, which reads from the structured Servicing Intent Package, not from model output).

The three layers operate independently — defeating one layer does not defeat the others. A fraudulent agent (A1) that somehow bypasses Layer 1 still faces the consumer-defined scope in Layer 2. A consumer being deceived about the action (A4) still sees the specific action in plain language at the Layer 3 IVC confirmation, generated from the structured package. This layered independence is what makes the Authorization Scope Envelope more robust than a single-layer policy enforcement approach.

Section 7.4 answers a different question than AFSP-10: which actions is this credential permitted to sign, not whether the human operating it is still present. AFSP-10 governs continuous presence verification and operates independently of tier scope.

TIER	CREDENTIAL CONSTRAINT	CONSUMER-DEFINED SCOPE	IVC REQUIRED?	EFFECTIVE ENVELOPE
Authorization Tier 1 (READ, DISCOVER, PRESENT)	EKU permits only Authorization Tier 1 payloads. INITIATE or EXECUTE signatures fail validation.	Which institutions visible, which product categories active, session duration.	No — information only. No financial action to confirm.	Observe and present within authorized categories. No financial action can be validly signed with an Authorization Tier 1 key.
Authorization Tier 2 (INITIATE)	EKU permits only INITIATE payloads. EXECUTE signatures fail validation.	Which specific institution selected, which product, biometric reconfirmation event.	Yes — specific payee, amount, account identifier, timing confirmed before signing.	Act precisely within what the consumer just confirmed at IVC. Payload scope limited to confirmed action instance.
Authorization Tier 3 (EXECUTE)	EKU permits only EXECUTE payloads.	Institution PROCEED_TO_KYC clearance received, consumer confirmed activation.	Yes — funding parameters, amount ceiling, account identifiers confirmed before signing.	Complete exactly what the institution cleared and the consumer confirmed. Payloads that exceed the confirmed parameters fail validation.

8. AFSP-02: Financial Profile Attestation

8.1 Overview

AFSP-02 defines the privacy-preserving behavioral attestation (Signal S4) that establishes the consumer's financial identity without transmitting raw account data or transaction records. The attestation is derived from consumer-permissioned financial data accessed under Section 1033 of the Dodd-Frank Wall Street Reform and Consumer Protection Act.

> **Data minimization principle:** Raw account balances, transaction records, account numbers, and personally identifiable financial records SHALL NOT be transmitted as part of an AFSP-02 attestation. Only derived behavioral signals as specified in Section 8.3 may be included.

8.2 Six-Stage Privacy-Preserving Pipeline

AFSP-02 attestation is generated through the following six-stage pipeline executed within a certified AFSP Signal 4 provider's infrastructure:

Stage 1 — Raw data ingestion: Consumer-permissioned multi-institution financial data is received through a Section 1033-compliant API connection. Raw data is ingested into ephemeral TEE memory. At no point does raw data leave the certified provider's controlled environment.

Stage 2 — Normalization: Raw financial data is normalized to a canonical internal format. Merchant names, transaction descriptions, and other free-text fields that could enable prompt injection attacks are removed at this stage.

Security note (A7 — Prompt Injection Defense): The Stage 2 normalization step is the primary defense against adversary class A7 (Section 3.1). By stripping all free-text fields — merchant names, transaction descriptions, payee references, and memo fields — before any signal derivation occurs, Stage 2 ensures that instructions embedded in financial data cannot reach or influence the S4 attestation. Stage 2 does not govern what the agent reads under the READ action class (Section 3.1, A7). No signal produced by AFSP-02 carries any free-text content from the consumer’s transaction history.

Stage 3 — Signal derivation: Six behavioral attestation signals are derived from the normalized data (see Section 8.3).

Stage 4 — Differential privacy application: Laplace-mechanism noise is applied to continuous signals with epsilon=1.0 as the preferred embodiment. K-anonymity enforcement at k=5 is applied to categorical signals, with generalization fallback if threshold is unmet. The attestation contains only coarse derived signals; individual transaction records cannot be reconstructed from it.

Stage 5 — Cryptographic binding: The attestation is bound to the specific inquiry event via SHA-256 hash of the inquiry event record, inquiry parameters, and session timestamp. A Merkle-tree commitment is made to the authorized recipient institution set.

Stage 6 — TEE zeroization: All raw data and intermediate processing results are cryptographically zeroized within 30 seconds of attestation generation. No raw data or intermediate results are retained.

8.3 Attestation Signal Definitions

The following six behavioral signals constitute an AFSP-02 attestation:

```
{
  "attestation_provider": "[certified AFSP Signal 4 provider identifier]",
  "attestation_version": "0.1",
  "inquiry_event_hash": "[SHA-256 hash of inquiry event record]",
  "signals": {
    "account_longevity": {
      "value": "[THIN | ESTABLISHING | ESTABLISHED | MATURE]",
      "confidence": "[0.0 to 1.0]"
    },
    "income_stability": {
      "value": "[IRREGULAR | VARIABLE | CONSISTENT | HIGHLY_CONSISTENT]",
      "frequency_component": "[0.0 to 1.0]",
      "regularity_component": "[0.0 to 1.0]",
      "magnitude_component": "[0.0 to 1.0]",
      "confidence": "[0.0 to 1.0]"
    },
    "balance_range": {
      "value": "[LOW | MID_LOW | MID | MID_HIGH | HIGH]",
      "confidence": "[0.0 to 1.0]"
    },
    "transaction_regularity": {
      "value": "[IRREGULAR | MODERATE | REGULAR | HIGHLY_REGULAR]",
      "behavioral_category": "[TRANSACTOR | REVOLVER | SAVER | MIXED]",
      "confidence": "[0.0 to 1.0]"
    },
    "relationship_breadth": {
      "value": "[SINGLE | LIMITED | MODERATE | BROAD]",
      "confidence": "[0.0 to 1.0]"
    },
    "product_eligibility": {
      "categories_indicated": ["SAVINGS", "CHECKING"],
      "minimum_balance_compatible": "[true | false]",
      "confidence": "[0.0 to 1.0]"
    }
  },
  "behavioral_consistency_score": {
    "score": "[0.0 to 1.0, aggregated per Section 8.4]",
    "markers_evaluated": "[integer]"
  },
  "privacy_parameters": {
    "method": "DIFFERENTIAL_PRIVACY_LAPLACE",
    "epsilon": 1.0,
    "k_anonymity_k": 5
  },
  "anti_replay": {
    "nonce": "[256-bit cryptographic nonce]",
    "sequence_number": "[monotonically increasing per-consumer sequence number]"
  },
  "temporal_validity": {
    "not_before": "[ISO 8601 UTC]",
    "not_after": "[ISO 8601 UTC, max 3600 seconds]"
  },
  "signature": {
    "algorithm": "ECDSA-P256",
    "certificate_chain": "[X.509 certificate chain for provider signing key]"
  }
}
```

8.4 Behavioral Consistency Evaluation Requirements

The Signal 4 differential privacy pipeline (Section 8.2) must evaluate the following behavioral consistency markers. These markers detect AI-generated synthetic financial profiles, which can pass point-in-time identity checks but cannot replicate the temporal consistency of genuine long-term financial behavior. Markers produce scored signals, not binary pass/fail tests. The pipeline aggregates scores into a Behavioral Consistency Score component of the attestation. The receiving institution applies its own thresholds — AFSP does not make a synthetic/genuine determination.

Marker 1 — Temporal Consistency: Life Stage to History Depth.

The pipeline must evaluate whether claimed or implied life stage is consistent with cross-institutional financial history depth. Flags include: consumer aged 45 or older with no financial account history predating 24 months; claimed homeowner with no mortgage, property tax, or homeowner's insurance premium history; consumer aged 60 or older with no retirement account or long-term savings history. [Fair-lending note: age-correlated flags require review against ECOA/Reg B disparate-impact considerations before implementation.]

Marker 2 — Cross-Institutional Account Opening Velocity.

Synthetic identities deploy in narrow fraud windows. Flags include: account opening events clustered within a 12-month window with no prior institutional relationships; more than 60% of total account relationships opened within the preceding 12 months; account closures in close temporal proximity to openings indicating cycling rather than sustained relationships.

Marker 3 — Income-to-Behavior Consistency.

Flags include: absent or irregular income deposit history for a consumer claiming stable employment; spending pattern distribution inconsistent with stated income level; consumer whose income history implies savings capacity but who has no observable savings account activity over a period consistent with their stated age.

Marker 4 — Geographic Behavioral Coherence.

Flags include: persistent transaction activity in locations geographically inconsistent with stated address without explanation; institutional relationships geographically scattered without pattern — real consumers typically establish banking relationships near home, workplace, or school.

Marker 5 — Nested Financial Relationship Depth.

Real financial lives exhibit nested relationships — accounts that fund other accounts, recurring transfers, insurance and investment relationships developing over time. Flags include: multiple accounts with no observable recurring transfers between them; absence of automated payment relationships (recurring bills, subscriptions, insurance premiums); absence of any financial relationship expected given implied life stage and profile.

Attestation Output.

Behavioral consistency evaluation produces a Behavioral Consistency Score component in the Signal 4 attestation. This score is not a fraud determination. It is processed within the TEE and zeroed after attestation assembly. No raw marker data or individual flag values are transmitted — only the aggregated score.

9. AFSP-03: Reverse Inquiry Protocol and Fit-Based Marketplace

9.1 Overview

AFSP-03 defines the reverse inquiry protocol through which a consumer's AI agent submits an anonymous structured financial product inquiry to a plurality of participating institutions simultaneously, receives ranked product offers, and presents them to the consumer for selection — without disclosing the consumer's identity to any institution prior to the consumer's explicit selection.

> **Consumer identity protection:** Under no circumstances SHALL the consumer's legal name, Social Security Number, date of birth, address, or any personally identifying information be transmitted as part of the reverse inquiry payload. Identity disclosure occurs only after the consumer explicitly selects a specific product offer and explicitly initiates the account opening process with the selected institution.

9.2 Tiered Participation States

Every reverse inquiry operates within one of three consumer-controlled participation states:

DEFAULT SCOPE STATE: The structured inquiry is transmitted exclusively to institutions satisfying both conditions: (a) the institution is an active network participant in the participation registry; and (b) a verified existing relationship between the consumer and the institution has been confirmed through cross-reference of the consumer's Section 1033-permissioned financial data against participation registry records. This is the default state absent explicit consumer scope expansion authorization.

EXPANDED SCOPE STATE: Activated exclusively by a biometric-authenticated consumer action producing a signed scope expansion token. In Expanded Scope, the structured inquiry is transmitted to all current active network participants. The scope expansion token is valid exclusively for the duration of the current inquiry session and does not constitute a standing authorization for subsequent sessions, each of which defaults to the Default Scope State.

NON-PARTICIPANT DISCLOSURE STATE: When a consumer inquires about a specifically named institution that is not a current active network participant, a real-time factual disclosure is generated comprising: (a) a statement that the named institution is not currently a participant in the inquiry network; and (b) three consumer routing options (see Section 9.6). Stated-preference data is logged.

9.3 Reverse Inquiry Payload Schema

The structured financial product inquiry transmitted to participating institutions:

```
{
  "inquiry_id": "[256-bit unique inquiry identifier]",
  "inquiry_timestamp": "[ISO 8601 UTC]",
  "authorization_token_hash": "[SHA-256 of consumer authorization token]",
  "participation_state": "[DEFAULT_SCOPE | EXPANDED_SCOPE]",
  "product_request": {
    "categories": ["SAVINGS", "CHECKING", "CD", "MMA"],
    "geographic_constraints": "[optional]"
  },
  "consumer_profile_attestation": {
    "[AFSP-02 financial profile attestation object]"
  },
  "consumer_preference_weights": {
    "w_rate": 0.35,
    "w_fee": 0.25,
    "w_balance_compatibility": 0.20,
    "w_features": 0.20
  },
  "consumer_identity_included": false
}
```

9.4 Institution Response Schema

Each participating institution's response to a reverse inquiry SHALL conform to the following schema:

```
{
  "inquiry_id": "[matching inquiry identifier]",
  "institution_id": "[institution's registered identifier in participation registry]",
  "response_timestamp": "[ISO 8601 UTC]",
  "product_offers": [
    {
      "offer_id": "[unique offer identifier]",
      "institution_legal_name": "[legal name of the offering institution]",
      "deposit_insurer": "[FDIC | NCUA]",
      "product_category": "[SAVINGS | CHECKING | CD | MMA]",
      "annual_percentage_yield": "[decimal — deposit products]",
      "apy_accurate_as_of": "[ISO 8601 date]",
      "rate_type": "[FIXED | VARIABLE]",
      "term_months": "[integer — certificates of deposit only]",
      "annual_percentage_rate": "[decimal — credit products]",
      "fee_schedule": [
        {
          "fee_type": "[MONTHLY_MAINTENANCE | TRANSACTION | OVERDRAFT |
                       EARLY_WITHDRAWAL | WIRE_TRANSFER]",
          "fee_amount": "[decimal currency value]",
          "waiver_conditions": "[plain-language waiver conditions, if any]",
          "fee_frequency": "[MONTHLY | ANNUAL | PER_TRANSACTION | ONE_TIME]"
        }
      ],
      "minimum_balance_requirement": "[decimal currency value]",
      "minimum_opening_deposit": "[decimal currency value]",
      "maximum_balance_limit": "[decimal currency value — if applicable]",
      "feature_set": ["MOBILE_DEPOSIT", "ATM_NETWORK_ACCESS", "JOINT_ACCOUNT",
                      "OVERDRAFT_PROTECTION", "BILL_PAY", "ZELLE"],
      "eligibility_confirmation": "[ELIGIBLE | LIKELY_ELIGIBLE | REQUIRES_REVIEW]",
      "approval_probability_signal": "[HIGH | MEDIUM | LOW | UNKNOWN]"
    }
  ],
  "institution_signature": {
    "algorithm": "ECDSA-P256",
    "signed_with": "[institution's current signing key from federated registry]"
  }
}
```

9.5 Fit-Based Ranking Engine

Thin-profile routing: Where Signal 4's behavioral consistency evaluation (Section 8.4) indicates a thin or temporally shallow financial profile — not necessarily synthetic, but insufficiently established for products requiring demonstrated financial history — the fit-based ranking must route the consumer toward products appropriate for their actual profile. This routing is consumer protection, not rejection: the consumer is presented with products they can genuinely qualify for rather than products that will result in declined applications.

The AFSP platform's fit-based ranking engine computes a FitScore for each responding institution's offer and ranks offers for presentation to the consumer. The ranking engine is structurally isolated from promotional weighting inputs.

> **Promotional weighting prohibition:** No institution may purchase, negotiate, or otherwise obtain a ranking position improvement through any commercial arrangement with the AFSP platform. Even if an institution has a commercial relationship with the AFSP platform operator, that relationship SHALL NOT influence the computed FitScore or the ranking of offers. The ranking algorithm computes scores independently of promotional inputs.

The FitScore formula:

```
FitScore(i) = SUM over k of (w_k × normalize(metric_k(i)))

Where:
  i          = responding institution offer index
  k          = metric dimension index
  w_k        = weight for metric dimension k (consumer-specified or default)
  normalize() = min-max normalization mapping metric value to [0, 1]
                across all responding institution offers (1.0 when fewer than two offers respond or all offers share the same value)
  metric_k(i) = value of metric dimension k for institution i

Default weights (consumer may specify custom weights summing to 1.0):
  w_rate     = 0.35  (rate competitiveness: APY for deposits, APR for credit)
  w_fee      = 0.25  (inverse of total annualized fee burden)
  w_balance  = 0.20  (compatibility of minimum balance with attested balance range)
  w_features = 0.20  (Jaccard similarity of feature set to consumer preferences)

Fee normalization:
  annualized_fee(i) = SUM of (monthly_fees × 12) + annual_fees
                      + (per_transaction_fees × attested_transaction_frequency)
  normalize(fee)(i) = 1 - min-max(annualized_fee(i))  [inverted: lower fee = higher score]

Balance compatibility:
  compat(i) = max(0.0, 1.0 - (min_balance(i) / consumer_attested_midpoint))
  If min_balance(i) > consumer_attested_range_upper: compat(i) = 0.0
```

9.6 Non-Participant Disclosure

When a consumer specifically inquires about a named institution that is not a current active network participant, the AFSP platform SHALL generate a Non-Participant Disclosure comprising:

(a) A factual, neutral statement that the named institution is not currently a participant in the AFSP inquiry network.

(b) A factual, neutral description of what network participation enables — without promotional language.

(c) Three consumer routing options:

Option A: Contact the institution directly through traditional channels.

Option B: Expand to Expanded Scope State — receive offers from all active network participants.

Option C: Return to Default Scope State — existing-relationship institutions only.

The AFSP platform SHALL log the non-participant inquiry event — including the named institution, the consumer's behavioral profile category, and the outcome (which routing option the consumer selected and, if Option B, which institution ultimately received the account opening) — as aggregate anonymized data for the Non-Participant Disclosure analytics product. No individual consumer data is stored or transmitted in this log.

Non-participating institutions MAY transition to active participant status at any time. Upon status change, the Non-Participant Disclosure for that institution is no longer generated.

10. AFSP-04: Clearance Return Protocol and Session Token

10.1 Overview

AFSP-04 defines the post-selection handshake that occurs after a consumer selects a specific product offer. The consumer's identity is disclosed to the selected institution for the first time only after the consumer selects an offer and the institution returns PROCEED_TO_KYC, through institution-controlled channels. The clearance return protocol governs how the institution evaluates the PCAA package, returns a signed clearance decision, and enables the AFSP platform to mint a session-scoped authorization token.

> **Platform absence from KYC:** The AFSP platform is explicitly absent from the Know-Your-Customer verification flow, the account opening decision, and all subsequent regulatory processes. The AFSP platform acts as a trust identity provider prior to KYC initiation. It does not receive, transmit, or process KYC data. The receiving institution's KYC collection occurs directly between the institution and the consumer through institution-controlled channels.

10.2 Clearance Handshake Sequence

The following sequence occurs after consumer product selection:

Step 1 — Consumer selection: Consumer explicitly confirms selection of a specific product offer from a specific institution.

Step 2 — PCAA package delivery: The AFSP platform transmits the complete PCAA package to the selected institution's registered API endpoint. The PCAA package is cryptographically bound to the selected institution's identifier and is invalid if presented elsewhere.

Step 3 — Institution evaluation: The receiving institution independently evaluates the PCAA package against its configured validation requirements and risk-tiered decisioning parameters. The institution returns one of three responses: PROCEED_TO_KYC, REQUEST_ADDITIONAL_VERIFICATION, or DECLINE_SESSION.

Step 4 — Session token minting (PROCEED only): Upon PROCEED_TO_KYC response, the AFSP platform validates the institution's signed clearance response and mints a session-scoped authorization token. This token is the sole mechanism by which the AI agent may initiate account opening with the institution.

Step 5 — Identity disclosure and KYC: The selected institution collects the consumer's legal identity (name, SSN or equivalent identifier, date of birth, address) directly through institution-controlled channels. No other institution receives the consumer's identity, and the AFSP platform does not receive or transmit it. The AI agent presents the session token to the institution's account opening API. The AFSP platform is absent from this step.

Step 6 — Other outcomes: If the institution returns REQUEST_ADDITIONAL_VERIFICATION, it specifies and conducts the verification through its own channels, and the session proceeds only if the institution then returns PROCEED_TO_KYC. If the institution returns DECLINE_SESSION, no session token is minted. In either case, the CLO informs the consumer as required by Section 5.9 (CD-6).

10.3 Session Token Structure

```
{
  "token_type": "AFSP_SESSION",
  "version": "0.1",
  "package_id": "[matching PCAA package identifier]",
  "session_id": "[session correlation identifier]",
  "institution_clearance": {
    "institution_id": "[institution identifier]",
    "clearance_decision": "PROCEED_TO_KYC",
    "clearance_timestamp": "[ISO 8601 UTC]",
    "institution_signature": "[ECDSA-P256 signature by institution signing key]"
  },
  "validity": {
    "not_before": "[ISO 8601 UTC]",
    "not_after": "[ISO 8601 UTC, max 1800 seconds]",
    "single_use": true
  },
  "platform_signature": {
    "algorithm": "ECDSA-P256",
    "signed_by": "platform.afsp.org"
  }
}
```

Single-use enforcement. The “single_use”: true field is normative. Receiving institutions SHALL maintain a registry of consumed session tokens and SHALL reject any token presented more than once, regardless of whether the token has otherwise passed signature and validity checks.

Validity window. The “not_after” field SHALL NOT exceed 1800 seconds from “not_before”. A session token presented after “not_after” SHALL be rejected as expired; the AI agent must obtain a new session token by returning through the clearance handshake (Section 10.2) rather than reusing an expired token.

Conditional minting. A session token is minted only upon a “PROCEED_TO_KYC” clearance decision (Section 10.2, Step 4). No session token is minted, issued, or made available to the AI agent for any other registered clearance decision (Section 4.3: “REQUEST_ADDITIONAL_VERIFICATION” or “DECLINE_SESSION”); the account opening flow SHALL NOT proceed in the absence of a validly minted session token.

Signature verification. Before accepting a session token as valid, the receiving institution's account opening API SHALL verify the “platform_signature” against the AFSP platform's published public key, and SHALL verify that “institution_clearance.institution_signature” matches the institution's own previously-issued clearance response (Section 10.2, Step 3). A session token whose “platform_signature” cannot be verified, or whose embedded “institution_signature” does not match the institution's own signing key, SHALL be rejected.

10.4 Audit Record

Upon session token minting, the AFSP platform generates a signed, tamper-evident audit record. The audit record:

Contains: Timestamps for each protocol stage, cryptographic references to each trust signal source identified by functional role (not entity name), the clearance validation outcome, and the session token identifier.

Does NOT contain: Personally identifiable information, raw financial data, or biometric data. Records are linked to sessions by pseudonymous identifiers.

Retention: Audit records are retained for a minimum of seven years, longer than the general five-year retention period for BSA records.

Accessibility: Audit records are available to authorized regulators and to the receiving institution upon written request.

11. AFSP-05: Behavioral Session Intelligence Interface

11.1 Overview

AFSP-05 defines a vendor-agnostic abstraction layer through which certified behavioral intelligence providers contribute a normalized session-initiation score (Signal S5) to the PCAA attestation package. The abstraction layer enables multiple competing providers to participate in the Signal 5 ecosystem without exposing their proprietary methodologies or their relative weighting to the receiving institution.

AFSP-05 governs the session-initiation composite score. Continuous per-transaction moment-level behavioral monitoring is governed by AFSP-10 (AFSP v1.0 roadmap — IP filed). AFSP-05 answers: was this session human-operated and behaving normally when it was initiated? AFSP-10 answers: is the session still human-operated at this specific moment?

11.2 Signal 5 Abstraction Layer Interface

Certified AFSP Signal 5 providers SHALL submit behavioral session intelligence scores via the following standardized interface:

```
{
  "provider_id": "[certified AFSP Signal 5 provider identifier]",
  "session_correlation_id": "[matching session identifier]",
  "score": {
    "raw_score": "[floating-point on provider native scale]",
    "normalized_score": "[floating-point 0.0 to 1.0]",
    "interpretation": {
      "0.0": "Lowest confidence — automated or coerced session strongly indicated",
      "1.0": "Highest confidence — genuine human-operated session"
    }
  },
  "signal_types_contributing": [
    "KEYSTROKE_DYNAMICS",
    "TOUCH_PATTERNS",
    "DEVICE_MOTION",
    "MOUSE_MOVEMENT",
    "NAVIGATION_TIMING",
    "DEVICE_HARDWARE",
    "AMBIENT_ENVIRONMENT_CONSISTENCY" // optional, Section 11.4 ECL
  ],
  "collection_duration_ms": "[passive signal collection duration in milliseconds]",
  "active_challenge_response_used": false,
  "attestation": {
    "algorithm": "ECDSA-P256",
    "signed_by": "[provider's registered signing key identifier]",
    "statement": "Score derived exclusively from passive behavioral observation
                  during the specified session. No active challenge-response
                  interaction was required."
  }
}
```

11.3 Composite Score Computation

The AFSP platform receives Signal 5 submissions from one or more certified providers and computes a composite score using a weighted ensemble methodology:

```
composite_score = SUM over p of (ensemble_weight(p) × normalized_score(p))

Where:
  p                = certified Signal 5 provider index
  ensemble_weight(p) = provider's assigned weight (proprietary platform configuration)
  normalized_score(p) = provider's submitted normalized score [0.0, 1.0]

Requirements:
  SUM of ensemble_weight(p) over all providers = 1.0
  Individual provider scores are discarded after composite computation
  Individual provider identities and weights are NOT transmitted to the
  receiving institution or recorded in the public audit log
  The receiving institution receives only the composite_score value
```

The Environmental Consistency Layer (Section 11.4) is not one of the ensemble providers in this formula and does not contribute a weighted term to composite_score. It operates as an independent gate: its Environmental Similarity Score triggers mandatory HITL step-up directly, on its own threshold, regardless of the composite_score value. This preserves ECL's ability to force step-up on a single strong signal rather than having that signal diluted by averaging against unrelated providers.

11.4 Environmental Consistency Layer (ECL)

The Environmental Consistency Layer (ECL) is a Signal 5 subsystem that evaluates the physical device ecosystem surrounding the session device as an independent trust signal. Where the operating system permits and the consumer consents, it uses ambient radio signals available to the application. Availability varies by platform, and some platforms randomize or restrict these signals. All environmental data is processed within the Trusted Execution Environment. Raw device identifiers are never stored or transmitted; only derived signals and flags leave the secure environment.

This layer addresses the physical device compromise attack vector: a stolen device operated from a foreign location with GPS spoofed to assert false domestic presence. This attack defeats Signal 2 (the attacker has the physical device) and is not detectable by Signal 5 behavioral monitoring alone if the attacker mimics normal usage patterns. The Environmental Consistency Layer detects this attack through three environmental cross-checks designed to make location spoofing harder to sustain without physical presence in the claimed location.

Component 1 — Cryptographic Environment Profile.

A differential-privacy-processed fingerprint of the device's typical ambient device ecosystem — the cluster of Bluetooth and NFC device identifiers that consistently appear across multiple sessions. Over a minimum of thirty (30) days of AFSP session activity, Signal 5 passively collects a hashed representation of nearby device identifiers processed through a differential privacy mechanism within the TEE, representing the statistical distribution of the environmental cluster rather than individual identifiers. At each session initiation, the current ambient cluster is compared against the established profile, producing an Environmental Similarity Score (0–1). Below the device-type-calibrated threshold, the session triggers mandatory HITL step-up. Upon successful HITL step-up in a mismatched environment, a temporary authorized variance is recorded (default 72 hours); sessions within an active variance window do not re-trigger for the same environmental cluster.

Component 2 — Device-Type-Aware Threshold Calibration.

Environmental mismatch thresholds are calibrated by device type, reflecting different legitimate mobility profiles. Desktop computers have fixed physical locations by definition, so any environmental mismatch beyond de minimis tolerance is a high-confidence anomaly — the strictest threshold. Laptop computers move between primary and established secondary environments; a laptop in a location matching no established cluster triggers HITL, while one in a previously certified secondary environment does not. Mobile phones accompany the consumer everywhere, so the threshold requires complete absence of any familiar signal within a configurable lookback window — not merely an unfamiliar environment. After thirty (30) days of session history, a continuously updated Environmental Mobility Index (EMI) personalizes threshold calibration beyond device-type defaults, with device type serving as a floor constraint.

Component 3 — Geographic Triangulation and Spoof Detection.

A device's GPS-asserted location is spoofable; the ambient cluster's geographic inference is harder to spoof, since matching it generally requires physical presence in a location where the surrounding device ecosystem matches the claimed geography. This component cross-references three independent geographic signals: (A) device-asserted GPS location — spoofable, used as the claimed baseline; (B) ambient cluster geographic inference — nearby Bluetooth device identifier prefixes, regional device distribution patterns, and WiFi network characteristics that imply physical location independently of GPS assertion; (C) network infrastructure geographic inference — autonomous system numbers, BGP routing characteristics, and carrier identifier patterns, partially spoofable with a VPN but not fully consistent with a fabricated cluster geography.

Geographic Risk Tier Classification: Geographic Risk Tier 1, Consistent domestic — standard session proceeds. Geographic Risk Tier 2, Domestic but unfamiliar environment — HITL step-up with temporary variance certification. Geographic Risk Tier 3, International with an established pattern — confirmation prompt, no full re-certification. Geographic Risk Tier 4, International with no prior pattern — full HITL step-up and temporary variance certification required. Geographic Risk Tier 5, High-risk jurisdiction or GPS-cluster mismatch — the session is suspended and routed to out-of-band reauthorization through the receiving institution, using channels independent of the session device; the session cannot resume through a prompt on the session device, since the mismatch itself indicates potential device compromise. Suspension pending reauthorization is consistent with the Non-Decline Principle (Section 1.5). [Note: these Geographic Risk Tiers are a distinct classification from the Authorization Tiers defined in Section 7.4 — the two systems govern different things and should not be conflated.]

Environmental Consistency Layer — v0.1 Status.

SHOULD requirement in v0.1. Threshold parameters will be defined as MUST in v0.2 after working group calibration from real deployment data. The differential privacy processing, TEE confinement, and no-raw-identifier-storage requirements are MUST in v0.1 for any implementation of this component.

12. AFSP-06: Non-Participant Disclosure Standard

12.1 Overview

AFSP-06 defines the governance mechanism that ensures consumers are accurately informed when a financial institution they inquire about is not an active AFSP network participant — and that preserves their ability to exercise their stated preferences, including choosing to contact a non-participating institution directly.

12.2 Participation Registry

The Participation Registry is governed by the AFSP Foundation under the Network Participation Agreement (Section 17.2) and operated on the AFSP platform. It is a distributed data structure mapping institution identifiers to their current certification status, signing keys, API endpoints, supported protocol versions, and product category declarations. The participation registry is the authoritative source for routing decisions and Non-Participant Disclosure triggers.

Each participation registry entry contains:

```
{
  "institution_id": "[globally unique institution identifier]",
  "institution_name": "[legal institution name]",
  "status": "[ACTIVE | INACTIVE | PENDING]",
  "status_change_date": "[ISO 8601 UTC]",
  "api_endpoint": "[registered AFSP API endpoint URL]",
  "certification": {
    "pathway": "[STANDARD | PLATFORM_OPERATOR]",
    "certified_date": "[ISO 8601 UTC]",
    "kyc_rigor_status": "[CURRENT | LAPSED]", // Section 17.2: public-record review and officer attestation
    "next_reverification_date": "[ISO 8601 UTC, annual]"
  },
  "signing_key": {
    "key_id": "[key identifier]",
    "public_key": "[ECDSA-P256 public key]",
    "rotation_timestamp": "[ISO 8601 UTC of last key rotation]"
  },
  "protocol_versions_supported": ["0.1"],
  "product_categories": ["SAVINGS", "CHECKING", "CD"],
  "geographic_markets": ["US-CA", "US-TX", "US-FL"]
}
```

12.3 Non-Participant Disclosure Content Requirements

Non-Participant Disclosures SHALL be:

Factual: Stating only verifiable facts about participation status.

Neutral: Neither promotional of network participants nor disparaging of non-participants.

Accurate: Reflecting the institution's current participation status at the time of the inquiry. Institutions may join the network at any time; Non-Participant Disclosures SHALL reflect current status, not historical status.

Consumer-empowering: Presenting the consumer's routing options without steering toward any particular choice.

12.4 Stated-Preference Data Logging

The AFSP platform logs each Non-Participant Disclosure event as part of the aggregate anonymized Non-Participant Disclosure analytics dataset. Logged data includes:

(a) The named non-participating institution identifier.

(b) The consumer's behavioral profile category (not individual consumer identifier).

(c) The consumer's selected routing option.

(d) If Option B (Expanded Scope) was selected: the institution that ultimately received the confirmed account opening, derived from the Return Trip API.

This aggregate dataset constitutes the Non-Participant Disclosure analytics product — made available to non-participating institutions as a participation incentive. No individual consumer data is included in any analytics product.

13. Return Trip API

13.1 Overview

The Return Trip API is the mechanism by which receiving institutions return signed confirmation of account opening outcomes to the AFSP platform. Return Trip API data closes the attribution loop between pre-opening inquiry data and post-opening behavioral outcomes, enabling the AFSP analytics product suite.

13.2 Return Trip Confirmation Schema

```
{
  "return_trip_type": "ACCOUNT_OPENING_CONFIRMATION",
  "session_id": "[matching session identifier]",
  "package_id": "[matching PCAA package identifier]",
  "institution_id": "[institution identifier]",
  "outcome": "[OPENED | DECLINED | PENDING | ABANDONED]",
  "outcome_timestamp": "[ISO 8601 UTC]",
  "opening_details": {
    "product_category": "[SAVINGS | CHECKING | CD | MMA]",
    "account_opened": true,
    "funded_within_7_days": "[true | false | unknown]",
    "active_within_30_days": "[true | false | unknown]"
  },
  "institution_signature": {
    "algorithm": "ECDSA-P256",
    "signed_by": "[institution signing key identifier]"
  }
}
```

Return Trip API confirmation is required for Level 2 Certified Implementation Partner status and Level 3 Active Network Participant status. Level 1 Founding Endorsers are not required to implement the Return Trip API. Founding Endorsers that operate as receiving institutions may implement it voluntarily, and must do so upon advancing to Level 2 or Level 3.

Return Trip API data is used exclusively for: (a) per-opening fee confirmation and settlement; (b) FitScore calibration and network intelligence; and (c) Non-Participant Disclosure analytics. Return Trip API data is not sold, shared with third parties, or used for any purpose other than those enumerated.

14. Certification and Participation Levels

14.1 Three Levels of Participation

AFSP participation exists on a spectrum of commitment and implementation. No organization is required to implement all components before endorsing the standard's direction.

Level 1 — Founding Endorser

Organizations that endorse AFSP as the right governance framework for the agentic banking channel and commit to working collaboratively toward implementation. Level 1 requires:

- A written letter of endorsement stating support for the AFSP governance framework

- Commitment to participate in the AFSP Technical Working Group
- No production implementation required at this level

Level 2 — Certified Implementation Partner

Organizations that have implemented one or more AFSP components and received AFSP certification, administered by the AFSP Foundation or, before its formation, by Primitive as interim protocol administrator. Level 2 requires:

- Completed technical implementation of assigned AFSP component(s) meeting specification requirements
- Successful completion of AFSP certification testing
- Execution of AFSP Certified Implementation Partner Agreement
- Ongoing compliance with certification program requirements

Level 3 — Active Network Participant

Organizations with live production implementations generating real AFSP-governed sessions. Level 3 requires:

- Level 2 certification for all applicable components
- Live production API endpoint registered in the participation registry
- Return Trip API implementation active
- Network Participation Agreement and Platform Access Agreement in effect (Section 17.2)

14.2 Performance Requirements

AFSP v0.1 implementations SHALL meet the following performance requirements for conformance:

Package assembly time: Total PCAA package assembly (all five signals) SHALL complete within 1,000 milliseconds (1 second) under normal operating conditions. Signals SHALL be assembled in parallel, not sequentially.

Attestation API response time: The AFSP platform attestation API SHALL respond to receiving institution queries within 100 milliseconds at the 95th percentile.

Agent registry query time: Agent provenance attestation API responses SHALL be delivered within 50 milliseconds at the 95th percentile.

Availability: The AFSP platform SHALL maintain 99.9% uptime availability on a rolling 30-day basis.

15. Regulatory Considerations

The three foundational principles governing AFSP's relationship to regulatory frameworks — Consumer Sovereignty, the Regulated Institution as Trust Anchor, and Protocol as Regulatory Reinforcement — are stated in Section 1.3. This section addresses specific regulatory frameworks and compliance considerations.

15.1 Bank Regulatory Accountability

AFSP is designed to preserve and reinforce — not diminish — the regulated financial institution's accountability for all consequential financial decisions. The following principles apply:

KYC accountability: The receiving financial institution retains sole KYC decision-making authority and full regulatory accountability for Customer Identification Program compliance under BSA/AML.

Credit decisioning: AFSP does not make credit, underwriting, or eligibility decisions. Receiving institutions make those decisions exclusively, using their own infrastructure and policies. Where an institution chooses to consider AFSP signals, it does so within its own decisioning and compliance framework.

Fair lending: The AFSP fit-based ranking engine is structurally isolated from promotional weighting and uses objective behavioral compatibility metrics. Institutions should satisfy themselves through their own fair lending analysis that AFSP-sourced inquiries comply with ECOA, Regulation B, and applicable fair lending requirements.

Model risk: AFSP is not intended to be a model within the meaning of SR 26-2 / OCC Bulletin 2026-13, which superseded SR 11-7 / OCC 2011-12, or successor guidance. It is designed as a pre-access attestation service that delivers trust signals; it does not generate decisions, recommendations, or consumer risk ratings for receiving institutions. Each receiving institution determines, under its own policies, whether and how to use AFSP signals, including how to treat them within its own risk management framework. AFSP providers will make reasonable documentation available on request to support an institution's due diligence.

15.2 CFPB and Section 1033

AFSP Signal 4 attestations are derived from consumer-permissioned financial data accessed under Section 1033 of the Dodd-Frank Act. AFSP:

Does not create new data access rights: AFSP operates within the consumer data rights framework established by Section 1033 and implementing regulations. Consumers authorize AFSP Signal 4 data access through the same consent flows required by Section 1033.

Implements data minimization: Signal 4 attestations carry only coarse derived signals, never raw balances, transaction records, or account numbers.

Consumer sovereignty: The agent ranks offers using the published FitScore formula and the consumer's own preference weights, explains the basis of the ranking, and never selects a product on the consumer's behalf. Implementers should confirm with counsel how their specific deployment is characterized under applicable law.

15.3 EU, Canada, and UK Extension (Future)

AFSP v0.1 is anchored on US regulatory frameworks (Section 1033, BSA/AML, FFIEC). Future versions are planned for EU (PSD3/PSR, DORA, GDPR, eIDAS 2.0), Canada (OSFI), and UK (UK Open Banking Standard v4.0, FCA AI governance guidance) markets. EU, Canadian, and UK implementations will require adaptation of the regulatory anchor language while preserving the core technical architecture.

16. AFSP v1.0 Roadmap

The following six components make up the AFSP v1.0 roadmap. They will be developed in the Technical Working Groups with founding endorsers and other participants, refined through implementation experience, and published as v1.0 standard components. Any organization may contribute through the extension process in Section 4. Provisional patent applications covering AFSP-07 through AFSP-11 were filed between May 1 and May 29, 2026, and applications covering AFSP-12 (Post-Opening Session Continuity) and other enhancements were filed on September 4, 2026. Section 17.5 lists the applications

and sets out the patent license commitment for implementers of this specification. See Section 5.7 for the AFSP-12 architectural framework.

COMPONENT	NAME	DESCRIPTION
AFSP-07	Agent Provenance and Certification Registry	Full certification authority infrastructure: developer registration, build submission and hash verification, automated static analysis plus human review, certified build registry with revocation history, real-time provenance attestation API. Three-tier credential hierarchy (Discovery/Transactional/Exception) with cryptographically enforced scope constraints.
AFSP-08	Federated Credential Verification	Receiving institution verification of agent credentials via federated registry without bilateral relationships or KYC data transmission. Directly answers S1 and S3 through cryptographic verification sequence.
AFSP-09	Pre-Execution Revalidation Protocol	Synchronous per-transaction revalidation intercepting every agent-initiated execution before it occurs. Multi-factor check against scope, velocity, behavioral envelope, and revocation. Reinforces downstream payment protocols by providing a tamper-detection baseline from session opening — any drift detectable before payment execution.
AFSP-10	Behavioral Envelope Engine (Moment-Level)	Per-credential ML behavioral baseline computed at every transaction execution event — not once at session initiation. Non-decline principle: anomalies route to human reauthorization. Cross-credential federated behavioral dataset. This is the IP home for moment-level continuous detection, complementing AFSP-05's session-initiation score.
AFSP-11	Tamper-Evident Regulatory Audit Log	Cryptographically chained append-only log of every revalidation request, response, and execution outcome across credential lifecycle. Each entry hashed to the preceding entry. 7-year retention (exceeds the BSA's general five-year period). Extends AFSP-04 from single clearance event to full credential lifecycle chain.
AFSP-12	Post-Opening Session Continuity	Three-leg triangulation re-authentication for agent-mediated servicing of AFSP-opened accounts: biometric continuity (same TEE root key), behavioral continuity (Signal 5 baseline comparison), relationship confirmation (Return Trip API cross-reference). Session Continuity Token; Re-Entry PCAA Package; SERVICE action class extension. See section 5.7.

17. Governance, Versioning, and Feedback

17.1 Current Governance

AFSP v0.1 is published and governed by Primitive as the founding author of the standard. Primitive invites the formation of an AFSP Technical Working Group comprising founding endorsers, implementing institutions, AI platform operators, data aggregators, and regulatory observers.

It is the intent of Primitive to transfer governance of the AFSP specification to an independent non-profit standards body — the Agentic Financial Services Protocol Foundation (the AFSP Foundation) — upon achievement of sufficient founding endorser participation. The Foundation will govern subsequent versions of the specification and administer certification under the Network Participation Agreement; until the Foundation is formed, Primitive serves as interim protocol administrator. Primitive will retain commercial rights to the reference implementation and to the commercial platform services governed by the Platform Access Agreement, including per-inquiry and per-opening economics, as the founding commercial contributor to the standard (Section 17.2).

17.2 Receiving Institution Certification Requirements

All Receiving Institutions seeking AFSP certification must satisfy the following eligibility requirements, available through Standard Certification or Platform Operator Certification pathways. Both produce equivalent AFSP certification status. AFSP certification is the single eligibility gate — every certified institution, whether receiving PCAA packages or issuing Signal 3 federated referrals, has satisfied these requirements.

1. Deposit Insurance.

Active FDIC insurance coverage confirmed against the FDIC BankFind Suite, or active NCUA insurance coverage confirmed against the NCUA Credit Union Locator. Coverage must be current and in good standing. The Foundation reverifies annually.

2. Valid Banking Charter.

A valid banking charter issued by the OCC, a state banking regulatory authority, the NCUA, or equivalent federal or state authority. Non-bank entities without deposit insurance are not eligible for Receiving Institution certification.

3. BSA/AML Program Attestation.

A signed attestation by the institution's Chief Compliance Officer that the institution maintains a BSA/AML compliance program meeting FinCEN requirements, including a Customer Identification Program (CIP), Customer Due Diligence (CDD) procedures, and Suspicious Activity Reporting protocols. This attestation is self-certified with legal liability attached.

4. Payment System Participation.

For institutions using AFSP's account funding architecture: active ACH network membership through a NACHA-approved financial institution, or equivalent real-time payment network participation. Institutions that will not use the account funding architecture may attest to this limitation and are exempt.

5. Regulatory Standing, Including KYC Rigor Condition.

The institution must not be subject to an active formal enforcement action, consent order, or cease-and-desist order from its primary federal or state regulator that would materially impair its ability to open new accounts or onboard new customers.

KYC Rigor Condition: The institution must not be subject to any active public formal enforcement action, including a consent order, cease-and-desist order, formal agreement, or civil money penalty, specifically related to Bank Secrecy Act compliance, Customer Identification Program adequacy, Customer Due Diligence standards, or identity verification practices, regardless of whether the action restricts new account opening. The Foundation will review publicly available enforcement records of FinCEN, the OCC, the Federal Reserve, the FDIC, the NCUA, and state banking regulators as part of the certification review. In addition, the institution's Chief Compliance Officer or BSA Officer shall attest annually, to the extent permitted by law, that the institution is not subject to any such action. AFSP does not request, receive, or rely on Confidential Supervisory Information (CSI), and nothing in this condition requires an institution to disclose information that it is prohibited by law from disclosing.

Design intent: AFSP does not define or assess how banks perform KYC — that is the bank's sovereign determination and regulatory obligation. This condition relies only on regulators' public enforcement actions and the institution's own attestation. If a regulator has taken public formal action concerning an institution's KYC practices, AFSP certification eligibility reflects that action while it remains in effect.

6. Network Participation Agreements.

Full participation in the AFSP network involves two separate agreements with two separate counterparties:

Network Participation Agreement (NPA) — with the AFSP Foundation. The NPA governs the institution's certified participation status, right to use the AFSP certification mark, inclusion in the Participation Registry, working group participation rights, and compliance with the specification's certification maintenance requirements. NPA fees — including initial certification fees, annual reverification fees, and working group participation fees — are payable to the Foundation. During the period prior to Foundation formation, NPAs are executed with CiDR Technologies, Inc. D/B/A Primitive as interim protocol administrator and shall be assignable to the Foundation upon formation without requiring institution counter-signature, subject to thirty (30) days written notice to the institution.

Platform Access Agreement (PAA) — with CiDR Technologies, Inc. D/B/A Primitive. The PAA governs the institution's access to Primitive's commercial platform services, including the PCAA Assembly Engine, CLO Registry, marketplace infrastructure, Return Trip API, and analytics products. Platform fees — including per-inquiry fees, per-opening fees, SCT session governance fees, and analytics subscription fees — are payable to Primitive. The PAA is a separate commercial agreement and is not governed by the Foundation.

The Foundation administers certification status through the NPA. Primitive operates the commercial platform through the PAA. Neither agreement substitutes for the other.

Platform Operator Certification Pathway.

Core banking platform operators that have obtained AFSP Platform Operator Certification enable their institutional clients to obtain AFSP certification through a simplified Downstream Institution Attestation rather than the full Standard Certification documentation package. Platform Operator Certification requires demonstrating that the platform operator's existing institutional client vetting process meets or exceeds all Standard Certification requirements above, including the KYC Rigor Condition. The platform operator carries the ongoing certification obligation; annual reverification is the platform operator's responsibility.

Downstream Institution Attestation requires: confirmation of membership in good standing on the certified platform; confirmation that no material change in regulatory status, charter, or compliance program has occurred since the platform operator's most recent vetting; and NPA execution. Platform operators must notify the Foundation within five (5) business days if any platform institution loses deposit insurance, has its charter restricted, becomes subject to a public KYC-related enforcement action, or is removed from the platform for compliance reasons.

Certification Maintenance and Revocation.

AFSP certification is subject to annual review. Certification may be suspended or revoked for: loss of deposit insurance or charter validity; new formal enforcement action restricting new account opening; new public KYC-related enforcement action triggering the KYC Rigor Condition; material BSA/AML compliance failure; NPA breach; or, for platform-certified institutions, removal from the certified platform for compliance reasons.

17.3 Version History

v0.1-draft (August 2026): Initial internal draft for discussion. Six v0.1 components (AFSP-01 through AFSP-06). Six v1.0 roadmap components identified (AFSP-07 through AFSP-12).

v0.1 Public Review Draft (September 28, 2026): Published for public comment. Adds roles and responsibilities (Section 1.6), institution-direct sessions (Section 5.5), intent verification checkpoint (Section 5.8), consumer disclosure requirements (Section 5.9), institution-bound nonce binding (Section 6.1), interim transport security (Section 6.3), enrollment portability (Section 7.2), the Authorization Tier architecture and Authorization Scope Envelope (Section 7.4), behavioral consistency evaluation requirements (Section 8.4), the environmental consistency layer (Section 11.4), the KYC Rigor Condition and the NPA and PAA participation structure (Section 17.2), items for further consideration (Section 18), and references (Appendix B).

17.4 Contributing to AFSP

Primitive welcomes contributions to AFSP from financial institutions, AI platform operators, data aggregators, card networks, core banking platform providers, and regulatory bodies. Contributions may take the following forms:

Comment on the v0.1 Public Review Draft: Submit written comments on any aspect of this specification to afsp@primitive.com or through the GitHub repository at https://github.com/primitive-os/AFSP. The comment period closes November 16, 2026. The specification and related materials are also available at agenticfinanceprotocol.org.

Join the Technical Working Group: Organizations interested in participating in the AFSP Technical Working Group should contact afsp@primitive.com. Working group participation is open to all organizations with a bona fide interest in the agentic banking channel.

Become a Founding Endorser: Organizations wishing to endorse AFSP v0.1 as a governance framework for the agentic banking channel should contact afsp@primitive.com (primary contact), or Derek White (CEO, Primitive) or Tom Wells (CRO, Primitive) directly.

17.5 Intellectual Property Notice

The AFSP specification is published under Fair, Reasonable and Non-Discriminatory (FRAND) licensing terms. Any organization may reproduce the specification without payment of royalties to Primitive, subject to attribution requirements.

The underlying technical mechanisms described in this specification are the subject of provisional patent applications (the “Primitive Patents”) filed by CiDR Technologies, Inc. D/B/A Primitive with the United States Patent and Trademark Office (USPTO), including at least the following applications: (1) Application No. 64/055,040, filed May 1, 2026, Attorney Docket 048627.00010; (2) Application No. 64/055,198, filed May 1, 2026, Attorney Docket 048627.00011; (3) Application No. 64/078,196, filed May 29, 2026, Attorney Docket 048627.00012; (4) Application No. 64/078,218, filed May 29, 2026, Attorney Docket 048627.00013; (5) Application No. 64/148,106, filed September 4, 2026. Non-provisional utility applications will be filed. Patent License Commitment (FRAND): Primitive is willing to license, on fair, reasonable, and non-discriminatory terms, the patent claims of the Primitive Patents essential to implementing this specification to any organization implementing AFSP in good faith and complying with the AFSP Specification. The terms of that license, including whether any royalty applies, will be published before v1.0. This FRAND commitment applies to the technical mechanisms specified in this document. It does not extend to: (a) access to the AFSP platform operated by Primitive; (b) the AFSP certification program or certification mark, which are governed by the Network Participation Agreement; (c) inclusion in the Participation Registry, which is governed by the Network Participation Agreement; (d) the Return Trip API infrastructure operated by Primitive; or (e) any analytics products or data services operated by Primitive. Organizations wishing to achieve AFSP certification must execute a Network Participation Agreement. Organizations wishing to obtain commercial services from Primitive must execute a Platform Access Agreement as described in Section 17.2.

Nothing in this specification should be construed as limiting any party's rights to develop alternative approaches to the problems addressed herein.

18. Items for Further Consideration

Version 0.1 establishes the AFSP architecture, trust signals, and information model. The following items are acknowledged for further consideration and refinement by the AFSP Technical Working Groups and will be developed through the extension process in Section 4 for incorporation in subsequent versions.

Protocol binding: Transport and endpoint definitions, message sequencing, error codes and retry behavior, message serialization and signature format, and certificate profiles, including the EKU value registry (Section 7.4).

Message formats: Normative formats for the Signal 1 agent provenance attestation, the Signal 3 federated referral attestation, the scope expansion token, the institution clearance response, the Servicing Intent Package, and the Discovery Declaration (Section 5.4).

Agent provenance for hosted agents: Mechanisms by which cloud-hosted agents demonstrate that the running agent matches its certified build, complementing build hash verification (VAL-06).

Trusted confirmation display: Rendering of the Intent Verification Checkpoint by components independent of the CLO's model and tool chain, and binding of the displayed text to the consumer's signature (Section 5.8).

Consumer lifecycle: Device replacement and re-enrollment, revocation of enrollment and data-access consent, and accessible alternatives to biometric user verification for consumers who cannot use biometrics.

Exception flows: Handling of REQUEST_ADDITIONAL_VERIFICATION, timeouts, and partial failures across the clearance handshake (Section 10.2).

Signal 4 privacy model: The privacy model for individual attestations, including treatment of repeated attestations across sessions, and the application of differential privacy to aggregate analytics products (Sections 8.2 and 12.4).

Environmental Consistency Layer: Feasibility and calibration across device platforms and operating system privacy controls (Section 11.4).

FitScore refinements: Consumer-specified hard constraints, modeling of fee waiver conditions, and the calibration methodology referenced in Section 13 (Section 9.5).

Conformance: Conformance requirements organized by role, covering CLOs, the AFSP platform, receiving institutions, and signal providers, with an accompanying conformance test suite.

Privacy considerations: A consolidated inventory of data flows, retention periods, and permitted uses across AFSP components.

Regulatory alignment: Alignment of Signal 4 data access with the Section 1033 framework as finalized, and regional profiles for EU, Canada, and UK markets (Section 15.3).

Institution-direct participation: A participation mode for institutions that use AFSP without receiving marketplace inquiries, including registry participation flags, product-level opt-in, and related agreement terms. Consistent with Principle 1 (Section 1.3) and Section 5.5, such a mode would govern only which products an institution offers through AFSP, and Non-Participant Disclosure (Section 12) would apply only to institutions outside the network.

Additional marketplace operators: Criteria and a certification process under which the AFSP Foundation could certify marketplace operators beyond the network's initial marketplace, including neutrality requirements such as published ranking, no paid placement, and consumer data protections (Section 5.5).

AFSP v0.1 Public Review Draft · Agentic Financial Services Protocol · Primitive · September 28, 2026 · agenticfinanceprotocol.org · https://github.com/primitive-os/AFSP · afsp@primitive.com

License: Fair, Reasonable and Non-Discriminatory (FRAND) terms · See Section 17.5

Appendix A — Non-Normative: End-to-End Worked Example

> **This appendix is non-normative.** The values shown are illustrative. They demonstrate the protocol flow but do not define conformance requirements. Conformance is determined solely by the normative sections of this specification.

This appendix traces a complete AFSP-governed session from initiation through confirmed account opening. It illustrates the full architecture including the discovery phase (Phase 1) and the authorization and attestation phase (Phase 2). The example uses a standalone AI assistant to demonstrate the AI platform CLO entry point type. The protocol flow is identical regardless of which AI platform or entry point type initiates the session.

A.1 Scenario

Consumer: Alice Chen, existing customer of her primary financial institution (an AFSP network participant). Enrolled in AFSP through her primary financial institution's mobile app six months ago.

Agent: the AI assistant, acting as a general-purpose AI assistant with AFSP financial product discovery enabled. The AI platform operator is a certified AFSP Consumer Layer Operator (AI Platform entry point type). Agent build certified in the AFSP agent registry (agent_id: agt_aip_assistant_20260801_prod, build hash: sha256:c3f9a1e2... truncated).

Session goal: Alice asks the AI assistant: “Can you find me the best high-yield savings account? APY above 4.5%, no monthly fees, minimum balance under $500.” The AI assistant identifies this as financial product discovery intent, checks its AFSP tool registration, and initiates a session.

Session state: Alice is using her AI assistant on her enrolled smartphone (device_id: dev_9f3a1c2e). Her AFSP biometric enrollment was established through her primary financial institution's app. The AI assistant, as the certified CLO, presents the AFSP authorization request.

A.2 Step 0 — Discovery (Phase 1, Pre-AFSP)

Alice asks her AI assistant: “Can you find me the best high-yield savings account? APY above 4.5%, no monthly fees, minimum balance under $500.” The AI assistant identifies this as financial product discovery intent and checks its registered AFSP financial services tool. The AI assistant also queries the AFSP Discovery Declaration published by her primary financial institution, confirming they are an active network participant. Having confirmed AFSP availability, the AI assistant initiates Phase 2.

A.3 Step 1 — Consumer Authorization (AFSP-01)

Alice reviews her agent's proposed action scope and authenticates with Face ID. Her TEE generates the authorization token:

```
{
  "token_type": "AFSP_CONSUMER_AUTH",
  "version": "0.1",
  "consumer_device_id": "dev_9f3a1c2e4b5d6f7a8b9c0d1e2f3a4b5c",
  "biometric_event": {
    "type": "FIDO2_BIOMETRIC",
    "timestamp": "2026-08-18T14:23:07Z",
    "tee_attestation": "eyJhbGciOiJFUzI1NiJ9..." // device TEE hardware attestation
  },
  "authorized_agent_id": "agt_aip_assistant_20260801_prod",
  "receiving_institution_id": "NETWORK_MARKETPLACE", // expanded scope requested
  "scope": {
    "action_classes": ["READ", "DISCOVER", "PRESENT"],
    "product_categories": ["SAVINGS", "MMA"],
    "consumer_preference_weights": {
      "w_rate": 0.50, "w_fee": 0.30,
      "w_balance": 0.10, "w_features": 0.10
    }
  },
  "temporal_validity": {
    "not_before": "2026-08-18T14:23:07Z",
    "not_after":  "2026-08-18T15:23:07Z"
  },
  "scope_expansion": {
    "default_state": "EXISTING_RELATIONSHIPS_ONLY",
    "expansion_requires": "BIOMETRIC_RECONFIRMATION"
  },
  "signature": "MEYCIQDx9a2..." // ECDSA-P256, TEE root private key
}
```

Alice’s session defaults to DEFAULT SCOPE given her existing relationship with her primary financial institution — meaning the reverse inquiry would normally go only to that institution. Alice instead chooses to expand scope to see all active network participants. She performs a second Face ID authentication, which produces a signed scope expansion token (not shown), activating EXPANDED SCOPE as defined in Section 9.2. This is why all seven network participants receive her inquiry in Step 4.

A.4 Step 2 — Signal 4 Assembly (AFSP-02)

Simultaneously, the AFSP platform queries Alice's consented certified Signal 4 provider. The provider assembles the financial profile attestation through its six-stage pipeline:


```
{
  "attestation_provider": "certified_s4_provider_001",
  "attestation_version": "0.1",
  "inquiry_event_hash": "sha256:7f3a9c2e...",
  "signals": {
    "account_longevity":     { "value": "MATURE",           "confidence": 0.94 },
    "income_stability":      { "value": "HIGHLY_CONSISTENT",
                               "frequency_component": 0.97,
                               "regularity_component": 0.95,
                               "magnitude_component": 0.91,
                               "confidence": 0.93 },
    "balance_range":         { "value": "MID_HIGH",         "confidence": 0.88 },
    "transaction_regularity":{ "value": "HIGHLY_REGULAR",
                               "behavioral_category": "SAVER","confidence": 0.91 },
    "relationship_breadth":  { "value": "MODERATE",         "confidence": 0.86 },
    "product_eligibility":   { "categories_indicated": ["SAVINGS","MMA"],
                               "minimum_balance_compatible": true, "confidence": 0.95 }
  },
  "behavioral_consistency_score": {
    "score": 0.89,
    "markers_evaluated": 5
  },
  "privacy_parameters": {
    "method": "DIFFERENTIAL_PRIVACY_LAPLACE", "epsilon": 1.0, "k_anonymity_k": 5
  },
  "anti_replay": {
    "nonce": "a3f7c2e9b1d4f6a8c0e2b4d6f8a0c2e4",
    "sequence_number": 847
  },
  "temporal_validity": {
    "not_before": "2026-08-18T14:23:08Z",
    "not_after":  "2026-08-18T15:23:08Z"
  },
  "signature": { "algorithm": "ECDSA-P256",
    "certificate_chain": "MIIBkTCB+..." }
}
```


A.5 Step 3 — Signal 5 Assembly (AFSP-05)

The certified Signal 5 provider returns a behavioral session intelligence score based on Alice's passive interaction patterns during the session initiation:


```
{
  "provider_id": "certified_s5_provider_001",
  "session_correlation_id": "sess_2e4f6a8c0b2d4f6a",
  "score": {
    "raw_score": 847.3,
    "normalized_score": 0.91
  },
  "signal_types_contributing": [
    "TOUCH_PATTERNS", "DEVICE_MOTION", "NAVIGATION_TIMING", "DEVICE_HARDWARE"  // ECL (Section 11.4) not
active this session — SHOULD, not MUST, in v0.1
  ],
  "collection_duration_ms": 4200,
  "active_challenge_response_used": false,
  "attestation": {
    "algorithm": "ECDSA-P256",
    "signed_by": "s5_provider_signing_key_2026_q3",
    "statement": "Score derived exclusively from passive behavioral observation."
  }
}
```

A.6 Step 4 — PCAA Package Assembly and Reverse Inquiry Submission

The AFSP platform assembles all signals into a signed PCAA package retained by the AFSP platform, and submits a reverse inquiry payload to the marketplace. The reverse inquiry contains the consumer's behavioral profile and preference parameters. The PCAA package is held by the AFSP platform pending institution selection in Step 5. The full PCAA package is not transmitted to any institution at this step:

```
{
  "afsp_version": "0.1",
  "package_id": "pkg_1a3c5e7f9b2d4f6a8c0e2b4d6f8a0c2e",
  "session_id": "sess_2e4f6a8c0b2d4f6a",
  // no institution nonce yet: it is added when the package is bound to the selected institution (Section 6.1)
  "assembly_timestamp": "2026-08-18T14:23:09Z",
  "routing_destination": "AFSP_MARKETPLACE_REVERSE_INQUIRY",
  "validity_window": {
    "not_before": "2026-08-18T14:23:09Z",
    "not_after":  "2026-08-18T15:23:09Z"
  },
  "signals": {
    "S1": { /* agent provenance attestation — agt_aip_assistant_20260801_prod verified, build hash matched */ },
    "S2": { /* Alice's authorization token and scope expansion token — Face ID binding confirmed */ },
    "S3": { /* federated referral from Alice's primary financial institution — Alice KYC'd 6 months ago */ },
    "S4": { /* certified S4 provider financial profile attestation — MATURE/HIGHLY_CONSISTENT/MID_HIGH */ },
    "S5": { /* certified S5 provider session score — 0.91 normalized; may include an optional Environmental Consistency Layer
subsystem component (Section 11.4) */ }
  },
  "signals_present": ["S1","S2","S3","S4","S5"],
  "platform_signature": {
    "algorithm": "ECDSA-P256",
    "signed_by": "platform.afsp.org",
    "signature": "MEUCIHx7a3..."
  }
}
```

A.7 Step 5 — Institution Responses and FitScore Ranking

Seven participating institutions respond with product offers within 800ms. The AFSP platform computes FitScores. The top three results (non-normative illustration):

> **Institution:** Sunrise Digital Bank APY: 5.10% | Monthly fee: $0 | Min balance: $100 | FitScore: 0.94 Institution: Pacific Savings Co-op APY: 4.95% | Monthly fee: $0 | Min balance: $250 | FitScore: 0.89 Institution: Harbor Credit Union APY: 4.80% | Monthly fee: $0 | Min balance: $0 | FitScore: 0.81 Note: Promotional weighting excluded. Rankings reflect the published FitScore formula. Alice's preference weights applied: rate 50%, fee 30%, balance 10%, features 10%.

Alice reviews the ranked offers through her agent's interface. She selects Sunrise Digital Bank. This triggers Step 6, the clearance handshake; her identity is disclosed to Sunrise Digital Bank only after it returns PROCEED_TO_KYC.

A.8 Step 6 — Clearance Handshake and Account Opening (AFSP-04)

Alice reviews the IVC plain-language confirmation (institution, product and key terms) and confirms, then completes a further Face ID authentication (Authorization Tier 2 / INITIATE action class requires IVC confirmation followed by step-up).

The Servicing Intent Package presented to Alice at the IVC confirmation (Section 5.8):


```
{
  "package_type": "SERVICING_INTENT_PACKAGE",
  "afsp_version": "0.1",
  "session_id": "sess_2e4f6a8c0b2d4f6a",
  "action_class": "INITIATE",
  "authorization_tier": "Authorization Tier 2 (INITIATE)",
  "action_type": "ACCOUNT_OPENING",
  "receiving_institution": "Sunrise Digital Bank",
  "product": "High-Yield Savings Account",
  "product_terms": { "apy": "5.10%", "monthly_fee": "0.00", "minimum_balance": "100.00", "deposit_insurer": "FDIC" },
  "amount": null,
  "currency": "USD",
  "timing": "immediate",
  "ivc_display_text": "Open a High-Yield Savings Account at Sunrise Digital Bank? APY 5.10%, no monthly fee, $100 minimum
balance, FDIC-insured.",
  "ivc_confirmed": true,
  "ivc_confirmation_timestamp": "2026-08-18T14:24:12Z",
  "generated_from_structured_fields": true
}
```


The AFSP platform then sends a bound PCAA package to Sunrise Digital Bank's registered endpoint. Sunrise Digital Bank validates the package (VAL-01 through VAL-04a, VAL-05 through VAL-10), evaluates it against its risk-tiered decisioning framework, and returns a clearance response within 180ms:


```
// Sunrise Digital Bank clearance response
{
  "package_id": "pkg_1a3c5e7f9b2d4f6a8c0e2b4d6f8a0c2e",
  "institution_id": "sunrise_digital_bank_afsp_001",
  "clearance_decision": "PROCEED_TO_KYC",
  "clearance_timestamp": "2026-08-18T14:24:15Z",
  "institution_signature": "MEQCIBx9..."
}
// AFSP platform mints session token
{
  "token_type": "AFSP_SESSION",
  "version": "0.1",
  "package_id": "pkg_1a3c5e7f9b2d4f6a8c0e2b4d6f8a0c2e",
  "session_id": "sess_2e4f6a8c0b2d4f6a",
  "institution_clearance": {
    "institution_id": "sunrise_digital_bank_afsp_001",
    "clearance_decision": "PROCEED_TO_KYC",
    "clearance_timestamp": "2026-08-18T14:24:15Z",
    "institution_signature": "MEQCIBx9..."
  },
  "validity": {
    "not_before": "2026-08-18T14:24:15Z",
    "not_after": "2026-08-18T14:54:15Z", // 30 min window
    "single_use": true
  },
  "platform_signature": {
    "algorithm": "ECDSA-P256",
    "signed_by": "platform.afsp.org",
    "signature": "MEUCIQDk7..."
  }
}
```


At this point, Sunrise Digital Bank’s own account opening API takes over the session: it initiates KYC collection directly with Alice through its own interface, and the account opening application flow proceeds entirely within Sunrise Digital Bank’s infrastructure. Alice’s legal identity is collected and verified by Sunrise Digital Bank — the AFSP platform has no visibility into or involvement in this step. Sunrise Digital Bank delivers its account agreement and required account disclosures through its own channel before opening the account (Section 5.9, CD-5). Alice completes KYC in about 100 seconds, and the account is approved and opened.

A.9 Step 7 — Return Trip Confirmation

Sunrise Digital Bank returns a signed confirmation to the AFSP platform's Return Trip API:

```
{
  "return_trip_type": "ACCOUNT_OPENING_CONFIRMATION",
  "session_id": "sess_2e4f6a8c0b2d4f6a",
  "package_id": "pkg_1a3c5e7f9b2d4f6a8c0e2b4d6f8a0c2e",
  "institution_id": "sunrise_digital_bank_afsp_001",
  "outcome": "OPENED",
  "outcome_timestamp": "2026-08-18T14:25:57Z",
  "opening_details": {
    "product_category": "SAVINGS",
    "account_opened": true,
    "funded_within_7_days": "unknown",   // not yet known at the time of opening
    "active_within_30_days": "unknown"
  },
  "institution_signature": {
    "algorithm": "ECDSA-P256",
    "signed_by": "sunrise_digital_bank_signing_key_2026"
  }
}
```

A.10 Session Timeline Summary

Happy Path New Account Opening (Discovery → Account Opened)
14:23:07Z — Alice authenticates via Face ID. Authorization token generated in TEE.
14:23:08Z — Alice reconfirms with Face ID to expand scope to all network participants; scope expansion token signed.
14:23:08Z — Signal 4 assembled from the certified S4 provider (parallel with S1, S2, S5 assembly).
14:23:09Z — PCAA package assembly and reverse inquiry submission and broadcast.
14:23:09Z — 7 institutions receive reverse inquiry payload.
14:23:10Z — Offers returned by all 7 institutions and ranked by FitScore.
14:24:10Z — Alice selects Sunrise Digital Bank.
14:24:12Z — Alice reviews the IVC plain-language confirmation and confirms.
14:24:13Z — Step-up Face ID authenticates the Authorization Tier 2 (INITIATE) action.
14:24:14Z — Bound PCAA package (with institution-specific nonce) delivered.
14:24:15Z — Sunrise Digital Bank returns PROCEED_TO_KYC in 180ms (including VAL-04a nonce check). Session token minted.
14:24:16Z — KYC collection begins via Sunrise Digital Bank's interface.
14:25:57Z — Account approved and opened. Return Trip confirmation received.

Total elapsed time from authorization to account open: 2 minutes, 50 seconds.
Identity disclosed to one institution. Consumer intent verified directly against the actual product selected.
AFSP platform absent from KYC and the account-opening decision.
The Bank remains the regulated trust anchor.

AFSP v0.1 Public Review Draft · Appendix A · Non-Normative · Primitive · September 28, 2026

Appendix B — References

The following documents are referenced in this specification.

Standards and specifications

FIDO Alliance and W3C. FIDO2: Web Authentication (WebAuthn) and Client to Authenticator Protocol (CTAP).

Financial Data Exchange. FDX API.

IETF. RFC 5280, Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile.

IETF. RFC 6749, The OAuth 2.0 Authorization Framework.

IETF. RFC 8446, The Transport Layer Security (TLS) Protocol Version 1.3.

IETF. RFC 9396, OAuth 2.0 Rich Authorization Requests.

ISO 20022. Financial services — Universal financial industry message scheme.

NIST. FIPS 180-4, Secure Hash Standard (SHS).

NIST. FIPS 186-5, Digital Signature Standard (DSS).

OpenID Foundation. FAPI 2.0 Security Profile.

Laws, regulations, and supervisory guidance

Bank Secrecy Act recordkeeping: 31 CFR 1010.430, Nature of records and retention period.

Deposit insurance advertising and signage: FDIC, 12 CFR Part 328; NCUA, 12 CFR Part 740.

Dodd-Frank Wall Street Reform and Consumer Protection Act, Section 1033 (12 U.S.C. 5533), and implementing regulations (12 CFR Part 1033).

Equal Credit Opportunity Act and Regulation B (12 CFR Part 1002).

FFIEC. Authentication and Access to Financial Institution Services and Systems (2021).

OCC Bulletin 2026-13 and Federal Reserve SR 26-2, Revised Guidance on Model Risk Management (April 17, 2026), superseding OCC Bulletin 2011-12 and SR 11-7.

Truth in Savings: Regulation DD (12 CFR Part 1030) and, for federally insured credit unions, 12 CFR Part 707.

AFSP v0.1 Public Review Draft · Appendix B · Non-Normative · Primitive ·
September 28, 2026