

AGENTIC PRIMITIVES · AGENTIC TRUST LABS

Production Readiness Assessment

Packages · Contracts · Estate · Capability areas against the field

18 September 2026 · repository HEAD `2c711e01`

77 packages · 50 contracts · 939 / 939 contract tests · 54 open findings (2 Critical · 5 High)

Self-assessment by the maintainers from the repository's own evidence. **Not a third-party audit.**

Verdict: pre-production, honestly labelled.

GO for testnet pilots and demonstrations on our own chain. NO-GO for a real person's private data, for real value under delegated payments, and for a public mainnet — each with named, dated closing conditions (\$2.3). Thirty of the thirty-four findings raised by the 17 September review were closed in source within a day; the five contract fixes are generation 2 and no environment runs generation 2 yet.

Next full review: close of Wave R4 or 17 November 2026, whichever is first.

Verdict: pre-production, honestly labelled. GO for testnet pilots and demonstrations on our own chain. NO-GO for a real person's private data, for real value under delegated payments, and for a public mainnet — each with named, dated closing conditions (§2.3).

This document is what we would hand a senior third-party auditor or a technical due-diligence team today. It is a **self-assessment** produced by the maintainers from the repository's own evidence — the findings ledger, the continuous gates, the test suites, and a line-by-line review of the deployed configuration — completed on 17 September 2026 and re-verified on 18 September after a hardening wave closed thirty of the thirty-four findings it raised. Where something is unproven we say so; where something is fixed in source but not yet deployed we say that too, because the difference is the whole point.

How to read this document

Section	Question it answers
§1 What Agentic Primitives is	What the substrate does and the one property it exists to enforce
§2 Executive assessment	The verdict, what held, what did not, and the segmented posture
§3 Packages, by capability area	An assessment of each of the 77 TypeScript packages
§4 Contracts	An assessment of each of the 50 Solidity sources and what is deployed
§5 The estate	An assessment of each deployed service and of the operating posture
§6 Capability areas against the field	For each area we cover: who else is in it and where we stand
§7 Open findings register	Every Critical and High, the Mediums by class, and the remediation order
§8 Method and how to verify	How the grades were reached and how a stranger re-runs the evidence

The grades

Grade	Meaning
Green	The component does what its specification claims, its tests exercise the claim, and no open finding names it.
Amber	The property exists, but a Medium finding names it, or it is early (skeleton / scaffold / deprecated alias), or its evidence is partial. Safe to build on; watch the note.
Red	A Critical or High finding names the component's own code.

Two distinctions govern every grade below. **Closed in source is not deployed** — a fix that has landed with tests but has not been rolled to the environments people use does not reduce live exposure, and we grade the estate on what is running. **A finding against how a component is configured in the estate is not a finding against the component** — those are graded in §5 and referenced, not double-counted, in §3.

1. What Agentic Primitives is

Agentic Primitives is a substrate — packages and contracts, not a platform — for building software agents that act under **authority the agent cannot exceed**. Every person, organization and service is an ERC-4337 smart account (a Smart Agent); authority to act flows from a principal to an agent as

a signed, attenuable, on-chain-revocable delegation (ERC-7710-shaped) whose caveats bind what may be done, to whom, for how much, and — through an intent digest — for which specific request. A planner may be probabilistic; the check that stands between the plan and the effect is not.

Five properties are enforced by contract rather than by policy code, and no product we have surveyed composes all five:

- 1. **Identity survives the runtime** — the agent is the account, not a token issued by whoever hosts it.
- 2. **Attenuation is a chain** — each hop of delegation can only narrow the previous one, verified on chain.
- 3. **Revocation is checked per step** — not at login, not per session; before each effect.
- 4. **The intent is a caveat** — the digest of what the person approved binds the delegation that acts on it.
- 5. **Single-use by construction** — a digest-bound redemption is consumed once, on chain.

Around that core sit the layers a working agent needs: naming and discovery, a private vault with per-record delegation scope, an authority-aware harness (Ask → Intent → Mandate → Plan → per-step Verify → Receipt), an admission edge for A2A and MCP, provenance and receipts, coordination between agents, and a formal ontology that the code binds to by IRI. The nine capability areas in §3 and §6 are these layers.

The repository is Ring 0: primitives only. Products (the Home, the discovery registry, the naming service) are built from it and deployed as the **estate** assessed in §5.

2. Executive assessment

2.1 The numbers

Packages	77 (@agenticprimitives/*), 97k lines of TypeScript; every package ships README , CLAUDE.md , AUDIT.md ; 76/77 have tests
Contracts	50 Solidity sources, 10.7k lines; 939 / 939 Foundry tests green at HEAD; 7 symbolic proofs (halmos); echidna nightly, medusa weekly, three invariant suites
Gates	131 check:* gates in CI; 57 live gates nightly with negative twins; A2A 1.0 official TCK green (MUST 88 / SHOULD 8 / MAY 4)
Findings ledger	54 open (2 Critical + 1 external tracker · 5 High · 30 Medium · 16 Low); machine-checked for freshness on every CI run
Last full review	17 September 2026 (this document's basis), 1,543 commits after the previous one (12 July 2026)
Hardening since	30 of 34 findings raised on 17 September closed in source by 18 September; all 9 new Highs among them
Deployment	private chain faithchain (34348) and Base Sepolia testnet; no public mainnet; contract fixes of 18 September are generation 2, no environment runs generation 2 yet

2.2 What held, and what did not

Held. The authority plane is not probabilistic in code. Every tool that declares a capability is verified per step against a signed, chain-live, intent-bound mandate before it runs; the planner picks from an exact tool set and cannot rename a tool, lower its risk, or mint a mandate; unattended runs present no mandate and park for the person; payments are single-use on chain; a resumed run re-verifies against current chain state; a replay never re-runs a tool. Identity keys are delegates of a

custodied identity, never the identity itself. The package dependency graph is acyclic with no back-edge found in a workspace-wide sweep. Records are portable: a person's data can be carried to a second Home and every record read by session alone, with the receipt anchor read back from chain.

Did not hold (17 September), now closed in source (18 September). One unmodelled contract boundary — the account's `execute` reaching contracts that treat the account itself as principal below the custody tier — produced both new contract Highs (a custodian could brick a custody-governed account; approved-hash delegations survived custody recovery). The delegation core's session binding was looser than doctrine (a session key could redeem any delegation its principal ever signed). Stewardship of an organization's agent had regressed to self-assertion. Three harness perimeters were prompt-only (planner-callable memory writes; an external server's `readOnlyHint` deciding that no mandate was needed; the mandate not binding the plan's other arguments). Several CI gates claimed more than they read. All of these are fixed with tests; the contract fixes await deployment.

Still not held (operations and posture). No governance, timelock or guardian is deployed on the chain that is the authority substrate, so the live stack cannot be halted or upgraded under governance. The custody master key runs from a local secret in every environment. The one record store's backup is unscheduled. There is no alerting. The managed-KMS caller identity is shared across two estates. No code change closes these; they are the work of the next wave and the reason the verdict has not moved.

2.3 Segmented posture

Scope	Verdict	Closes when
Testnet pilot, demonstrations, personas on faithchain	GO	—
A real person's private data in a vault	NO-GO	personas isolated from real people; a crypto-shred (right-to-erasure) path; connected-app wires under a managed key rather than one Worker secret
Real value under delegated payments (beyond mock USDC)	NO-GO	above + generation-2 contracts deployed; the local custody master retired to managed KMS; a per-estate KMS caller identity; tested vault-key recovery and a scheduled backup
Public mainnet / L2	NO-GO	above + governance deployed; a fixed-commit third-party contract audit (the contract layer is now ready to freeze once generation 2 is the deployed generation)

3. Packages, by capability area

Seventy-seven packages, graded on their own code. "Tests" counts test files in the package. "Open findings" lists ledger items whose declared concern is a path in the package; where the finding is really about how the estate configures the package, the note says so and the grade is Amber, not Red. Versions: eighteen packages sit on the `1.0.0-alpha.25` core lineage; the remainder on `0.0.0-alpha.x`. A package on the `0.0.0` lineage with fewer than two test files is graded Amber for that reason alone.

Summary: 48 Green · 29 Amber · 0 Red. No package is Red on its own code. The Ambers fall into four classes: a Medium finding names the package (14), an estate-scoped Critical/High references it (3: `key-custody`, `contracts`, `fabric`), early stage (3: `intent-resolver` W1 skeleton, `content-storage` scaffold, the deprecated `agent-skills` alias), or thin tests on the `0.0.0` lineage (9).

Identity (19 packages)

Package	What it does	Version	Tests	Open findings	Grade	Note
types	Cross-cutting branded types and chain primitives shared across @agenticprimitives/* packages	1.0.0-alpha.251	—	—	Green	Zero-dependency root of the graph; the branded Address/Hex types every package shares.
agent-account	ERC-4337 smart-account substrate	1.0.0-alpha.2511	—	—	Green	ERC-4337 account SDK; CA-1 timelocked-upgrade surface still absent from the SDK (NEW-AA-3, low).
account-custody	Custody-layer SDK: CustodyPolicy ABI, CustodyAction enum + arg builders, EIP-712 typed-data helpers, and Custodian/Trustee/CustodyCouncil ty	1.0.0-alpha.254	—	—	Green	CustodyPolicy ABI + EIP-712 helpers; custody-epoch mirror added 2026-09-18.
key-custody	Pluggable envelope encryption + signers + HMAC providers (local-AES / AWS KMS / GCP KMS / AKCS agentic-kms)	1.0.0-alpha.2517	Crit D-P0-1 · High D-P0-2	—	Amber	The provider seam (local-AES / AWS / GCP / AKCS). The two estate findings on it (D-P0-1, D-P0-2) are about how the ESTATE configures it — the local-aes custody master and KEK-version durability — not defects in the package.
browser-identity	Browser sign-in adapter seam	1.0.0-alpha.8 1	—	—	Green	
agent-naming	Agent Naming Service SDK: typed hierarchical naming for Smart Agents (.me .org .team .svc .workspace .treasury — a suffix names the derived	1.0.0-alpha.2510	—	—	Green	Typed suffix naming (.me .org .team ...). Reclaim/transfer leaves stale resolver rows on chain (NAMING-RECLAIM, medium) — a contract fix.
agent-profile	Agent identity SDK: AP-native Canonical Agent Profile (CanonicalAgentProfileV1) + signed profile bundle, CAIP-10 native helpers, endpoint	1.0.0-alpha.2517	Med S-3-conjunction	—	Amber	Canonical profile + signed A2A card projection; the cardΛproofΛregistration conjunction is not yet one call (S-3, medium).
agent-relationships	Agent relationships SDK: trust-fabric edge store (subject, object, relationshipType) with role taxonomy. Phase 1 ships edge-ID derivation +	1.0.0-alpha.255	—	—	Green	
agent-resolution	Provider-neutral agent resolution primitives	0.0.0-alpha.6 10	—	—	Green	
identity-directory	Evidence-backed read model over canonical agents and their facets (ADR-0015 / spec 223)	1.0.0-alpha.251	—	—	Green	
identity-directory-adapters	Port implementations for @agenticprimitives/identity-directory	1.0.0-alpha.251	—	—	Green	
fedcm-idp	FedCM IdP contract as pure builders + validators	1.0.0-alpha.8 1	—	—	Green	
fedcm-rp	Relying-party FedCM wrapper	1.0.0-alpha.8 1	—	—	Green	

Package	What it does	Version	Tests	Open findings	Grade	Note
connect	Agentic Connect — the SSO broker (spec 224 / ADR-0014)	1.0.0-alpha.251	Med	session-rotation-revocation-unwired	Amber	SSO broker. Rotated-out credential sessions are not yet checked against the S-2 revocation seam by any consumer (medium).
connect-auth	User authentication (passkey + SIWE + Google OAuth), JWT sessions, and pluggable signer interfaces	1.0.0-alpha.258	Med	session-rotation-revocation-unwired	Amber	Passkey + SIWE + OAuth. Same S-2 seam note as connect.
connect-client	Generic relying-party connect client	1.0.0-alpha.153	—	—	Green	
organization	Organization membership, enrollment, roles, revocation cascade (spec 324 W2)	0.0.0-alpha.9 6	—	—	Green	
related-agents	Private, holder-resident related-agent credentials (person↔org links as vault situation credentials, never on-chain edges — ADR-0025) + scop	0.0.0-alpha.213	—	—	Green	
situations	Substrate Situation primitive (SituationV2, RoleBinding, EntityRef) — spec 324 W2	0.0.0-alpha.9 1	—	—	Amber	

Authority (8 packages)

Package	What it does	Version	Tests	Open findings	Grade	Note
delegation	EIP-712 smart-account delegations with caveats	1.0.0-alpha.2522	Low	FAB-7	Green	The authority core: EIP-712 delegations, caveats, DEL-001 tokens, mandates. Hardened 2026-09-18: session leaf names its audience, strict per-step mode, plan digest, no bare vault wildcard. One low remains (FAB-7 defense-in-depth).
delegated-signer	Generic named delegated-signer resolution — compose agent-naming + agent-account + delegation + key-custody into a signer for a named identi	0.0.0-alpha.171	—	—	Amber	Named identity → delegate signer; root delegator must equal the named SA (anchor verified).
key-authorization	Policy-bound one-time key release (spec 277 §14)	0.0.0-alpha.9 2	—	—	Green	
agentic-authorization	Native Agentic authorization discovery	0.0.0-alpha.3 1	—	—	Amber	
entitlements	Durable resource/action/field/purpose/classification authorization over VC-compatible entitlement credentials (spec 277 §10)	0.0.0-alpha.4 2	—	—	Green	
tool-policy	Protocol-agnostic classification, risk tiers, and exact-call policy	1.0.0-alpha.258	—	—	Green	Risk tiers + exact-call policy; transport-agnostic.
vault-authority	Vault subject authority primitives	0.0.0-alpha.5 1	—	—	Amber	

Package	What it does	Version	Tests	Open findings	Grade	Note
admission	Agentic Primitives Edge admission contracts	0.0.0-alpha.6	5	—	Green	Fail-closed admission engine: all-DENY composition, audit as precondition, zero evaluators ⇒ DENY (verified).

Harness (agent runtime) (9 packages)

Package	What it does	Version	Tests	Open findings	Grade	Note
orchestration	The Ring-0 agentic orchestration core (ADR-0044)	0.0.0-alpha.4	14	—	Green	Ring-0 plan/execute loop. Exact tool set, fail-closed fan-out, resume integrity (P-4 closed 2026-09-18); every fail-open knob now registered.
orchestration-anthropic	Anthropic LLM binding for the @agenticprimitives/orchestration Planner port (the in-repo adapter, chain-state-viem pattern)	0.0.0-alpha.5	2	—	Green	
orchestration-openai-compat	OpenAI-compatible chat-completions binding for the @agenticprimitives/orchestration Planner / AnswerComposer ports (the in-repo adapter, cha	0.0.0-alpha.4	6	—	Green	
harness	The authority-aware agent harness (spec 350)	0.0.0-alpha.5	11	—	Green	Composition root: mandate verifier, risk ladder, receipts, playbooks by digest. planDigest binding landed 2026-09-18.
context	The two-tier context contract (spec 351 P0.11)	0.0.0-alpha.4	18	—	Green	Two-tier context: public KB vs private records; party resolution refuses to guess (no name heuristics — verified).
service-agent	Generic service-agent capability framework (spec 283)	0.0.0-alpha.1	22	—	Green	
runtime-member	ADMISSION of an existing general-purpose runtime as a workspace MEMBER, without a rewrite (spec 400 W1 — admission, not 'add coding')	0.0.0-alpha.2	3	—	Green	
acp	Agent Client Protocol (ACP) host binding	0.0.0-alpha.2	1	—	Amber	
surface-catalog	One SurfaceDescriptor per HTTP route / A2A skill / MCP tool, and the generators that derive A2A card skills, MCP tool declarations, OpenAPI	0.0.0-alpha.4	4	—	Green	

Edge and protocols (8 packages)

Package	What it does	Version	Tests	Open findings	Grade	Note
edge-runtime	Vendor-neutral edge ADMISSION layer	0.0.0-alpha.7	3	—	Green	

Package	What it does	Version	Tests	Open findings	Grade	Note
edge-cloudflare	Cloudflare adapter for @agenticprimitives/edge-runtime	0.0.0-alpha.7 2	—		Green	
rate-control	Vendor-neutral traffic-rate-limit + hard-usage-budget controls	0.0.0-alpha.4 1	—		Amber	
rate-control-cloudflare	Cloudflare adapter for @agenticprimitives/rate-control	0.0.0-alpha.4 1	—		Amber	
mcp-protocol	Conformant MCP protocol primitive (spec 293)	0.0.0-alpha.3 2	—		Green	Supports MCP 2026-07-28 and 2025-11-25.
mcp-oauth	MCP OAuth compatibility adapter (spec 277 §6-§8,§15)	0.0.0-alpha.7 2	—		Green	OAuth ingress adapter (envelope, never authority). One MCP revision behind: no CIMD, no RFC 9207 iss parameter (low).
mcp-runtime	Delegation-aware authorization middleware around the official MCP TypeScript SDK	1.0.0-alpha.255	—		Green	Delegation-aware MCP middleware; allow receipt now minted after scope is judged. Non-native token path binds no tool args (N-3, medium, long-standing).
a2a	Async, delegation-authorized Agent-to-Agent (A2A) task transport	0.0.0-alpha.2321	Med R917-A-6		Amber	Own A2A 1.0 implementation, official TCK green (JSON-RPC binding). One medium open: message bodies land in DO storage when no delegation is on the task (R917-A-6).

Registry and discovery (7 packages)

Package	What it does	Version	Tests	Open findings	Grade	Note
registry-kit	AP-native registry-building primitives	0.0.0-alpha.1310	Med S-3-conjunction		Amber	Registry-building primitives; admission receipts; evidence never score (ADR-0038).
registry-resolution	The composed typed resolver (spec 346 §6.3/§8/§9)	1.0.0-alpha.251	—		Green	
intent-resolver	Resolver layer skeleton (W1) — types + PassThroughResolver only	0.0.0-alpha.211	—		Amber	W1 skeleton (pass-through resolver only).
intent-marketplace	Direct Lane intent matchmaking	0.0.0-alpha.211	—		Amber	
intent-engagement	Intent engagement protocol core (spec 336 §8)	0.0.0-alpha.4 2	—		Green	
capability-claims	Capability claim credentials (SkillClaimCredential) + on-chain SkillDefinitionRegistry helpers (spec 251)	0.0.0-alpha.238	—		Green	
agent-skills	Deprecated alias — re-exports @agenticprimitives/capability-claims (ADR-0051)	0.0.0-alpha.230	—		Amber	Deprecated alias of capability-claims; retire.

Evidence and provenance (9 packages)

Package	What it does	Version	Tests	Open findings	Grade	Note
audit	Append-only audit event schema + sink interface	1.0.0-alpha.252	—	—	Green	
provenance	The Ring-0 provenance projector (spec 316 §6)	0.0.0-alpha.3	5	—	Green	
verification-receipts	Verifier-retained verification receipts (spec 303)	0.0.0-alpha.6	1	—	Amber	
witness	Evidentiary dispute attestation (spec 307 / G-11)	0.0.0-alpha.9	1	—	Amber	
attestations	AttestationRegistry SDK — EAS-aligned with bilateral consent	0.0.0-alpha.212	—	—	Green	
verifiable-credentials	W3C Verifiable Credentials envelope + Eip712Signature2026 proof + DOLCE+DnS Situation bases + ontology-shape schema registration	0.0.0-alpha.222	—	—	Green	
privacy-credentials	Privacy presentation layer (spec 305 / G-6)	0.0.0-alpha.6	1	—	Amber	
agreements	Commitment-only AgreementRegistry SDK	0.0.0-alpha.223	—	—	Green	
fulfillment	FulfillmentCase lifecycle + Task/Message/Artifact (re-exported from mcp-runtime/a2a) + HandoffPolicy + EvidenceCredential + OutcomeCredential	0.0.0-alpha.221	—	—	Amber	

Coordination and collaboration (4 packages)

Package	What it does	Version	Tests	Open findings	Grade	Note
coordination	Coordination core: Endeavor, OutcomeSpecification, CoordinationPlan, Participation — spec 332 W1	0.0.0-alpha.8	7	—	Green	Endeavor / CoordinationPlan core; cross-Home federation proofs (G4–G6) still unproven.
collaboration	Governed collaboration activities (spec 378 huddles)	0.0.0-alpha.4	1	—	Amber	
fabric	The Agentic Interaction Fabric (spec 316)	0.0.0-alpha.1448	High HOME-PORT-1 · Med FAB-1 · Med FAB-ARCH-1 · Low FAB-5	—	Amber	Operational owner of interactions; DO-side. Two mediums (FAB-1 receipt after transport admission; FAB-ARCH-1 role table) and HOME-PORT-1 (wires in Home KV) reference it.
home	Portable Agentic Trust Home contracts (spec 310)	0.0.0-alpha.144	—	—	Green	Portable Home contracts; storage-elimination program W6 done.

Ontology (1 packages)

Package	What it does	Version	Tests	Open findings	Grade	Note
ontology	Monorepo-wide formal vocabulary (RDFS/OWL T-box + SHACL/SKOS C-box + A-box fixtures) for agentic primitives	1.0.0-alpha.2513	—	—	Green	T-box/C-box with IRI-bound code and a build gate — the strongest single asset in the census.

Operations, data and developer kit (12 packages)

Package	What it does	Version	Tests	Open findings	Grade	Note
devkit	The Developer Kit core (spec 398 §10 / spec 399 §3)	0.0.0-alpha.6 5	—	—	Green	ap doctor / ap test / ap conform — the doctrine gates as CLIs; consumed by ap-home via agentic.lock.json.
create-app	The Developer Kit generator (spec 398 §10.3, spec 399 §3.6)	0.0.0-alpha.6 1	—	—	Amber	
evaluation	Truthfulness evaluation for agent answers (spec 358 W2)	0.0.0-alpha.3 3	—	—	Green	Deterministic truthfulness judge. Adversarial (injection / tool-misuse / memory) cases: none yet — the gap the 09-17 review named.
chain-state	The chain-read authority port	0.0.0-alpha.3 1	—	—	Amber	Fail-closed chain-read port; no eth_getLogs (ADR-0012).
chain-state-viem	viem-backed ChainProvider adapter for @agenticprimitives/chain-state	0.0.0-alpha.5 1	—	—	Amber	
ap-kms	Manifest-driven managed-KMS orchestrator	0.0.0-alpha.161	—	—	Amber	Managed-KMS orchestrator (AKCS pilot). Verify false-green anchor (EXT-KMS-1) holds.
vault	Agentic Delegated Data Vault seam (spec 277) — the Vault read/write/list interface + data-classification taxonomy + persisted object envelop	0.0.0-alpha.4 2	—	—	Green	Vault seam + classification taxonomy; the D1-backed implementation lives in the estate.
content-primitives	Verifiable Content Substrate SDK	1.0.0-alpha.265	—	—	Green	
content-storage	Managed storage for Content Artifacts (SKILL.md, .ttl, .md, JSON-LD, images/video)	1.0.0-alpha.4 8	—	—	Amber	Scaffold-stage storage port.
payments	PaymentMandate + ContextBinding + MandateConstraints + open/closed mode discrimination	0.0.0-alpha.217	—	—	Green	PaymentMandate + x402 v2 / wallet / sponsored-userop rails; allowUnsignedMandate defaults false (verified). Re-conformance to the x402 Foundation text pending.
geo-features	Off-chain geo CLAIM credentials + on-chain GeoFeatureRegistry helpers (spec 251)	0.0.0-alpha.211	—	—	Amber	

Package	What it does	Version	Tests	Open findings	Grade	Note
contracts	Solidity contracts + ABIs + flattened sources + per-network deployment addresses for the agenticprimitives stack	1.0.0-alpha.2567	2567	High OPS-KILLSWITCH-1 · Med EXT-DM-NONACCO UNT-1 · Low NEW-RB-1 · Low pay-enforcer-double-receipt	Amber	Solidity + ABIs + deployments. The open High here (OPS-KILLSWITCH-1) is that governance is not DEPLOYED, not a code defect. See §4.

4. Contracts

Fifty Solidity sources under `packages/contracts/src`. **939 / 939 tests** pass at HEAD. Thirty contracts (counting the account implementation behind the factory) are deployed on faithchain (and Base Sepolia); the rest are libraries, interfaces, abstract bases, a mock, and two governance contracts that are written and tested but **not deployed** — the single most consequential fact in this section.

Generation 2. The 18 September hardening changed six contracts (`AgentAccount`, `AgentAccountFactory`, `ApprovedHashRegistry`, `CustodyPolicy`, `DigestBindingEnforcer`, `ReceiptAnchorRegistry`). Those changes are graded here as fixed — the tests that reproduce the 17 September findings now pass — but every deployed environment declares `CONTRACTS_GENERATION = "1"`, and the runtime deliberately ships against the deployed generation. Their grade carries the mark **G2** to say: the code you would audit is not yet the code that is running.

Contract	Role	Deployed	Tests	Open findings	Grade	Note
<code>AgentAccount</code>	ERC-4337 / ERC-7579 modular smart account (UUPS behind ERC-1967)	impl. behind factory	242 refs	—	Amber G2	<code>custodyEpoch</code> now bound into <code>0x03 approved-hash sentinel</code> ; delegated calls refuse governance targets. Sessions + spend still inlined (module extraction pending — spec 351 <code>SessionKeyValidator</code>).
<code>AgentAccountFactory</code>	Deterministic (CREATE2) account deployment	yes	114	—	Amber G2	Network posture is a constructor fact (<code>relaxedT4Floor</code>), no longer a chain-id list; recorded in the deployments file; preflight refuses a mismatch.
<code>ApprovedHashRegistry</code>	Org batch-approval sentinel store	yes	22	—	Amber G2	Keys now (<code>hash</code> , <code>custodyEpoch</code>) — approvals die with the custody set that made them. Identical-call replay of quorum signatures (DM-2) remains a long-standing Medium at the call-hash layer.
<code>SmartAgentPaymaster</code>	ERC-4337 paymaster with per-agent budget	yes	50	—	Green	Deposit alert absent (estate, OPS-MONITOR-1).

Contract	Role	Deployed	Tests	Open findings	Grade	Note
UniversalSignatureValidator	ERC-1271 / 6492 / WebAuthn / P-256 validation	yes	4	—	Amber	Correct and fuzzed via callers; direct test count is thin. Cross-account signature reuse without address(this) / chainid binding on the ECDSA/WebAuthn paths (CA-2, Medium) is acknowledged.
IAgentAccount	Account interface	—	32	—	Green	
agency/DelegationManager	ERC-7710-shaped delegation redemption, on-chain revocation	yes	173	Med EXT-DM-NONACCO UNT-1	Amber	Not upgradeable (correct). Per-call rather than cumulative value caps (DM-1) and a generic redeem path that does not require the root delegator to be an account contract are the two known Mediums; both have designs, neither is scheduled before the fixed-commit audit.
agency/ICaveatEnforcer, IDelegationManager	Interfaces	—	—	—	Green	
agreement/AgreementRegistry	Commitment-only bilateral agreements	yes	41	—	Green	
attestation/AttestationRegistry	EAS-aligned attestations with bilateral consent	yes	41	—	Amber	Revocation is cosmetic against salt replay (ATT-2, Medium) — re-anchoring under a new UID is possible.
custody/CustodyPolicy	ERC-7579 executor: threshold, guardians, recovery, approvals	yes	376	—	Amber G2	onUninstall reverts StillInstalledOn while the account lists it; the account clears its flag before calling out. The largest and best-tested contract in the set.
custody/IERC7579Module	Interface	—	—	—	Green	
enforcers/AllowedMethodsEnforcer	Caveat: selector allow-list	yes	9	—	Green	
enforcers/AllowedTargetsEnforcer	Caveat: target allow-list	yes	11	—	Green	
enforcers/CallDataHashEnforcer	Caveat: exact calldata hash	yes	6	—	Green	
enforcers/CaveatEnforcerBase	Abstract base	—	—	—	Green	
enforcers/DigestBindingEnforcer	Caveat: intent-digest binding, single-use per step	yes	11	—	Amber G2	Now stateful and DelegationManager-only; consumed under (delegator, delegationHash, kind, stepNonce) . Generation 1 on chain is stateless.
enforcers/PaymentEnforcer	Caveat: payment limits + receipt	yes	23	Low pay-enforcer-double-receipt	Green	Two payment caveats on one chain each record a receipt (Low).
enforcers/QuorumEnforcer	Caveat: N-of-M co-signature	yes	24	—	Amber	Quorum typehash carries no nonce/ expiry (EN-13, Medium; same root as DM-2).

Contract	Role	Deployed	Tests	Open findings	Grade	Note
enforcers/ TimestampEnforcer	Caveat: validity window	yes	19	—	Green	
enforcers/ ValueEnforcer	Caveat: value cap	yes	14	—	Amber	Per-call, not cumulative (EN-22, Medium; same root as DM-1).
geo/ GeoFeatureRegistry	Public geo features (claims stay private)	yes	22	—	Green	
governance/ AgenticGovernance	Timelocked governance / guardian pause	no	13	High OPS-KILLSWITCH-1 (estate)	Red (undeployed)	Written, tested, not deployed on any chain. Privileged authority today is the disclosed deployer EOA. The single highest-leverage operations item.
governance/ GovernanceManaged	Abstract: governance-managed base	no	17	—	Red (undeployed)	As above.
governance/ IGovernance	Interface	—	—	—	Green	
identity/ AgentProfilePredicates	Library: profile predicates	—	24	—	Green	
identity/ AgentProfileResolver	On-chain profile / agent-type record	yes	7	—	Amber	Authority for <code>atl:agentType</code> ; thin direct tests.
libraries/ MultiSendCallOnly	Safe-pattern batch call (ported, no runtime dep)	—	2	—	Green	
libraries/ P256Verifier	P-256 precompile wrapper	—	3	—	Green	Fallback posture is a constructor fact of the factory (G2).
libraries/ SignatureSlotRecovery	Signature packing / slot recovery	—	24	—	Green	
libraries/ WebAuthnLib	WebAuthn assertion verification	—	48	—	Green	Symbolic proofs cover this path.
mocks/ MockUSDC	Test token	—	10	—	n/a	Never deployed to production; value pilots use it, which is why "real value" is NO-GO.
naming/ AgentNameAttributeResolver	Name attribute records	—	16	—	Green	
naming/ AgentNamePredicates	Library: typed-suffix predicates	—	51	—	Green	
naming/ AgentNameRegistry	Hierarchical typed naming, forced-unique	yes	54	—	Amber	Reclaim / owner-transfer does not clear resolver rows (NAMING-RECLAIM, Medium): a prior owner's forward+reverse binding persists until overwritten.
naming/ AgentNameUniversalResolver	Composed forward/reverse resolver	yes	3	—	Amber	Thin direct tests; exercised through the SDK.

Contract	Role	Deployed	Tests	Open findings	Grade	Note
naming/ PermissionlessSubregistry	Open subregistry (demo.agent-style)	yes	20	—	Amber	Front-running (SUB-1) and sybil/homoglyph squatting (SUB-2) are Mediums by design of "permissionless"; a commit-reveal is the named fix.
ontology/ AttributeStorage	Abstract: subject attribute store	—	12	—	Amber	No subject-ownership gate in the base (ONT-7, Low); concrete stores add their own.
ontology/ OntologyTermRegistry	On-chain term anchors	yes	30	—	Green	
ontology/ ShapeRegistry	SHACL shape anchors	yes	43	—	Amber	Trusts a caller-supplied store address (ONT-4, Low).
payments/ PaymentEscrow	Escrowed settlement	—	25	—	Amber	Tested, not deployed; not on any live rail.
payments/ PaymentReceiptRegistry	Settlement receipts	yes	10	—	Green	
provenance/ ReceiptAnchorRegistry	Anchor a receipt digest on chain	yes	7	—	Amber G2	Rows now keyed (anchoredBy, digest); a carried record's anchor is read as (agent, digest). Generation 1 keys by digest alone. Deployed 17 September.
registry/ AgentRegistryBase	Registry-kit base: SA-anchored entries, pluggable membership	yes	44	Low NEW-RB-1	Green	First-create-wins registry ids and caller-chosen entry ids (Lows) are documented squatting edges for permissionless registries.
registry/ IRegistryMembershipPolicy	Interface	—	6	—	Green	
relationships/ AgentRelationship	Public relationship edges	yes	23	—	Green	Private person↔org links are vault credentials, never these edges (ADR-0025).
relationships/ AgentRelationshipPredicates	Library	—	58	—	Green	
relationships/ RelationshipTypeRegistry	Relationship vocabulary anchors	yes	7	—	Green	
skills/ SkillDefinitionRegistry	Capability-definition anchors (ids only)	yes	18	—	Green	

Verification coverage, stated honestly. 939 unit and integration tests; three invariant suites; echidna nightly and medusa weekly on the delegation and custody paths; seven halmos symbolic proofs — all over onlySelf and WebAuthn paths, **none over the redemption path end to end.** Static analysis (slither, aderyn) runs in CI; aderyn is advisory. **No third-party audit of the current code exists.** The most recent external assessment (June 2026) returned No-Go for real funds, durable authority and PII, and its blocking items are tracked in the ledger; a fixed-commit audit is the mainnet closing condition, and the contract layer is now in a state to freeze for it once generation 2 is deployed.

5. The estate

The estate is the product deployment built from these packages: the Home a person signs into, the runtime that acts for them, the vault that holds their records, the edge that admits traffic, and the chain that holds authority. Four environments run today — faithnet and faithnet-b (faithchain), production (Base Sepolia), and the ap-home product mirror. Each piece is graded on **what is running**, not on what is in source.

Piece	What it is	Grade	Assessment
faithchain (chain 34348)	Private EVM chain; the authority substrate for every delegation, revocation and receipt anchor	Red	No governance, timelock or guardian deployed; the deployer EOA holds factory/governance authority; nothing can halt the system under governance. Contracts on it are generation 1. This is the top operations item.
Base Sepolia (84532)	Public testnet deployment of the same contracts	Amber	Same governance gap; useful as the public-testnet proof and the eventual public-L2 rehearsal.
Agent runtime (demo-a2a → agent-runtime)	The A2A service: interactions, harness runs, mandates, tasks, org agents	Amber	Authority plane verified per step (held under audit). Hardened 18 September: stewardship derived from the org's side, in-Worker marker not HTTP-borne, connector perimeter refuses private address space, self-authorized writes need the person's turn, plan-bound mandates. Custody master key from a local secret (ALLOW_LOCAL_MASTER_KEY=true) in every environment — Critical, unchanged since July.
Vault (demo-mcp)	The person's/org's encrypted record store; per-record delegation scope; the only RECORD store	Amber	Per-record scope enforced; allow receipts minted after scope is judged (fixed 18 Sep). Vault KEKs in managed KMS per person. Backup unscheduled ; KEK-version destruction is irreversible data loss; the key-binding table is the sole owner→key map — three carried findings (one Critical, two High) that are operations, not code. /mcp* public ingress is a labelled non-conformant interop demo, gateway assertion deliberately off there.
Home MCP (home-mcp)	Claude.ai as a client of the person: OAuth 2.1 AS + MCP server that reaches her agent under her own delegation	Amber	TOKEN_SECRET now fails closed. Every connected person's ask-as-me wire is sealed under one Worker secret and the delegate key is a raw secret rather than a managed-KMS signer (R917-A-9, Medium) — the Real-PII blocker on this piece.
Edge (demo-edge)	Admission boundary: HTTPS, application auth, canonical resolution, admission — before anything	Amber	Admission engine fail-closed (verified). workers_dev=true on faithnet exposes bypass URLs around the edge (Medium); ap-home has it off.
Chain RPC gateway	Read-only chain access for the estate	Green	Read-only; no eth_getLogs in product paths (ADR-0012).
Home (demo-sso-next → ap-home)	The person's control plane: sign-in, ceremonies, cards, vault views, operator view	Amber	Holds no privileged state by design (storage elimination W6 complete; records carry to a second Home with the anchor read from chain). Two Home KV allow-list items remain: member→org and relying-app wires are orphaned in KV (High, portability) and a cache prefix classified by name. SSO recognition cookie is JS-managed (Medium). Vercel-deployed and therefore outside the Cloudflare preflight scan (Medium).
Discovery (indexer · MCP · A2A · GraphDB)	The public knowledge base: only on-chain-derivable facts (ADR-0040)	Green	Indexer is the only writer; consumer services never write viewer-derived facts back; the claim "reading the KB reveals nothing reading the chain would not" holds by construction and is gate-checked. Admin SIWE gate is client-side only (Low).
Managed KMS (AKCS pilot)	Per-identity HSM signing keys and vault KEKs on a confidential VM	Amber	Correct architecture (keys are delegates, never the identity). One caller identity shared across two estates — the one open trust item the operators name themselves.
Operations	Alerting · backups · runbooks · on-call · SLOs	Red	No alerting; no paymaster-deposit alert; no post-deploy canary; RECORD-store backup unscheduled and its runbook names a KMS the estate no longer uses; SLOs stated as goals, not measured. The nightly live gates now fail the workflow when red (fixed 18 Sep).

Piece	What it is	Grade	Assessment
Release and supply chain	Publishing, provenance, pinning	Amber	CycloneDX SBOM on release; CI actions SHA-pinned (18 Sep); release gate now runs the full check set minus a dated, reviewed skip list. npm provenance is off because the repository is private; consumers cannot yet verify a tarball came from this source. Toolchain split (wrangler 3/4, two viem lineages) is a Medium.

Estate configuration truth table (the flags that decide posture; every row verified against the deployed configuration files on 17 September):

	faithnet	faithnet-b	production (Base Sepolia)	ap-home
Runtime signer / envelope KMS	managed (AKCS)	managed	GCP KMS	managed (AKCS)
Custody master key	local secret	local secret	local secret	local secret
Vault KEK	managed, per person	managed	GCP, per person	managed
Gateway assertion required	runtime yes · vault no (labelled demo)	yes	runtime yes · vault no	runtime yes · vault no
workers_dev edge bypass	yes	yes	yes	no
Demo personas enabled	yes	yes	yes	yes
Governance / timelock deployed	no	—	no	no
KMS caller identity	own	shared	n/a	shared with faithnet
Alerting · backup · on-call	none	none	none	none

6. Capability areas against the field

For each area the substrate covers: the products a reviewer will already know, what we ship, and where we stand. "Ahead" means a property they do not have and we enforce; "at par" means comparable; "behind" means they have it and we do not. Dates are as of September 2026. The field survey behind this table is maintained as a spec in the repository and refreshed each review.

6.1 Agent identity

Products in the area. Okta Agent SSO (GA Aug 2026) and Okta for AI Agents · Microsoft Entra Agent ID (GA Apr 2026) · Auth0 Agent as Principal · WorkOS, Descope, Hush, Cymphony · ERC-8004 identity registry (mainnet Jan 2026, 170k registrations) · Kite Agent Passport

What we ship. A Smart Agent (ERC-4337 account) per person, organization and service — the address is the canonical identifier; credentials (passkey, SIWE, hardware) rotate under a custody policy while the identity persists; typed hierarchical naming (.me .org .team .svc .workspace .treasury .registry); canonical profile with a signed A2A card projection; FedCM-first browser sign-in; a private, holder-resident record of person↔org relationships.

Where we stand. Ahead on the model, behind on the bridge. Every enterprise product models an agent as a directory object with a human owner and a short-lived token — identity that exists at the pleasure of the issuer. Ours survives the vendor, the runtime and the credential. What we lack is the binding from their world: no signed link from an Okta/Entra/Auth0 agent token to a Smart

Agent, so an enterprise cannot yet bring its directory agents to our authority model. ERC-8004 has scale we do not publish to; a projection is the named gap.

6.2 Delegated authority and policy enforcement

Products in the area. Auth0 nested `act` claims · AWS AgentCore Identity + Policy (Cedar, "verified by automated reasoning") · Google Agent Gateway (IAM + Model Armor on HTTP/MCP/A2A, `ext_authz` hook) · agentgateway (CNCF) + Cerbos / OPA · OpenID AuthZEN · MetaMask Smart Accounts Kit (ERC-7710 delegations, ERC-7715 permission requests) · Lit Vincent · Smart Sessions

What we ship. EIP-712 delegations with composable caveat enforcers (targets, methods, calldata hash, value, time, quorum, payment, intent digest); on-chain revocation checked per step; a mandate = delegation + intent digest, verified before every effect; the whole authority chain assembled and verified against chain state, not as claims about the past; a fail-closed admission engine; a protocol-agnostic risk ladder and exact-call tool policy.

Where we stand. Ahead on semantics; absent on the wire; behind on formal proof. Nothing in the field verifies an attenuation chain against live chain state per step — theirs is a token that was true when minted, or an ACL the platform holds. But their gateways cannot call our verifier: we publish no AuthZEN or `ext_authz` decision endpoint, so a Google or agentgateway deployment has no way to ask us "may this proceed?". Cedar's "formally verified" phrasing is now a reviewer's expectation; our enforcers are fuzzed and seven proofs exist, none over the redemption path end to end. Interop with MetaMask's 7710 objects is aligned by design and untested in practice; ERC-7715 requests are not implemented.

6.3 Custody, recovery and keys

Products in the area. Safe (multi-sig) · Privy, Turnkey, Fireblocks (custodial key infrastructure) · Coinbase AgentKit wallets · Kite (custodian in the loop) · Lit Protocol

What we ship. Our own custody module (ERC-7579 executor): threshold, guardians, timelocked recovery, custody epochs; credential rotation that never changes the identity or invalidates its delegations; a pluggable key-custody seam (local, AWS, GCP, our managed-KMS pilot); service keys that are delegates of a custodied identity, never the identity. No third-party multi-sig dependency.

Where we stand. At par on custody design, behind on operating maturity. The design is stronger than a wallet-provider model because recovery is governance, not a delegation, and a compromised service key is a revocable delegate. The estate does not yet live up to it: the custody master key is a local secret everywhere, and the managed-KMS pilot shares one caller identity across two estates. Fireblocks-grade operations — HSM-backed everything, key ceremonies, audited backups — are the gap, and they are operations, not architecture.

6.4 Agent runtime, harness and durable execution

Products in the area. LangGraph / LangSmith Deployment · Microsoft Agent Framework + Foundry (Durable Task) · Google ADK 2.0 · OpenAI Agents SDK (re-architected Apr 2026; Agent Builder retiring Nov 2026) · Anthropic Claude Agent SDK and Claude Managed Agents · AWS AgentCore Runtime · Temporal (\$12.55B, Sep 2026), Restate, DBOS, Orkes · Cloudflare Agents SDK · Agno, Strands, Mastra, CrewAI, Pydantic

What we ship. An authority-aware harness: Ask → Intent → Mandate → Plan → per-step Verify → Receipt; a Ring-0 orchestration loop with planner and tool-invoker as ports (Anthropic and OpenAI-compatible bindings; five providers budget-routed in the runtime); playbooks (SKILL.md packages) pinned by digest; checkpoints where a resume re-verifies authority; replay that never re-runs a tool; fan-out that is N mandates; a Durable Object / Workflows binding for durable steps; admission of existing runtimes (Claude Code, goose, Codex) as workspace members over ACP.

Where we stand. Ahead on "may this run?" and "what ran?"; at par on durable execution with one binding; behind on observability tooling and hosting. Every framework above has guardrails as policy code around a probabilistic loop; ours puts a contract between the plan and the effect. LangGraph's time-travel and LangSmith's traces, Foundry's durable tasks and OTel conventions are the reference points; we match checkpoint/resume and exceed them on what a resume must re-prove, but our durable-step port has a single (Cloudflare) adapter, we export no OpenTelemetry, and we offer no multi-tenant hosted runtime for third parties. The visual builder we deliberately did not build is the one OpenAI is retiring.

6.5 Edge, protocols and admission

Products in the area. MCP (revision 2026-07-28; auth hardening: CIMD, RFC 9207) · A2A 1.0 (Mar 2026; JSON-RPC, HTTP+JSON, gRPC) · Google Agent Gateway · agentgateway · Cloudflare AI Gateway · Kong / Apigee AI gateways · SPIFFE / mTLS meshes

What we ship. Our own A2A 1.0 implementation — official TCK green (MUST 88 / SHOULD 8 / MAY 4) on the JSON-RPC binding; MCP protocol primitive negotiating 2026-07-28 and 2025-11-25; delegation-aware MCP middleware where OAuth is an ingress envelope and Web3 is the authority; a vendor-neutral admission edge (HTTPS required, mTLS optional evidence, admission always) with a Cloudflare adapter; rate control per Smart Agent as a hard budget.

Where we stand. At par on A2A (JSON-RPC), one revision behind on MCP auth, ahead on what admission means. The field treats a gateway as a place to apply an ACL; ours is where transport evidence, application authentication and canonical identity resolution are ordered and nothing skips a stage. Gaps: the HTTP+JSON and gRPC A2A bindings are not served; the MCP OAuth adapter lacks Client ID Metadata Documents and the RFC 9207 issuer check; no third-party runtime has yet been recorded reaching one of our agent cards.

6.6 Registry, discovery and naming

Products in the area. ERC-8004 (170k registrations; 3-15 % with a live endpoint) · GoDaddy ANS · Hashgraph HCS registries · MCP Registry (12.7k components) · Project NANDA · AGNTCY · DNS-AID · ENS

What we ship. A registry kit, not a registry: Smart-Agent-anchored entries, pluggable membership and validation hooks, signed cards with binding proofs, admission receipts and a lifecycle log, evidence-only trust (never a score); typed contextual resolution (`alice@workspace`); a public discovery knowledge base holding only chain-derivable facts; an ARD / ACP-compatible read tier; private, recipient-bound resolution grants for agents that are unlisted.

Where we stand. Ahead on doctrine, absent on distribution. The registries that exist are mostly placeholders (ERC-8004's own numbers say so), and reputation-as-score is what we refused to build. But we publish to none of them: there is no projection of our entries to the 8004 identity registry or the MCP Registry, so an agent built here is findable only where we run a registry. Private resolution — "may this party discover how to reach this agent?" as a grant separate from "may it use it?" — has no peer we know of.

6.7 Evidence, provenance and observability

Products in the area. LangSmith · Langfuse · Arize · OpenTelemetry GenAI semantic conventions (still Development) · W3C PROV-O (standard) · PROV-AGENT (research) · Datadog / Honeycomb LLM observability

What we ship. Three layers: OpenTelemetry-shaped spans with a firewall, W3C Trace Context as the cross-agent spine, and PROV-O saying what the trace means — a span is a PROV activity; a signed receipt per act; verifier-retained verification receipts (commitment-split evidence that

signature, scope and revocation were checked at time T); receipt anchors on chain; an operator view over the evidence; SD-JWT selective disclosure and pairwise pseudonyms for presentations; dispute attestation by deterministic walk.

Where we stand. Ahead on evidence, no exporter, unmapped for compliance. The field records traces for debugging; ours are evidence a third party can verify against chain state. But we export nothing to the OTel ecosystem the field lives in, so a customer's Datadog cannot see our runs, and the artifacts a regulator wants (EU AI Act Art. 12 logging, Art. 14 oversight; ISO 42001; NIST AI RMF) are produced by construction with no written crosswalk that lets a compliance officer see it.

6.8 Payments and mandates

Products in the area. AP2 (in FIDO since Apr 2026; v0.2 Human-Not-Present) · Mastercard Verifiable Intent · x402 Foundation (LF, Jul 2026; 40 members incl. the card networks) · Stripe MPP · Visa TAP · OpenAI/Stripe ACP · Virtuals ACP · Skyfire, Nevermined, Crossmint · MetaMask Smart Accounts Kit paying x402 via 7710

What we ship. A payment mandate bound to its context; the intent digest as a delegation caveat with a single-use on-chain `PaymentEnforcer`; three rails (x402 v2, direct wallet, sponsored UserOp); settlement receipts on chain; a paymaster with per-agent budget; the two-flip offer algebra that treats listings, carts and auctions as one shape.

Where we stand. Same idea, no projection; at par on x402. AP2's Intent/Cart/Payment mandates and Mastercard's Verifiable Intent are the artifact we call the intent-digest caveat — and we have written no adapter that emits or accepts theirs, so a merchant on AP2 cannot honour our mandate. x402 v2 is implemented; conformance against the Foundation's text has not been re-checked since the transfer. Card rails (MPP, TAP) are not ours to build; the mandate must project onto them. Real-value pilots run on mock USDC.

6.9 Coordination between agents

Products in the area. Microsoft Agent Framework handoffs · LangGraph multi-agent · Google A2A task model · Salesforce Agentforce · Virtuals ACP (commerce lifecycle) · Buzz.xyz · Slack / Notion / Lindy agent workflows

What we ship. Two planes kept apart: **coordination** between agents (durable Endeavors, shared CoordinationPlans adopted by decision, participants resolved by capability through discovery, signed contribution commitments against an exact plan revision) and **orchestration** within one agent (the harness). Interactions as a governed root aggregate with immutable signed exchanges; huddles as governed collaboration activities; workspace and club contexts with standing grants.

Where we stand. Ahead in kind, unproven across Homes. Every framework's "handoff" is a function call between agents the same operator runs; ours is a signed commitment between principals who may not trust each other's runtime, and neither plane is authority. The claim that this works across two independently operated Homes — the federation proofs (G4-G6) — is written down and not yet demonstrated; that is the gap. Buzz-style approvals are a click; ours is a signature — a differentiation we have, and a UX cost we carry.

6.10 Knowledge, context and memory

Products in the area. Claude Managed Agents (versioned memory, per-user credential vaults) · AWS AgentCore Memory · Mem0, Zep, Letta · LangGraph stores · Notion / Glean enterprise context · Inrupt Solid pods

What we ship. Two tiers that are never joined in an engine: a public knowledge base (RDF, only chain-derivable facts) and a person's private vault (encrypted, per-record delegation scope,

questions compile to selectors evaluated in the store). Memory as vault-resident records the person owns; connectors (any MCP server) admitted under her grant, never the platform's; documents as evidence her agent reads. A truthfulness evaluation set seeded from live incidents.

Where we stand. Ahead on ownership and portability, at par on capability, behind on adversarial evaluation. Every hosted memory product is a platform store; ours travels with the person to another Home and is scoped per record. Self-authorized writes now require the person's turn (18 Sep) — the prompt-only perimeter the review found. What we lack is an external adversarial corpus (injection, tool poisoning, memory poisoning — the OWASP ASI categories) with "unauthorized effect = 0" as the published metric; our evaluation set is truthfulness-only.

6.11 Ontology and semantic grounding

Products in the area. Palantir Foundry ontology · Schema.org / OASF · W3C PROV-O, VC, DOLCE · enterprise knowledge-graph vendors

What we ship. One formal vocabulary (RDFS/OWL T-box, SHACL/SKOS C-box) grounding identity, authority, interaction, coordination, execution and provenance; code binds to it by IRI and a build gate fails when it does not; domain relationships live in the ontology, not in prompts or tables; on-chain term and shape anchors.

Where we stand. Ahead; no direct peer at the substrate level. Palantir's Action Types are the closest analogue and are a product feature; ours is the thing the products are compiled from — one capability model generates the UX affordance, the Ask's vocabulary, the A2A card and each act's mandate requirement. The risk is the opposite of the field's: sophistication that a stranger must learn before contributing.

6.12 Security, operations and compliance posture

Products in the area. OWASP Top 10 for Agentic Applications (2026) · NIST AI 100-2 / AI RMF · EU AI Act (Art. 50 in force Aug 2026; Annex III Dec 2027) · ISO/IEC 42001 · SOC 2 · agent-security vendors (Zenity, Straiker, HiddenLayer, AIR)

What we ship. A findings ledger machine-checked for freshness on every CI run; 131 doctrine gates; 57 nightly live gates with negative twins that now fail the build; forge/echidna/medusa/halmos on contracts; a consolidated OWASP ASI01-10 map with an enforcement point per category; red-team scope and evidence checklist; a written methodology; this document.

Where we stand. Ahead on engineering culture, behind on external proof. Few projects at this stage have a ledger a stranger can audit or nightly gates that assert the negative case. But there is no third-party audit of the current code, no funded bounty, no SOC 2 for the hosted estate, no ISO 42001 / NIST mapping, no alerting, and governance is not deployed on the chain that holds authority. Security truth is now current (this review); posture is the gap.

6.13 Developer experience and ecosystem

Products in the area. Vercel AI SDK · LangChain ecosystem · shadcn-style registries · Cursor / Copilot / Devin (build tooling) · npm provenance / SLISA

What we ship. A Developer Kit: `create-app` writing a product repository, `agentic.lock.json`, `ap doctor` / `ap test` / `ap conform` as the doctrine gates in CLI form, rules projected by digest, a read-only Developer MCP; 210 specs and 65 ADRs as institutional memory; an `/llms.txt` and starter.

Where we stand. Behind, and honestly so. Every consumer is first-party; the repository is private; there are no download statistics, no external contributors, no public issue tracker, no published npm provenance. The kit has exited a clean machine once, measured across one

developer. This is the dimension a diligence team writes as a closing condition for any thesis that depends on ecosystem, and no amount of architecture substitutes for it.

7. Open findings register

Every Critical and High, and the Mediums by class. Identifiers are the ledger's; a stranger with repository access can look each one up. Nothing here is an exploit recipe; the ledger entries that would be are held with the code.

7.1 Critical

Id	Title	Class	Status
D-P0-1	Rotating the runtime master key silently orphans every OIDC-custodied Smart Agent	key management / identity continuity	open — needs a documented rotation-with-re-association ceremony; operations
NEW-C2	Master signing key and session-envelope key run from local secrets in production	key management	open, narrowed to the custody master (runtime signer and vault KEKs already in managed KMS); the "real value" blocker
VCD-EXTERNAL	Tracker for a first-party consumer repository's own findings	external	tracked in that repository

7.2 High

Id	Title	Class	Status
D-P0-2	Destroying a per-person vault key version is permanent, irreversible data loss	backup / key lifecycle	open; needs tested recovery + scheduled backup
D-P1-3	The vault key-binding table is the sole owner→key map	backup / single point	open; same wave
HOME-PORT-1	Member→org and relying-app delegation wires orphaned in Home KV	portability	open; the last Home storage-elimination item
NEW-H3	Channel/directory read gate does not re-check member-access revocation on chain	authority (fabric read path)	open; design named
OPS-KILLSWITCH-1	No on-chain governance / kill-switch deployed	operations / governance	open; the highest-leverage item

Nine further Highs raised on 17 September — two contract, three delegation-core, one stewardship, three harness-perimeter — were **closed in source on 18 September** with tests. The two contract Highs await generation-2 deployment.

7.3 Medium, by class (30)

Class	Count	Representative items
Contract semantics (long-standing, designed, unscheduled before the fixed-commit audit)	9	per-call not cumulative value caps (DM-1 / EN-22); no nonce on quorum/approved-hash signatures (DM-2 / EN-13); 1271 paths without account/chain binding (CA-2); attestation revocation vs salt replay (ATT-2); subregistry front-running and squatting (SUB-1/2); stale resolver rows after reclaim; generic redeem path not requiring an account root

Class	Count	Representative items
Estate configuration and operations	9	<code>workers_dev</code> edge bypass; no monitoring; runbooks absent; no crypto-shred path; RECORD-store backup unscheduled; preflight does not scan the Vercel-deployed Home; Home MCP wires under one Worker secret; toolchain split; A2A bodies DO-local without a delegation
Authority seams awaiting a consumer	5	session-rotation revocation seam unwired; non-native MCP token path binds no tool arguments; treasury dispatch not binding delegator == served treasury; fabric receipt signed after transport-only admission; fabric roles as an ACL table
Product truthfulness and portability	4	an org's own agent cannot read its membership; Home KV cache classified by prefix; SSO recognition cookie JS-managed; <code>cardΛproofΛregistration</code> not conjoined
Release gate	1	release gate's dated skip list (now enforced with review dates)
Other	2	

Sixteen Lows are tracked in the ledger and are not enumerated here.

7.4 Remediation order

Wave	Content	State
R1 — authority plane	stewardship derived from the org's side; strict per-step mode; session leaf audience; plan-bound mandates; person's turn for self-authorized writes; <code>readOnlyHint</code> never decides; connector perimeter; in-Worker marker; <code>TOKEN_SECRET</code> fail-closed	done in source, 18 Sep
R2 — contracts	custody-tier boundary; custody-epoch approvals; single-use digest binding; constructor-fact network posture; anchors by anchorer	done in source (generation 2); not deployed
R3 — gates that tell the truth	global coverage floor; release gate skip list with review dates; nightly fails the build; every fail-open knob registered; API-surface snapshots over every subpath; SHA-pinned actions	done, 18 Sep
R4 — operations and posture	deploy generation 2 with governance; retire the local custody master; per-estate KMS caller identity; scheduled backup + tested recovery; alerting; <code>workers_dev</code> off; crypto-shred; Home MCP signer to managed KMS; SLO ledger; fixed-commit third-party audit	open — the verdict's remaining distance

8. Method, and how to verify

Who. The maintainers, with five parallel read-only review passes (contracts; authority packages; estate services; harness and knowledge plane; architecture, operations and supply chain) and a lead verification of every new High at its source line. No finding was accepted from a summary; each names a file and a symbol in the ledger.

What was read. All 77 packages and both domain packages; all 50 contracts; the deployed configuration of every environment; the product mirror; the CI workflows and the gates they run; the previous review (12 July 2026) and every open ledger item in scope, each re-verified at HEAD.

What was run. `forge test` (939/939); the audit-freshness gate (every open finding's anchor present in source); the production-deploy preflight (fails closed on the three named posture items); the full `check:all` set; the A2A TCK; and the 57 nightly live gates with negative twins.

What this is not. It is not a third-party audit. It is not a penetration test of the live estate. It is not a proof over the redemption path. It records those absences as findings rather than omitting them.

How a stranger verifies it. The findings ledger (`docs/audits/findings.yaml`) is the record; each entry names its concern paths, its anchor symbol and its enforcement level (`code-present`, `test-covered`, `production-enforced`, `deployed`). `pnpm check:audit-freshness` fails when an anchor is gone. The full internal review and its 18 September addendum, with file-and-line evidence, are in `docs/`

audits/. Deployed contract addresses and the contract generation each environment runs are in `packages/contracts/deployments-*.json` and the runtime configuration. The next full review is due at the close of Wave R4 or 17 November 2026, whichever is first.

Agentic Primitives is a project of Agentic Trust Labs. This assessment is published as a statement of where the substrate stands, including where it falls short. Corrections are welcome; the ledger is the place they land.