SOLARA ASSOCIATES · COMPARISON BRIEF

SYNTHERA VS GOOGLE · IBM · DATABRICKS

Everyone secures the agent. Almost no one secures the action.

Each platform below is genuinely strong at what it was built for. We state that plainly, in its own terms, before we state the one gap that VAID and SYNTHERA exist to fill.

Version 1.0 Date 6 July 2026 Owner Solara Associates

Claims last verified 2026-07-03 Recheck Quarterly

Naming competitors is a credibility signal, not a marketing risk — if it's done fairly. Every claim here is grounded in each vendor's own public materials and dated, so it can be re-checked on a cadence (see Claims & sources). Non-disparagement is deliberate: each comparison states a documented capability as of the date checked, not a permanent limitation.

01 · GOOGLE

Agent Identity in the Gemini Enterprise Agent Platform

WHAT IT DOES WELL

Perimeter governance, and it's real and partly shipping. An Agent Gateway applies IAP/IAM authorization over a SPIFFE-based Agent Identity floor — short-lived X.509 SVIDs, enforced with mTLS and DPoP — that cryptographically confirms an agent is who it claims to be within its Google Cloud trust domain. It authenticates and authorizes the agent in your environment, and logs its actions to your audit trail.

THE GAP THAT REMAINS

It answers "is this agent authenticated and authorized inside my environment, right now." It does not produce a portable, per-action proof of authorship that a party outside your trust domain can verify independently, without a prior federation agreement. Cross-organisation still means federating trust domains — SPIFFE bundle exchange, or an admin-configured trusted-trust-domain setting — negotiated per pair.

Worth noting fairly: A2A's Signed Agent Cards prove an agent's identity at discovery, not per action; and AP2 offers portable signing, but scoped to payments. Neither is general per-action attestation.

02 · IBM

Verify + Vault for agentic identity

WHAT IT DOES WELL

Real, necessary capabilities, aimed at a gap IBM itself names — the "accountability gap" in how enterprises govern agents. IBM registers agents with unique identities, replaces standing credentials with short-lived scoped tokens, and ties agent actions back to the human who authorized them, with signed audit trails. Verify anchors the human side, Vault (formerly HashiCorp Vault) the workload side, on open standards — OAuth, OIDC, token exchange, SPIFFE — with IBM Verify as the issuing authority.

THE GAP THAT REMAINS

IBM markets verifiable credentials that let an outside party confirm an agent's identity and authority are genuine. That is credential verification, not action verification. As of mid-2026, IBM's shipped agentic offering does not give a third party a way to independently confirm that a specific action was actually performed by the agent it claims to be, without validating tokens IBM Verify issued or federating with IBM Verify. The accountability model is scoped tokens plus signed audit trails anchored to IBM Verify — centralised and standards-based, not a portable per-action proof any external verifier can check on its own.

03 · DATABRICKS

Unity AI Gateway

WHAT IT DOES WELL

A genuinely strong runtime governance layer (the evolution of Mosaic AI Gateway, in Beta as of July 2026). Built on Unity Catalog, it unifies policy, spend control and observability for models, agents and MCP tools inside the Databricks estate, with an open partner ecosystem bringing in identity providers (Okta, Ping, SailPoint, Saviynt) and security vendors for defence in depth. Its in-estate audit is real and strong — inference tables, Unity Catalog logs, and a namespace where governance travels across a customer's accounts, regions and clouds.

THE GAP THAT REMAINS

It answers "what is this agent allowed to do, and how much is it costing me," governed centrally inside one platform's runtime path — and that governance spans a customer's whole Databricks footprint. What it doesn't answer is the cross-boundary question: when an action needs to be verified by a party that isn't running inside Unity Catalog at all — a customer's system, a partner's agent, an auditor with no Databricks access — there's no portable, independently-checkable proof that travels with the action itself.

04 · THE SHARED GAP, STATED ONCE

Authenticated inside a platform is not the same as verifiable anywhere.

Every one of these answers a version of "is this actor authenticated and authorized inside my platform, right now," and logs it to an in-platform audit trail. All three are good at that. None of them produces a portable, independently verifiable proof that travels with a **specific action** — checkable by anyone, with no shared trust domain, no federation agreement, and no integration into the issuing platform's stack.

> **That's the layer VAID and SYNTHERA occupy: above workload identity, above platform-bound policy engines, at the level of the individual signed action.**

05 · CLAIMS & SOURCES

Grounded, dated, and built to be re-checked

These vendors ship fast, so this brief has a shelf life. Every claim above is grounded in each vendor's own public materials and isolated here for re-checking on a cadence.

Last verified 2026-07-03 Owner Solara Associates Cadence Quarterly

GOOGLE

Agent Identity is SPIFFE-based; the Agent Gateway enforces IAP/IAM authorization; SVIDs are short-lived X.509, enforced via mTLS/DPoP, scoped to a Google Cloud trust domain.

Google Cloud docs — Govern / Agent Gateway / Agent Identity · checked 2026-06-30

GOOGLE

Cross-domain verification requires SPIFFE federation (bundle exchange), configured per pair of trust domains.

SPIFFE Federation spec; SPIRE federation docs · checked 2026-06

IBM

IBM markets short-lived scoped tokens replacing standing credentials and audit trails tying agent actions to the authorizing human; IBM itself names the "accountability gap".

ibm.com/solutions/agentic-ai-identity-management; IBM Think Insights, "The accountability gap in autonomous AI" · checked 2026-06

IBM

IBM Verify verifiable-credential marketing covers verifying an agent's identity/authority ("any verifier can independently confirm … the agent's credentials"), not per-action verification.

ibm.com/solutions/agentic-ai-identity-management · checked 2026-06

DATABRICKS

Unity AI Gateway (evolution of Mosaic AI Gateway) was unveiled at Data + AI Summit June 2026 and is in Beta per Databricks' own docs as of July 2026; identity partners include Okta, Ping, SailPoint, Saviynt.

databricks.com blog (2026-06-16/17); docs.databricks.com/aws/en/ai-gateway (updated 2026-07-01) · checked 2026-07-01

Comparison Brief · SYNTHERA vs Google, IBM, Databricks.

Non-disparagement is deliberate: each comparison states a documented capability as of the date checked, not a permanent limitation. If a vendor ships portable per-action verification, this brief changes.