PROCUREMENT COMPARISON

Fonteum vs Definitive Healthcare

A 12-dimension procurement comparison for AI-buyer and data-team buying decisions.
Version v1 · Q2 2026 · published 2026-05-29.

CONTENTS

EXECUTIVE SUMMARY

The decisive differences are not features

Definitive Healthcare, IQVIA OneKey, and HealthVerity are mature commercial platforms built on proprietary aggregation and claims data. Fonteum is built on primary federal sources cross-resolved on a shared identity backbone, with a field-level provenance contract on every record. For a team wiring provider data into agents and pipelines, the decisive differences are provenance, attestation, reproducibility, public pricing, and time-to-first-result.

PROCUREMENT CHECKLIST

- Does every field cite a source row, snapshot date, and digest?
- Are historical snapshots immutable and citable?
- Is the scoring methodology pinned to a published version?
- Are commercial-use rights and embargo flags stated up front?
- Is there a FHIR-conformant API with public docs and a sandbox key?
- Is there an MCP server and agent card for agent-native access?
- Can a third party reproduce a published claim?
- Is pricing public, and how fast is the first successful query?

DIMENSION 1 OF 12

Provenance contract

Can every field cite its source row and snapshot?

FONTEUM

Each displayed fact is written to the provider_field_provenance layer with its source name, the date the pipeline last reconciled it, and any known limitation. A compliance auditor can trace any rendered value back to the originating federal record — a CMS dataset ID, an OIG HHS file, or an HRSA portal URL — rather than to an opaque aggregation step. The contract is a fixed 14-tuple shape that travels with the record into exports and the FHIR API.

COMPETITORS

Definitive Healthcare publishes a high-level description of its data model and sourcing on its public site, but does not surface a per-field source citation in exports. IQVIA OneKey and HealthVerity similarly treat linkage and resolution as proprietary; none of the three publicly documents a field-level provenance contract.

Source: https://www.definitivehc.com/resources

WHY IT MATTERS FOR AI BUYERS

An AI agent or data team building on provider data needs to know which federal record backs each value before it acts on it. Without a field-level contract, a wrong or stale field is indistinguishable from a correct one, and there is no audit trail to defend a downstream decision.

DIMENSION 2 OF 12

Snapshot immutability

Are historical snapshots retained, addressable, and citable?

FONTEUM

Every published snapshot is stored as an immutable artifact identified by its SHA-256 digest, with Ed25519 witness co-signatures recorded in snapshot_witness_signatures. A claim cited from a 2026-Q1 snapshot resolves to exactly the bytes that backed it, even after the live data has moved on.

COMPETITORS

The incumbent platforms refresh their data on rolling cadences; none publicly documents an immutable, content-addressed snapshot that a third party can pin a citation to. Historical point-in-time access, where offered, is a contract feature rather than an addressable public artifact.

Source: https://www.definitivehc.com/resources

WHY IT MATTERS FOR AI BUYERS

Reproducible analysis and defensible audit trails require that a figure published today still resolves to the same underlying state next year. Mutable-in-place data cannot support a citation that survives the next refresh.

DIMENSION 3 OF 12

Methodology versioning

Is the scoring methodology pinned to a published version?

FONTEUM

Methodology versions are pinned string constants, never derived at runtime, and each has a published page describing inputs, transforms, and known limits. When a methodology changes, the version bumps and the prior version stays addressable, so a consumer always knows exactly which logic produced a score.

COMPETITORS

Definitive Healthcare, IQVIA OneKey, and HealthVerity each treat their modeling and resolution methodology as proprietary intellectual property. Public documentation describes capabilities at a marketing level but does not pin a citable methodology version that a buyer can reference in their own audit.

Source: https://www.definitivehc.com/resources

WHY IT MATTERS FOR AI BUYERS

A model that consumes a vendor score needs to know whether the scoring logic changed between runs. Unversioned methodology turns a silent vendor change into an unexplained shift in your own outputs.

DIMENSION 4 OF 12

Federal source coverage

How many primary federal sources are cross-resolved?

FONTEUM

Fonteum ingests 23 federal source families straight from their government portals — CMS NPPES, PECOS, Care Compare (eight facility modules), PBJ staffing, SNF All Owners, OIG LEIE, HCRIS, Open Payments, QPP MIPS, HRSA HPSA and UDS, BLS, BEA, and Census — and cross-resolves them on the NPI and CCN identity backbone. 2,703,357 rows are ingested with each family documented at /sources by tier, cadence, and redistribution posture.

COMPETITORS

Definitive Healthcare assembles facility, technology-install, and contact intelligence from claims and proprietary sourcing layered over public data; the public framing is coverage of facilities and executives, not a count of cross-resolved primary federal sources. IQVIA OneKey and HealthVerity are built on licensed reference and claims universes rather than a federal-source-first model.

Source: https://www.definitivehc.com/resources

WHY IT MATTERS FOR AI BUYERS

Primary-source coverage is what lets a buyer reason about freshness, gaps, and legal posture per source. A blended proprietary universe is convenient but hides which signal came from where.

DIMENSION 5 OF 12

Update cadence

How often is the data refreshed, and is the date on the field?

FONTEUM

Cadence is a property of each source family, not a platform-wide promise. NPPES refreshes weekly, OIG LEIE monthly, Care Compare and PBJ quarterly, and so on — and every rendered field carries the date the pipeline last reconciled it, so staleness is visible rather than assumed.

COMPETITORS

The incumbent platforms cite regular refresh cycles in marketing material, but do not surface a per-field last-checked date in their delivered data. A consumer cannot tell from a record alone how old a specific value is.

Source: https://www.definitivehc.com/resources

WHY IT MATTERS FOR AI BUYERS

Provider data ages unevenly — an exclusion flag matters the day it lands, while a taxonomy code rarely moves. A per-field date lets an agent weight freshness instead of treating the whole record as one age.

DIMENSION 6 OF 12

Cryptographic attestation

SLSA-style build provenance, SHA-256, signed artifacts?

FONTEUM

Each snapshot is digested with SHA-256 and co-signed with Ed25519 witness signatures recorded in snapshot_witness_signatures (public read). Exports and API responses can carry integrity headers (for example X-Fonteum-SHA256) so a consumer can confirm the bytes they received match the attested artifact. A public chain key is published at /.well-known/chain-public-key.

COMPETITORS

None of Definitive Healthcare, IQVIA OneKey, or HealthVerity publicly documents cryptographic attestation, content digests, or signed data artifacts. Integrity, where addressed, is handled at the transport and access-control layer rather than the artifact layer.

Source: https://www.definitivehc.com/resources

WHY IT MATTERS FOR AI BUYERS

An autonomous agent acting on data it did not fetch itself needs a way to confirm the payload was not altered in transit or substituted. Artifact-level attestation is the only way to close that gap without trusting every hop.

DIMENSION 7 OF 12

License clarity

Are commercial-use rights and embargo flags stated?

FONTEUM

The federal records Fonteum redistributes are US Government works and are not copyrightable (17 U.S.C. § 105), so the structured, provenance-tagged versions are openly redistributable. Each source family states its redistribution posture and any embargo at /sources, so a buyer knows the commercial-use position of every field before relying on it.

COMPETITORS

Definitive Healthcare, IQVIA OneKey, and HealthVerity are licensed commercial products; their data is governed by per-contract terms that are negotiated rather than publicly posted. Commercial-use rights and any redistribution limits are determined in the agreement, not surfaced on the record.

Source: https://www.definitivehc.com/resources

WHY IT MATTERS FOR AI BUYERS

A team shipping a product on top of provider data needs unambiguous commercial-use rights. Per-contract opacity means legal review on every new use case, where a public-domain base plus stated flags is decidable up front.

DIMENSION 8 OF 12

API quality

Stripe-grade docs, OpenAPI, SDKs, sane rate limits?

FONTEUM

The API implements HL7 FHIR R4 US Core 6.1.0 with five USCDI v3 Provider resources, a CapabilityStatement at /api/fhir/metadata, SMART Backend Services auth, and HL7 Bulk Data ($export). Docs follow a three-column Stripe-style layout, and an unauthenticated sandbox key (pk_dx_sample, 100 requests/hour) lets a developer call a real federal record before talking to anyone.

COMPETITORS

Definitive Healthcare offers a REST API to enterprise customers; its public documentation is gated behind sales and it is not FHIR-conformant. IQVIA OneKey delivers data through integration partners and licensed connectors; HealthVerity delivers via its own pipelines. None publicly posts an open OpenAPI specification or a free sandbox key.

Source: https://www.definitivehc.com/resources

WHY IT MATTERS FOR AI BUYERS

API quality is the difference between an afternoon integration and a quarter-long one. For EHR-vendor pipelines specifically, FHIR conformance and a discoverable CapabilityStatement are table stakes that a proprietary REST API does not meet.

DIMENSION 9 OF 12

Agent / MCP support

First-class MCP server and pre-built agent integrations?

FONTEUM

Fonteum publishes an MCP server descriptor at /.well-known/mcp.json and an agent card at /.well-known/agent.json with a full skills inventory, so Google ADK, LangGraph, and BeeAI consumers can discover and call it. The /for/ai-agents surface documents the agent-facing integration, and the FHIR layer is reachable by tool-using models directly.

COMPETITORS

Definitive Healthcare, IQVIA OneKey, and HealthVerity do not publicly document an MCP server or an agent-card skills inventory. Their integration model is human-operated dashboards and enterprise connectors rather than agent-native discovery.

Source: https://www.definitivehc.com/resources

WHY IT MATTERS FOR AI BUYERS

AI-buyer infrastructure is being assembled by agents, not just analysts. A platform with no MCP descriptor and no agent card is invisible to the multi-agent frameworks that are doing the buying.

DIMENSION 10 OF 12

Reproducibility

Can a third party recreate a published claim?

FONTEUM

Every research study at /research ships the underlying CSV and JSON, a methodology page, and the snapshot identifiers that backed the figures. Because the inputs are federal public records and the methodology version is pinned, a third party can pull the same source files and recreate a published number independently.

COMPETITORS

Definitive Healthcare, IQVIA OneKey, and HealthVerity publish findings and reports derived from proprietary data and models. Because the inputs and methodology are not open, a third party cannot independently recreate a published figure; the result must be taken on trust in the vendor.

Source: https://www.definitivehc.com/resources

WHY IT MATTERS FOR AI BUYERS

Reproducibility is the difference between a citable fact and a vendor assertion. For research, regulatory, and diligence work, a number that cannot be recreated cannot be defended.

DIMENSION 11 OF 12

Pricing transparency

Public pricing, or opaque enterprise sales?

FONTEUM

All published research snapshots and datasets are free to access and cite with attribution — no account, no key for the static files. The paid pilot tier is publicly posted from $2,000/mo at /pricing and adds custom export scoping, production API access, and methodology-versioning commitments, with a 30-day no-penalty exit.

COMPETITORS

Definitive Healthcare and comparable enterprise platforms route buyers through a sales process; pricing is by quote and not posted publicly. Independent reviews and procurement write-ups commonly describe five- to six-figure annual licenses, but the vendor sets the figure per account.

Source: https://www.definitivehc.com

WHY IT MATTERS FOR AI BUYERS

Public pricing lets a team size a build before committing to a sales cycle. Opaque enterprise pricing front-loads weeks of procurement before the data can even be evaluated.

DIMENSION 12 OF 12

Time-to-first-result

From signup to the first successful query.

FONTEUM

The free research datasets need no account, and the unauthenticated sandbox key (pk_dx_sample) returns a real FHIR record on the first request, so a developer reaches a successful query within minutes of arriving at /api. Production access is a pilot conversation, but evaluation is immediate.

COMPETITORS

For Definitive Healthcare, IQVIA OneKey, and HealthVerity, the first successful query follows a sales process, a signed agreement, and onboarding. Evaluation access, where offered, is a scheduled demo rather than a self-serve key.

Source: https://www.definitivehc.com

WHY IT MATTERS FOR AI BUYERS

Time-to-first-result is the single best proxy for how a platform treats builders. Minutes-to-query means a team can prove value before procurement; weeks-to-query means the opposite.