



ARCHITECTURE BRIEF · OCTOBER 2026

# Healthcare Quality Control System

From retrospective analytics to real-time outcome control for HEDIS, Star Ratings and CMS reporting.

DATA

POLICY

AI

ACTION

STARS

Nartis is a healthcare quality platform designed and prototyped by Adam Fiejdasz. It evaluates quality measures continuously, ranks open gaps by value, routes interventions to people, and keeps audit-grade lineage from data source to CMS-facing outcome. This brief describes the design and a working prototype that runs on synthetic data.

## THE CENTRAL SHIFT

# Healthcare quality needs a control system, not another dashboard.

**Legacy pattern**

Data is extracted, transformed, warehoused and visualized, then handed to people to interpret and act on. Gaps surface days or weeks after the visit that could have closed them.

**Nartis pattern**

Each clinical event triggers policy evaluation, gap detection, ranked action and outcome measurement. People approve clinical steps; the system keeps the evidence.

## CLOSED-LOOP MODEL

Observe → Decide → Act → Measure → Learn

The platform connects evidence, policy, prediction, intervention and audit in one operating model. Measure logic lives outside the pipelines as versioned, tested policy-as-code. Models rank and explain; they never decide whether a member meets a measure. Every request carries the identity chain of user, agent, purpose and data class.

**Event-driven**

Measures re-evaluate when a lab, fill or visit arrives, not at the next batch.

**Closed-loop**

Actions and their outcomes flow back so prioritization improves.

**Audit-ready**

Every decision is traceable to the rule version and the evidence used.


# Reference architecture

## LEGACY ANALYTICS VS NARTIS

| LEGACY ANALYTICS: SYSTEM OF INSIGHT              | LAYER      | NARTIS: SYSTEM OF CONTROL                                                 |
|--------------------------------------------------|------------|---------------------------------------------------------------------------|
| Claims-centric data, batch feeds daily to weekly | DATA       | FHIR/EHR, claims, pharmacy, labs, ADT and scheduling as events plus batch |
| Batch ETL and warehouse, periodic processing     | PROCESSING | Event processing that updates member state in seconds to minutes          |
| Measure logic embedded in SQL, hard to test      | POLICY     | Policy-as-code: versioned, tested measure rule packs                      |
| Descriptive dashboards: what happened?           | ANALYTICS  | Gap ranking, closure likelihood, next best action                         |
| Static, point-in-time Star reporting             | OUTCOMES   | Distance to cut points and the bonus math behind them                     |
| Care managers work lists outside the platform    | ACTION     | Orchestrated outreach and alerts, with human approval for clinical steps  |
| Role-based access, report-level logs             | GOVERNANCE | User, agent, purpose and data class checked on every request              |
| Static executive and analyst views               | EXPERIENCE | Role-based work queues: ranked gaps, provider action lists                |



# Six-layer operating model



**Rules decide. Models rank, explain and draft. People approve. Everything is logged.**

# Positioning

SYSTEM OF INSIGHT TO SYSTEM OF CONTROL

| CAPABILITY    | LEGACY ANALYTICS           | NARTIS                                         |
|---------------|----------------------------|------------------------------------------------|
| Core model    | System of insight          | System of control                              |
| Timing        | Batch, retrospective       | Event-driven, near real time                   |
| Measure logic | Embedded rules and ETL     | Policy-as-code with tests                      |
| Stars view    | Reporting and analysis     | Cut-point simulation and intervention planning |
| Action layer  | Depends on manual workflow | Orchestrated, human-approved                   |
| AI role       | Add-on analytics           | Ranking and explanation, bounded by rules      |
| Governance    | Traditional access control | User, agent, purpose and data-class controls   |
| Auditability  | Report lineage             | End-to-end evidence trail per decision         |

**Positioning:** legacy analytics helps organizations understand performance. Nartis is designed to improve it continuously.

## THE TRUST LAYER

**"The model never sees data unless both the user and the agent are authorized for that purpose, risk context, and data class."**

AI is treated as a governed actor, not a feature. Each request preserves the identity chain across user, agent, data, policy, decision, action and audit event.

# The 4.0 line is a step, not a slope.

- **Who is rated.** CMS Star Ratings cover Medicare Advantage and Part D plans only, on a 1 to 5 scale in half stars.
- **What goes in.** HEDIS measures, the CAHPS member survey, the Health Outcomes Survey, Part D measures such as medication adherence, and CMS administrative data, each scored against cut points and weighted.
- **The bonus.** Plans at 4 stars or higher receive a Quality Bonus Payment: a 5% increase to their county benchmark, 10% in qualifying double-bonus counties. The same 5% applies at 4.0, 4.5 and 5.0.
- **Rebates.** Ratings also set the share of bid savings a plan keeps as rebates, so they affect benefits and premiums.

### ILLUSTRATION

50,000 members × \$1,200 benchmark PMPM × 12 = \$720M benchmark base. A 5% QBP is worth up to \$36M a year. The realized amount depends on the plan's bid.

### RATE MATH

One closure adds 1 / denominator to a measure rate. In a 2,400-member denominator that is about 0.04 points, so moving from 80.2% to an 81.0% cut point takes about 20 closures.

**A measure may cover a subset of members, but crossing the 4.0 line changes payment for the whole contract.**

# Target architecture

AWS, HIPAA-ELIGIBLE SERVICES

## SOURCES

Claims  
837 / 835

EHR  
FHIR R4, ADT

Pharmacy  
PBM fills

Labs  
LOINC results

Enrollment  
834

## INGEST

Kinesis or MSK  
event stream

Glue and Lambda  
batch normalize

Amazon HealthLake  
FHIR R4 store

S3 + Iceberg  
Lake Formation

## STATE AND POLICY

DynamoDB  
member-measure state

Policy engine  
versioned rule packs, test  
decks

EventBridge  
gap opened / closed

## INTELLIGENCE

SageMaker  
closure propensity

Bedrock + Guardrails  
gap explanation, outreach  
drafts

Spec-to-code assistant  
human-reviewed

## ACTION

Step Functions  
human approval steps

Amazon Connect  
outreach

API Gateway + Cognito  
role-based UI

## GOVERNANCE ACROSS EVERY LAYER

IAM with purpose-bound access · Lake Formation controls · KMS · VPC endpoints and PrivateLink · CloudTrail and CloudWatch · append-only audit log in S3 Object Lock · Macie for PHI discovery

## Deterministic

- Measure status: denominator, numerator, exclusions, lookbacks
- Rate math, PDC, cut-point distance, bonus formula
- Eligibility, dates, audit

## Generative

- Measure spec to code drafts, reviewed and tested
- Gap explanations citing rule and evidence
- Outreach drafts for a person to review

### HIPAA and PHI posture

- **No PHI to a model without a BAA.** Production model calls run on Amazon Bedrock under the AWS BAA, inside the VPC.
- **BAA-eligible services only** for anything that touches PHI.
- **Minimum necessary.** Each model call gets only the fields its measure needs. No names, IDs or free-text notes.
- **De-identification** with tokenized member IDs for analytics and model training.
- **Audit trail** of user, agent, purpose, data class and decision, written append-only.

### The prototype today

- **Synthetic data only:** 10,004 generated members, dates relative to the current measurement year.
- **Real code:** 34 simplified measure rules, rate math from the population, the Stars bonus formula, an in-memory audit trail.
- **Simulated:** fixed risk constants, heuristic prioritization, illustrative cut points and weights.
- **One live model:** Ask Claude explains an open gap and drafts outreach through a validated, rate-limited serverless function.
- **Built AI-first with Claude Code** in about a week in May 2026; refreshed October 2026.

#### NEXT TO HARDEN

Licensed HEDIS value sets and test decks in CI · a closure-propensity model trained on outreach outcomes · Bedrock with Guardrails plus an eval set that checks explanations against rule output · auth, persistent audit storage, edge rate limiting · tracing and cost monitoring for model calls.

Nartis is a design and prototype by Adam Fiejdasz ([linkedin.com/in/adamfiejdasz](https://linkedin.com/in/adamfiejdasz) · [atomfade.com](https://atomfade.com)). It is not a commercial offering and has no customers or deployments. Star Ratings facts reflect CMS program rules as generally described; cut points and weights in the prototype are illustrative.