

APSYS SOLUTIONS — EPIC ANALYTICS RESOURCES

The Epic Analytics Intelligence Guide

Clarity · Caboodle · Reporting · Revenue Cycle · Dashboard Governance

How health systems turn Epic data into decisions that move clinical and financial outcomes — and the specific places the architecture breaks.

Most Epic analytics problems are not Epic problems. They are architecture problems, governance problems, and wrong-tool-for-the-job problems. The health system that has 200 Clarity reports nobody trusts has a governance problem, not a data problem. The CFO who gets three different answers to the same revenue question from three different dashboards has an architecture problem. The analytics team that spends 80% of its time answering one-off report requests instead of building strategic intelligence has a prioritization problem.

This guide was written by practitioners who have personally rebuilt broken Clarity ETL pipelines at 2 AM while a seven-hospital network waited for its morning dashboards, optimized Caboodle data warehouses that were choking on TempDB contention, and redesigned SlicerDicer governance models after organizations discovered their analysts had built 400 conflicting reports measuring the same clinical quality metric. The patterns that create these problems are documented here, along with the architectural and governance decisions that prevent them.

It also covers the strategic shift that every health system analytics team is now navigating: Epic's move toward Cogito Cloud and Microsoft Fabric, which is not a feature upgrade. It is a structural change that forces organizations to decide now what their analytics architecture looks like for the next decade.

708

organizations using Epic Caboodle globally

Growing 15.55% year-over-year (Folio3, 2026)

11.8%

average initial claim denial rate, 2024

Up from 10.2% three years prior (MGMA 2024)

>50%

of health systems report denials exceeding 10%

MGMA 2024 Benchmarking Report

\$181

cost to rework one denied hospital claim

Experian Health State of Claims Survey 2024

CHAPTER 1 THE EPIC ANALYTICS ARCHITECTURE

Chronicles, Clarity, Caboodle, and the data flow that most organizations underuse

The Four Layers — And Why Misusing Any One of Them Breaks Everything

Epic's analytics stack is four distinct layers with four distinct purposes. Most of the analytics dysfunction in healthcare IT — slow reports, conflicting numbers, dark dashboards on go-live morning — can be traced to organizations using the wrong layer for the job, or failing to understand that these layers are complementary, not redundant.

LAYER 1 Chronicles The live system. Every order, note, result, and charge written as it happens.	LAYER 2 Clarity The reporting database. Nightly ETL from Chronicles into normalized relational tables.	LAYER 3 Caboodle The enterprise data warehouse. Dimensional model built on top of Clarity data.	LAYER 4 Cogito Cloud / Fabric The emerging architecture. Real-time analytics beyond overnight batch limitations.
---	---	--	---

▪ Real-time clinical transactions ▪ No reporting queries ever ▪ High-performance, not analytics-optimized	▪ Encounter-level clinical detail ▪ Audit-defensible quality measures ▪ Complex SQL analysis ▪ Thousands of normalized tables	▪ Executive and operational dashboards ▪ Star-schema for BI tool performance ▪ Non-Epic data integration ▪ Self-service SlicerDicer	▪ Near-real-time operational insights ▪ Microsoft Fabric integration ▪ AI/ML pipeline readiness ▪ Future analytics architecture
---	---	---	---

The Data Flow That Creates Most Reporting Problems

Data flows in one direction: Chronicles → Clarity → Caboodle → dashboards. This flow has one critical vulnerability: the nightly ETL between each layer. When the Clarity ETL fails — and it will fail, whether from TempDB contention, job sequencing design flaws, or SQL Server resource saturation — every downstream layer is reporting yesterday’s data. When nobody notices for two days, the organization is making clinical and operational decisions on data that is 72 hours old.

The second vulnerability is the Clarity-to-Caboodle ETL, which is a separate job from the Chronicles-to-Clarity ETL. Organizations that monitor only one of these two jobs discover the other is broken when a department head notices the executive dashboard hasn’t updated. The question “is our data current?” requires a confirmed answer from both ETL jobs, not one.

The most dangerous failure mode in Epic analytics

Silent partial loads. The ETL completes without an error code, logs a ‘success’ status, but writes incomplete or inconsistent data. This is not an edge case. It happens when TempDB is misconfigured, when job sequencing has implicit dependencies that aren’t honored, and when SQL Server resources are saturated during the ETL window. The downstream result: reports return data without throwing errors, but the numbers are wrong. The CFO gets a revenue report, the denial rate looks fine, and the underlying data has been quietly incorrect for three weeks. This is exactly the failure pattern APSIS documented in our Caboodle ETL Stabilization case study, where 23 Caboodle tables had accumulated inconsistent records from partial loads over 11 weeks of nightly failures.

Clarity vs. Caboodle: When to Use Which

This is the question that generates the most confusion in Epic analytics teams, and the answer is not “it depends.” It is specific. Using the wrong tool for the question is expensive — either in performance (a Caboodle question asked of Clarity requires an analyst to write and debug SQL while the clinical department head waits) or in data quality (a Clarity-level question routed through Caboodle may lose the encounter-level granularity needed to make the answer defensible).

CLARITY USE CASES	CABOODLE USE CASES
When to use CLARITY	When to use CABOODLE
You need encounter-level or record-level clinical detail	You need a department head’s daily operational dashboard
The question requires audit-trail defensibility (quality measures, research, compliance)	The question is “what is our 30-day readmission rate by service line?”
You are building a cohort with complex inclusion/exclusion logic	You are connecting Epic data with non-Epic sources (claims, registry, HR feeds)
You need timestamp-level precision (when was the order placed, when was it resulted?)	You need a report that a non-technical operational leader can run themselves
You are feeding a downstream data science or ML pipeline	You are building Power BI or Tableau dashboards for executive consumption

You are producing a regulatory submission or legal-hold audit trace	You need sub-10-second query performance on aggregated data
---	---

The Conflicting Numbers Problem

The single most trust-destroying event in healthcare analytics is when the CMO and the CFO are looking at the same metric — say, 30-day readmission rate — on two different dashboards and see two different numbers. Both dashboards were built on Epic data. Both numbers cannot be right. And now every other number on both dashboards is suspect.

This problem has a technical root cause and a governance root cause. The technical cause: Clarity and Caboodle use different aggregation logic for some metrics, particularly for metrics that involve time windows, patient attribution, or encounter categorization. A readmission report built directly on Clarity PAT_ENC tables may use different encounter type filters than the Caboodle READMIT_FACT table. They will produce different numbers for the same question.

The governance cause: nobody owns the metric definition. There is no documented, approved definition of “30-day readmission rate” that every report in the organization is required to use. Two analysts, working independently from two different departments, built two reports using two different logic sets, and both got approved because nobody was in charge of whether the definitions matched.

The fix requires both layers. The technical layer: a certified metrics library where every high-stakes metric has one documented SQL definition used consistently across all reports. The governance layer: a defined approval process for new reports that verifies metric alignment before any report goes into production. Neither layer alone solves the problem.

CHAPTER 2 CLARITY PERFORMANCE ENGINEERING

The SQL Server decisions that determine whether reports run in minutes or hours

Why Epic Clarity Performance Is a Different Problem

Clarity runs on Microsoft SQL Server. Most SQL Server performance optimization techniques apply. But Clarity has a constraint that standard SQL Server environments don’t: Epic controls the schema. You cannot modify the underlying Clarity table structure, add columns to existing tables, or restructure the data model. This means the standard DBA toolkit — restructure the table, normalize differently, add a computed column — is not available.

The performance optimization available in Clarity is real and impactful, but it operates within Epic’s architecture constraints. Organizations that do not understand this distinction often bring in generalist DBAs who apply standard optimization techniques, some of which make things worse in a Clarity environment because they are optimizing against the wrong constraint.

The Five Clarity Performance Issues That Appear Most Frequently

#	ISSUE	ROOT CAUSE	THE FIX
01	Missing and stale indexes on high-volume Clarity tables	Epic’s default index strategy is optimized for transactional writes, not analytical reads. Clarity tables for PAT_ENC, CLARITY_MEDICATION, and ORDER_PROC are queried constantly by reports but may have index gaps that cause full table scans. Statistics on high-volume tables go	Index audit across the 20 highest-volume Clarity tables referenced in your critical reports. Add covering indexes where full table scans appear in execution plans. Rebuild statistics on a schedule (not the SQL Server default) and confirm the auto-update statistics

		stale quickly and the query optimizer makes increasingly wrong execution plan choices as data grows.	threshold is appropriate for Clarity's data volumes.
02	TempDB misconfiguration causing I/O contention during ETL window	Single-file TempDB creates severe I/O contention when multiple Clarity ETL jobs run concurrently. The contention creates lock waits that cascade across all jobs in the ETL window, causing some to fail silently or complete with partial data. This was the root cause of every multi-week ETL failure APSIS has diagnosed.	Split TempDB into one data file per logical processor (8 files for an 8-core SQL Server instance) with equal sizing. Place TempDB on a dedicated, fast disk subsystem isolated from Clarity data files. Verify that tempdb file autogrowth is disabled — initial sizing should be sufficient, not dependent on autogrowth.
03	Sequential job execution where parallelism would be appropriate	Many Clarity ETL configurations run all jobs serially because that is the default and nobody has reviewed the job schedule since implementation. High-volume, independent jobs — ones with no data dependency on each other — can run in parallel, cutting total ETL window time significantly.	Map the job dependency tree. Identify jobs that have no upstream data dependency on each other. Configure parallel execution for independent jobs. Confirm that parallelism does not create resource contention by monitoring CPU and I/O utilization during the parallel window in a test run.
04	Reports written against Clarity tables when Caboodle would be appropriate	Reports that aggregate data across large time windows — rolling 90-day readmission rates, annual quality measure calculations — are being run as Clarity queries against normalized tables. This forces multi-table joins across millions of rows. Caboodle's dimensional model answers the same question with one or two joins against pre-aggregated fact tables.	Identify reports that take longer than 5 minutes and review them against the Clarity vs. Caboodle decision framework. Rebuild aggregation-heavy reports against Caboodle facts. Reserve Clarity for encounter-level, audit-defensible, or complex cohort logic where Caboodle's pre-aggregated model loses precision.
05	Uncontrolled report proliferation degrading shared SQL Server resources	Every report running on the Clarity server consumes SQL Server resources. An organization with 300 active Clarity reports, 40% of which run overnight, is competing for the same SQL Server CPU and I/O as the ETL jobs. Reports that are not governed — built ad hoc, never retired, running on schedules nobody has reviewed — create resource contention that shows up as 'slow ETL' when the real problem is 'too many reports running at the same time'.	Conduct a report usage audit: which reports have not been run in 90 days? Which run on schedules that overlap the ETL window? Which have 0 subscribers? Retire inactive reports, move all non-critical scheduled reports to off-peak windows, and implement a governance gate for new scheduled report creation.

A word on execution plans

The execution plan is the most direct diagnostic tool in Clarity performance work. Every slow Clarity report has an execution plan that will show you exactly where the time is being spent: a table scan where an index seek should be, a key lookup that could be eliminated by a covering index, a hash join that could become a nested loop join with better statistics. If your Clarity analyst cannot read an execution plan, they cannot diagnose Clarity performance problems. They can only guess. The difference between a generalist DBA and an Epic Clarity performance specialist is knowing which execution plan problems are fixable within Epic's schema constraints and which require a different approach entirely.

CHAPTER 3 DASHBOARD STRATEGY & GOVERNANCE

Building a dashboard portfolio that leaders actually trust and use

The Dashboard Problem That Most Analytics Teams Won't Admit

Most healthcare organizations have more dashboards than they use. The dashboards were built in response to requests — a department head asked for it, a quality initiative required it, a new CMO wanted to see something specific — and each one was treated as a net-new deliverable rather than an addition to a governed portfolio. Three years later, the analytics team is maintaining 150 dashboards, 40% of which haven't been viewed in the past 90 days. The ones that are being viewed are producing conflicting numbers because nobody governed the metric definitions. The analysts spend most of their time either fixing broken reports or fielding requests for new ones.

Epic's own guidance on this is worth quoting directly. On their analytics intelligence page, Epic writes: "The purpose of analytics is not to run reports frequently. The purpose is to make us better at what we do. And the way to do that is to ask users: why do you want this information, and what do you plan to do with it?" That is a governance philosophy, not a technical specification. Organizations that treat every analytics request as a development request never ask that question. The ones that build trusted, high-use analytics portfolios do.

The Four-Tier Dashboard Architecture

Not all dashboards are the same problem. The governance failure that creates 150 dashboards nobody trusts usually comes from treating a C-suite strategic dashboard the same way as an analyst's operational data pull. A structured four-tier architecture assigns different governance standards, different update frequencies, and different ownership models to different types of analytics output.

TIER	AUDIENCE	DATA SOURCE	REFRESH	GOVERNANCE STANDARD
Tier 1 Strategic	Board, C-suite, VPs	Caboodle fact tables	Daily or weekly	Certified metrics only. Every metric has one approved definition. Any change requires analytics governance committee sign-off. No self-service modifications.
Tier 2 Operational	Department leads, directors, service line chiefs	Caboodle + selected Clarity	Nightly or near-real-time	Designated metric owner per report. Metric definitions documented. New metrics require Tier 1 review if they appear on any existing strategic dashboard.
Tier 3 Analytical	Analysts, quality teams, research	Clarity (analyst-built)	On-demand or scheduled	Code review required for production deployment. SQL reviewed against performance standards. Execution plan documented for queries running >2 minutes. 90-day usage review triggers retirement process.
Tier 4 Self-Service	Clinical and operational end users	SlicerDicer (Caboodle)	Real-time within SlicerDicer	Data domain access controlled by role. PHI access subject to HIPAA minimum necessary review. Users may not save reports that use metric names conflicting with Tier 1 or Tier 2 certified metrics.

SlicerDicer: The Self-Service Tool That Becomes a Governance Crisis Without a Model

SlicerDicer is genuinely useful. It allows clinical and operational users to explore Epic data without writing SQL, answer population health questions in seconds, and surface trends that would otherwise require an analyst engagement. The problem is that "self-service" in a healthcare analytics context is not the same as "ungoverned."

The governance failure pattern is consistent across organizations: SlicerDicer is turned on without a data access model, without metric naming standards, and without any process for reviewing what users are saving and sharing. Within 12 months, the organization has dozens of saved SlicerDicer sessions with names like “Readmissions v3 FINAL NEW” and “ED throughput – Maria’s version,” many of which use metric logic that contradicts the Tier 1 certified metrics on the executive dashboard. A physician sees both and asks why the numbers disagree. The analytics team has no good answer.

- Access control before activation: define which data domains each role can see in SlicerDicer. Not all clinical data should be available to all users. The HIPAA minimum necessary standard applies to SlicerDicer access the same way it applies to any ePHI access.
- Naming governance for saved sessions: organizations with mature SlicerDicer governance prohibit users from saving sessions that use the same metric names as Tier 1 or Tier 2 certified reports without analytics team review. This prevents the “readmission rate (Jane’s version)” proliferation.
- Quarterly access audit: who has saved what, and are any saved sessions being shared broadly that should be reviewed by the analytics governance committee?

CHAPTER 4 REVENUE CYCLE INTELLIGENCE

Why the Epic data you already have contains your denial prevention strategy

The State of Healthcare Denials in 2025 and 2026

Claims denial is getting worse, and it is getting more expensive. Initial claim denials hit 11.8% in 2024, up from 10.2% three years earlier, with more than half of U.S. healthcare organizations reporting denial rates above 10%. The average amount denied per inpatient claim rose 12% in 2025 and 14% for outpatient claims, according to MDaudit data from 4,500 facilities. Telehealth-related denials rose 84% in a single year. The conventional response to rising denials is more staff. More denial management specialists, more appeal coordinators, more coding reviewers. But the math doesn’t work indefinitely. The cost to rework a denied hospital claim is \$181. Organizations that rely on reactive denial management — deny happens, staff works it, gets reworked or written off — are fighting the denial spiral with a bucket. Predictive denial prevention using the analytics data organizations already have in Epic is how the most effective organizations are breaking the cycle.

What Epic Already Knows About Your Denials

Every health system running Epic has the data required to build a denial prediction model. The data already exists in Clarity and Caboodle. What most organizations don’t have is the analytics infrastructure to turn that data into upstream prevention rather than downstream rework.

- **46% of denials are caused by missing or incorrect patient information at registration, according to Experian Health’s 2024 State of Claims Survey. Clarity’s patient registration tables contain every field that payers check. An analytics report comparing your registration data completeness against your denial root cause codes will tell you exactly which registration gaps are causing which denial categories.** Registration data quality:
- **Denials tied to authorization failures are among the most preventable. Clarity’s authorization data, cross-referenced against your denial codes, identifies which service lines, which payers, and which order types are generating auth-related denials — and at what volume.** Prior authorization tracking:
- **Not all payers deny the same things. A Clarity analysis of denial root causes stratified by payer will surface patterns that the billing team knows anecdotally but has never quantified: payer X consistently denies outpatient imaging claims that require a specific modifier; payer Y has started auto-denying a category that was previously clean.**

Quantifying these patterns enables proactive response before the denial wave hits. Payer-specific denial patterns:

- Epic’s Caboodle contains both the claim submitted and the denial code received. Pairing these with Clarity encounter data enables a retrospective analysis that identifies which documentation patterns consistently precede which denial types — the foundation of a CDI program driven by analytics rather than intuition. Coding accuracy signals:

The Five Revenue Cycle Dashboards Every Health System Needs — And Most Don’t Have Right

DASHBOARD	WHAT IT SHOWS	EPIC DATA SOURCE	WHY MOST ORGS DON'T HAVE IT RIGHT
Denial Root Cause Matrix	Denial volume and dollar amount by root cause code, payer, service line, and month — trending 24 months	Clarity CLAIM_EX tables, ClaimEdit, HAR_DENIAL	Built as one-dimensional report; no payer stratification; no trend; no dollar weighting by root cause
Registration Completeness Scorecard	Missing field rates by registration type, site, and staff, trended against downstream denial rates	Clarity PAT_ENC, COVERAGE, PAYER tables correlated with denial data	Registration reporting lives in a separate system or is not correlated to denial outcomes at all
A/R Aging with Payer Behavior Analysis	Aging buckets by payer, with days-to-first-denial and rework cycle time metrics	Clarity HAR, PHX_CLAIM_EX, payment tables	A/R aging exists but has no payer behavior overlay; leadership sees aging without knowing which payer is causing it
CDI Opportunity Dashboard	Documentation patterns that precede denials vs. documentation patterns that result in clean claims, by DRG and physician	Clarity encounter documentation, Caboodle ENCOUNTER_FACT, denial history cross-reference	CDI program is managed by gut feel and anecdote; no systematic analytics feedback loop from denied claims to documentation improvement
Prior Auth Coverage Gap Report	Authorization rate by service type, payer, and site; flagging service types with rising auth failure rates before the denial wave	Clarity CL_AUTH, ORDER_PROC, COVERAGE tables	Auth tracking exists in Epic but is not connected to denial analytics; organizations discover auth policy changes from payers by seeing denial spikes, not by monitoring auth approval rates in advance

CHAPTER 5 DATA QUALITY
Why Epic data quality problems are not Epic problems

The Root Cause of Most Healthcare Data Quality Failures

Healthcare organizations routinely blame Epic for data quality problems that originate upstream of Epic, in the workflows and processes that determine what data gets entered into the system in the first place. Epic can only report what was documented. If the documentation is incomplete, inconsistent, or wrong, the reporting will be incomplete, inconsistent, and wrong — regardless of how sophisticated the analytics infrastructure is.

That said, Epic-layer data quality problems do exist and they are distinct from workflow documentation problems. Understanding which category a data quality issue belongs to determines who is responsible for fixing it and what the fix looks like.

Category 1: Workflow Documentation Problems (Not Fixable in Analytics)

These are problems where clinicians or operational staff are not documenting in Epic the way the analytics team expects them to document. A readmission dashboard that shows “unknown” in the discharge disposition field does not have a Clarity problem. It has a nursing documentation compliance problem. The fix is workflow redesign, training, and compliance monitoring — not SQL tuning.

- Inconsistent use of standard fields across sites or providers within the same organization
- Free-text documentation where structured discrete data should be captured (critical for quality measure calculation)
- Missing encounter type assignments creating encounter classification ambiguity in reports
- ADT workflow gaps — admits, discharges, and transfers not documented at the time they occur

Category 2: Epic Build and Configuration Problems (Fixable in Build)

These are problems where the Epic system is configured in a way that creates data quality gaps. They require Epic build expertise to fix, not analytics layer changes.

- Missing or incorrectly mapped interface values — data coming from third-party systems that doesn't map to Epic's standard value sets
- Crosswalk mapping errors between Epic and external code sets (ICD-10, SNOMED, CPT) that create incorrect categorization in reports
- SmartForm and order set build that allows optional fields which should be required for analytics purposes
- Registration workflow configuration that permits claim submission without required demographic fields

Category 3: Analytics Layer Problems (Fixable in Clarity/Caboodle)

These are genuine data quality problems at the analytics layer — the tables, ETL logic, and reporting definitions that process and present Epic data.

- Stale or incorrect Clarity ETL transforms that apply outdated business logic to current data
- Caboodle fact table calculations that use different aggregation logic than the Clarity queries they are supposed to match
- Silent ETL partial loads (the most dangerous category, described in Chapter 1) that write incomplete data without error codes
- Report filter logic that excludes encounter types, visit types, or locations in ways that are not documented or understood by the report consumers

The Data Quality Audit Framework

A meaningful data quality audit in an Epic environment works from the source backward: start with the analytics output, identify the specific discrepancy, trace it upstream to determine whether it originates in workflow documentation, Epic build, or the analytics layer, and assign remediation to the appropriate team. Assigning all data quality issues to the analytics team is the most common organizational error in healthcare analytics governance.

The question to ask when a report produces wrong data

Work through these in order and stop when you find the root cause. (1) Is the ETL current — did last night's job complete successfully with full data? (2) Is the SQL logic correct — does the report filter, join, and aggregate against the right field definitions? (3) Is the source data in Clarity accurate — does it match what's in Chronicles? (4) Is the Chronicles data accurate — does it reflect what actually happened clinically? If the answer to #4 is 'no,' the problem is not in analytics. It is in the

workflow that failed to document the clinical event correctly. Analytics cannot fix undocumented care.

CHAPTER 6 THE COGITO CLOUD AND MICROSOFT FABRIC DECISION

The structural shift every Epic analytics leader needs to make a plan for

What This Transition Actually Means

Epic’s move toward Cogito Cloud and Microsoft Fabric integration is not a product announcement. It is a structural realignment of the entire Epic analytics architecture. The overnight batch model that has defined Epic analytics since Clarity was introduced — Chronicles runs during the day, Clarity extracts at night, dashboards update by morning — is being supplemented by a near-real-time analytics pathway that runs on Fabric’s lakehouse architecture.

The implications are real and forcing. RSM’s analysis, published in April 2026, puts it directly: “Epic customers are being pulled toward Fabric by design, and the decisions made now will determine cost structures, analytics agility and artificial intelligence readiness for years to come.” This is not optional. Epic is building its analytics product roadmap around Fabric. Organizations that do not engage with this transition will find their analytics infrastructure increasingly disconnected from Epic’s investment direction.

The Specific Patterns RSM Documents in Organizations Resisting the Transition

- Analytics tool sprawl from years of layered BI and replication tools, each added to solve a problem that the previous tool created
- Escalating cloud and licensing costs driven by duplicated data movement across systems that were built independently
- Delayed insights tied to overnight Epic refresh cycles — organizations that need intraday capacity management data are running reports that are 12–16 hours old
- Conflicting vendor guidance leaving leadership uncertain about which migration path is defensible in terms of cost, compliance, and Epic support

The Three Questions Every Analytics Leader Should Answer Now

Question 1

What analytics use cases do I have today that the overnight batch model cannot satisfy?

Capacity management, ED throughput, real-time bed huddles, intraday census tracking — if these are currently populated with yesterday’s data, you already have the business case for near-real-time analytics. Document these use cases before evaluating architecture options.

Question 2

What is our current analytics technical debt, and does a Fabric migration reduce or compound it?

Organizations with significant custom Clarity logic, non-Epic data integrations, and existing Power BI infrastructure will have different migration complexity than organizations building greenfield on a recent Epic implementation. Inventory the debt before committing to a migration timeline.

Question 3

What is our AI/ML analytics strategy, and does it require data architecture decisions that Cogito Cloud enables?

Predictive denial prevention, sepsis scoring, readmission risk models — these analytics use cases require a data architecture that can support ML pipelines. If your organization has AI analytics on the roadmap, the architecture decision and the AI strategy need to be made together, not sequentially.

The measure of an analytics program is not the number of reports it produces.

It is the number of decisions the organization makes differently because of what the analytics program surfaced. A health system with 20 well-governed, trusted, consistently-accurate dashboards that clinical and operational leadership uses to make daily decisions is analytically more mature than one with 200 dashboards that nobody trusts. Building that 20-dashboard program — with certified metric definitions, reliable ETL infrastructure, governed self-service access, and a revenue cycle intelligence layer that turns denial patterns into prevention strategies — is the work. The architecture is the means. The decisions are the point.

Epic analytics gaps? APSIS can close them.

Clarity performance engineering, Caboodle optimization, revenue cycle intelligence, and dashboard governance — deployed in 24–48 hours.

apsisinc.com | 813-999-4853
info@apsisinc.com

About Apsis Solutions Inc.

Apsis Solutions Inc. is a Tampa-based healthcare IT consulting firm specializing in Epic analytics, Clarity/Caboodle performance engineering, revenue cycle intelligence, data governance, and IT workforce staffing. We have personally rebuilt ETL pipelines across multi-hospital networks, redesigned analytics governance programs that had lost stakeholder trust, and built denial prevention dashboards that turned reactive rework into proactive revenue protection. We serve health systems, academic medical centers, federal agencies, and prime contractors as a capable prime vendor, strategic subcontractor, and managed services partner. The agility of a boutique firm. The capability of a prime. — aphisinc.com | 813-999-4853 | info@apsisinc.com