

Enterprise Data Warehouse & Analytics Platform For A Large U.S. Eyecare Organization

Introduction

This case study describes the design and implementation of a large-scale enterprise Data Warehouse (DW) and analytics platform for one of the largest eyecare organizations in the United States, operating 550+ practices nationwide. The initiative answered critical questions around operational efficiency, clinical productivity, financial performance, revenue cycle effectiveness, and enterprise-level decision-making, using measurable improvements in reporting accuracy, performance, and adoption as evaluation criteria.

Key Questions Answered

01

How can enterprise-wide performance be measured consistently across 550+ decentralized practices?

02

How can operational, clinical, marketing, merchandising, RCM, and financial KPIs be unified into a single, trusted system?

03

How can leadership move from static, delayed reporting to near real-time, self-service insights?

04

How can analytics scale to support hundreds of KPIs without degrading performance?

Evaluation Criteria

01

Accuracy and reconciliation against source systems

02

Performance and query response times at enterprise scale

03

KPI standardization and governance

04

User adoption across roles (practice, regional, corporate, executive)

05

Ability to support future growth and acquisitions

Sources Used

01

Internal
Operational And
Financial Systems

02

Practice
Management
And Clinical Data
Sources

03

Marketing
Platforms And
Call-Tracking
Systems

04

Revenue
Cycle And
Billing Systems

05

Internal
Documentation,
KPI Definitions,
And Leadership
Interviews

The Situation

Before this engagement, the organization operated with highly fragmented reporting across hundreds of practices. Each function i.e. operations, finance, marketing, merchandising, and clinical leadership used separate tools, spreadsheets, and ad-hoc extracts. Reporting cycles were slow, often lagging business reality by days.

Baseline Observations Included:



Inconsistent KPI definitions across regions and teams



Manual reconciliation between finance, operations, and RCM



Limited visibility into doctor-level and practice-level performance

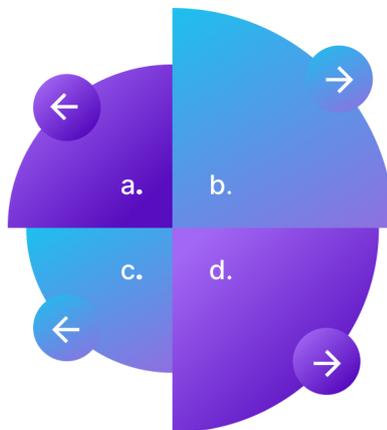


Heavy dependence on analysts for report creation

The Problem

Leadership lacked a single source of truth across the enterprise

Decision-making was reactive rather than proactive



Practice and doctor performance could not be compared reliably

Scaling reporting for acquisitions and new practices was not feasible with existing processes

This directly impacted margin control, doctor productivity optimization, marketing ROI, and leadership confidence in data.



What Had Been Tried



Spreadsheet-driven reporting maintained by regional teams



Siloed BI tools connected directly to source systems



Point-in-time extracts created by analysts

These Approaches Failed Due To:

01

Poor performance at scale

03

High maintenance cost and error risk

02

No centralized KPI governance

Constraints

01.

Large volume of data across hundreds of locations

02.

Multiple legacy source systems

03.

Need for high performance with 300+ kpis

04.

Enterprise security and access control requirements

05.

Minimal disruption to ongoing operations

Stakeholders



Practice Managers And Regional Operators



Marketing And Merchandising Teams



Executive Leadership And PE Sponsors

Doctors And Clinical Leadership



Revenue Cycle And Finance Teams



Scope And Approach

What was in scope



01 Design and implementation of an enterprise Data Warehouse in Microsoft SQL Server

02 OLAP and high-performance analytics layer using StarRocks

03 Definition, standardization, and implementation of 300+ KPIs

04 Development of 70+ role-based reports and dashboards

05 Secure API layer built in .NET Core

06 Web-based analytics UI built in Angular

07 Enterprise-wide data governance and KPI validation

What Was Out Of Scope



Replacement of source transactional systems



Operational workflow changes inside practice systems



Real-time clinical decision support (future phase)

Inputs

01

Stakeholder interviews across all business functions

02

Source system schema reviews and data profiling

03

Historical data audits and reconciliation exercises

04

Existing reports, spreadsheets, and KPI definitions

05

User feedback from pilot groups and leadership



Validation



KPI definition workshops with functional owners

01

02

Parallel run comparisons with legacy reports



Financial reconciliation with audited numbers

03

04

Performance and load testing at enterprise scale



Leadership sign-off at each milestone

05



Findings And Insights

What the data showed



Significant variation in practice and doctor performance hidden by averages



Marketing spend effectiveness varied widely by geography and channel



RCM delays and denials were correlated with specific operational patterns



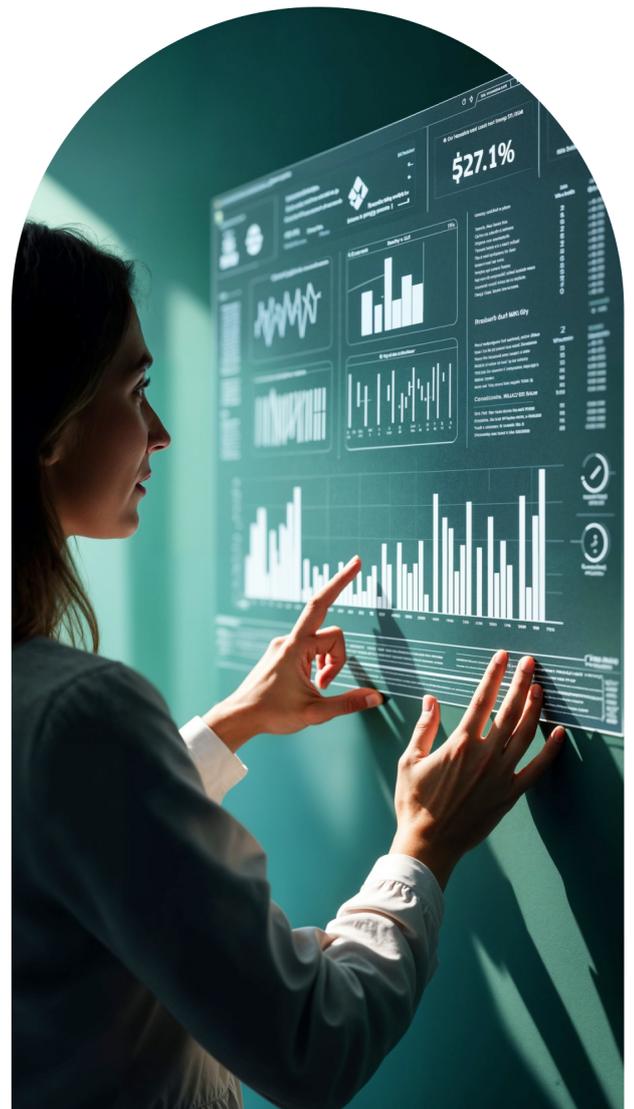
Merchandising performance depended heavily on localized inventory strategies

Root Causes

- > KPI definitions were inconsistent or undocumented
- > Data lived in silos with no enterprise model
- > Reporting tools were not designed for scale or self-service

Critical Insights

- > Standardized KPIs were more valuable than additional metrics
- > Performance depended on pre-aggregated, OLAP-optimized data - not raw queries
- > Adoption increased dramatically when reports were role-specific and intuitive



Solution

What was recommended

A layered analytics architecture that separated ingestion, transformation, analytics, and presentation - allowing scalability, governance, and performance without sacrificing flexibility.

What Was Built



Enterprise Data Warehouse (MSSQL)

- Centralized fact and dimension models
- Historical tracking and auditability



OLAP Layer (StarRocks)

- High-performance aggregates for enterprise reporting
- Support for complex, multi-dimensional KPIs



KPI Framework

300+ KPIs across:

- Practice Operations
- Doctor Operations
- Marketing
- Merchandising
- Revenue Cycle Management
- Finance
- Senior Leadership



API Layer (.NET Core)

- Secure, scalable access to analytics
- Decoupled UI from data logic



Analytics UI (Angular)

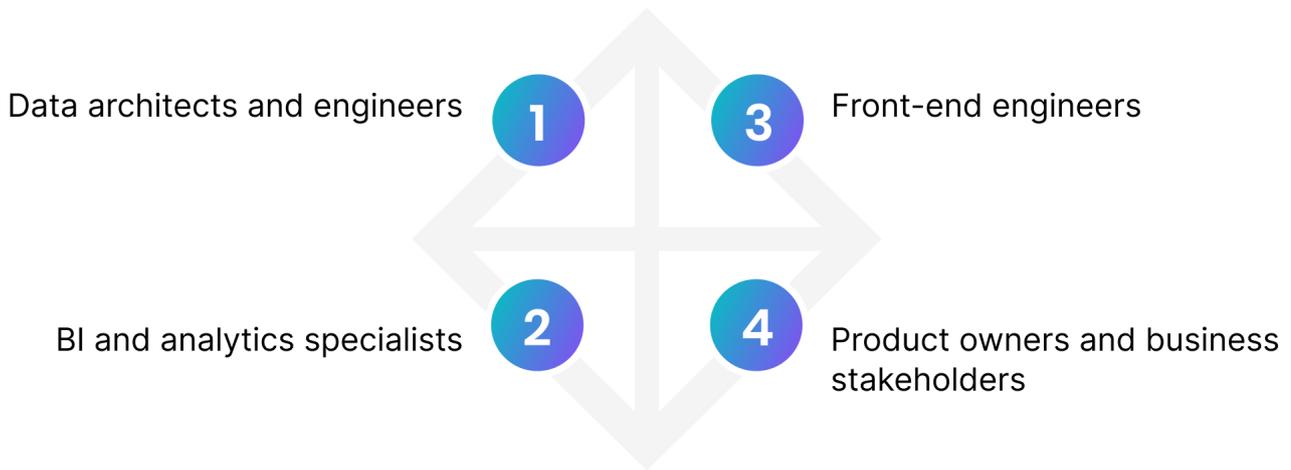
- 70+ dashboards and reports
- Role-based views for practice, regional, and enterprise users

How It Was Delivered

Timeline and milestones



Roles

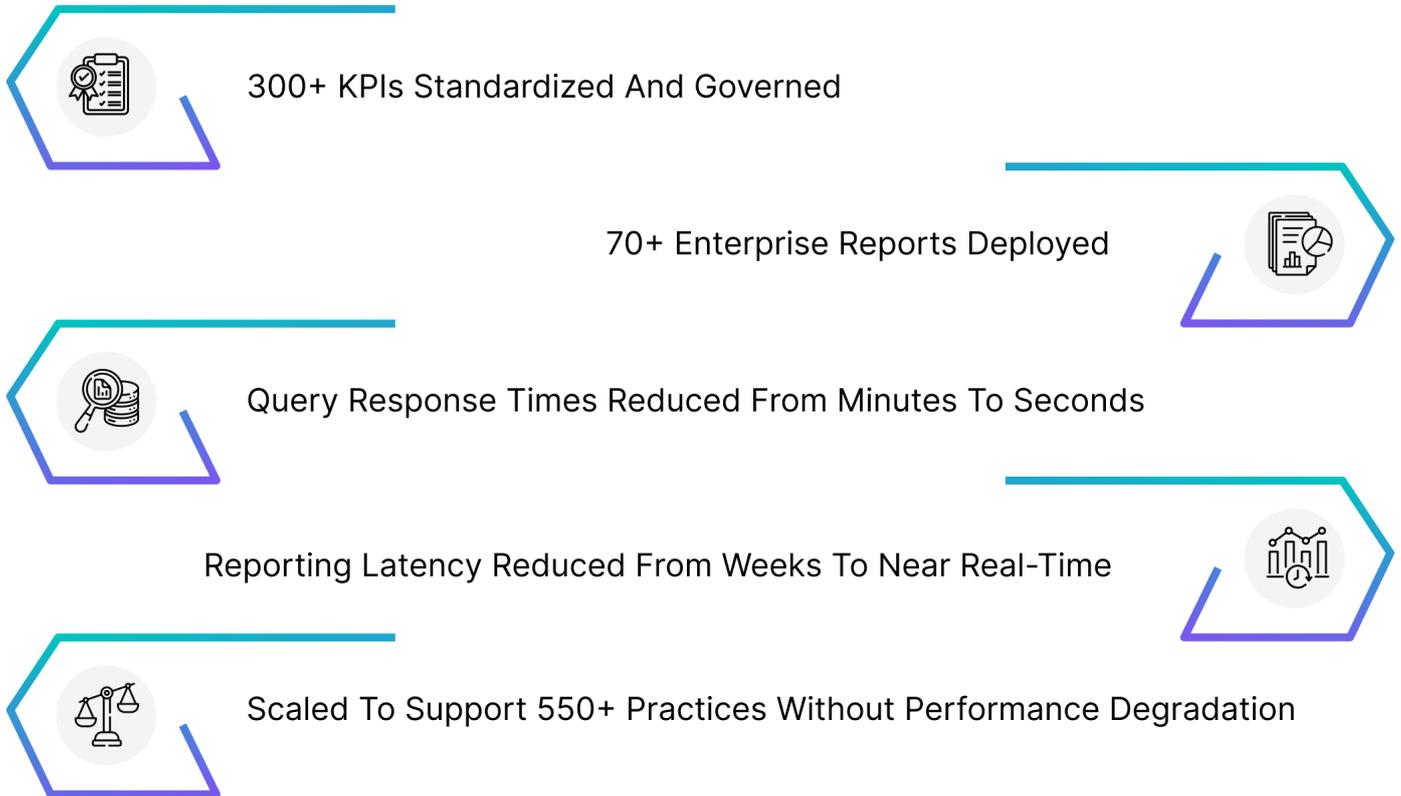


Dependencies And Integrations

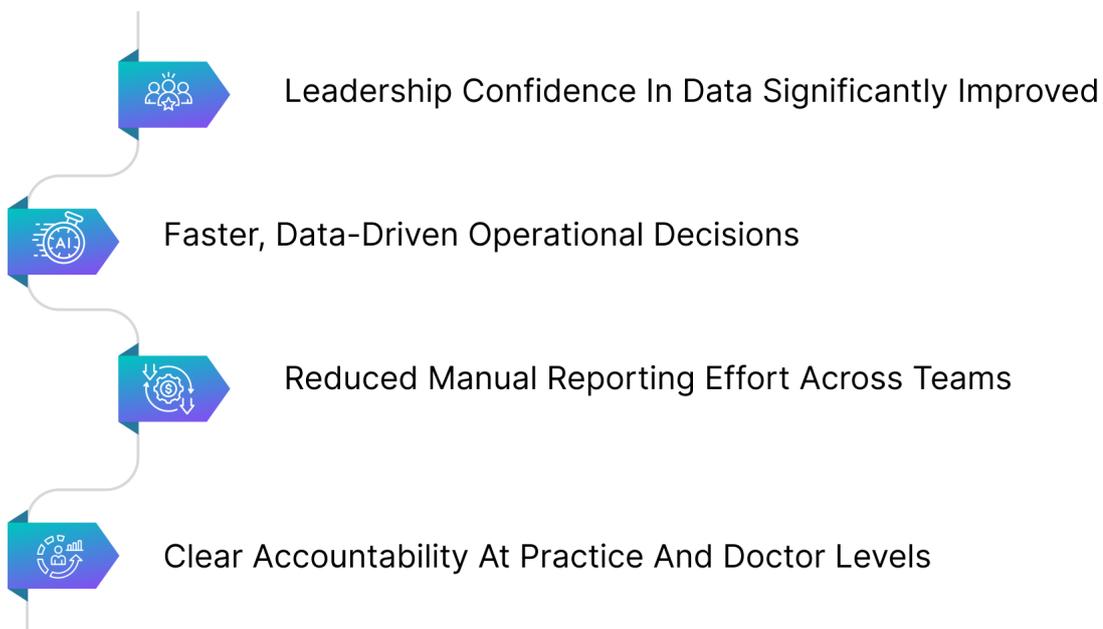


Results

Quantitative outcomes



Qualitative Outcomes



Adoption

Used daily by operations, finance, marketing, and leadership teams

Became the default system for enterprise performance reviews

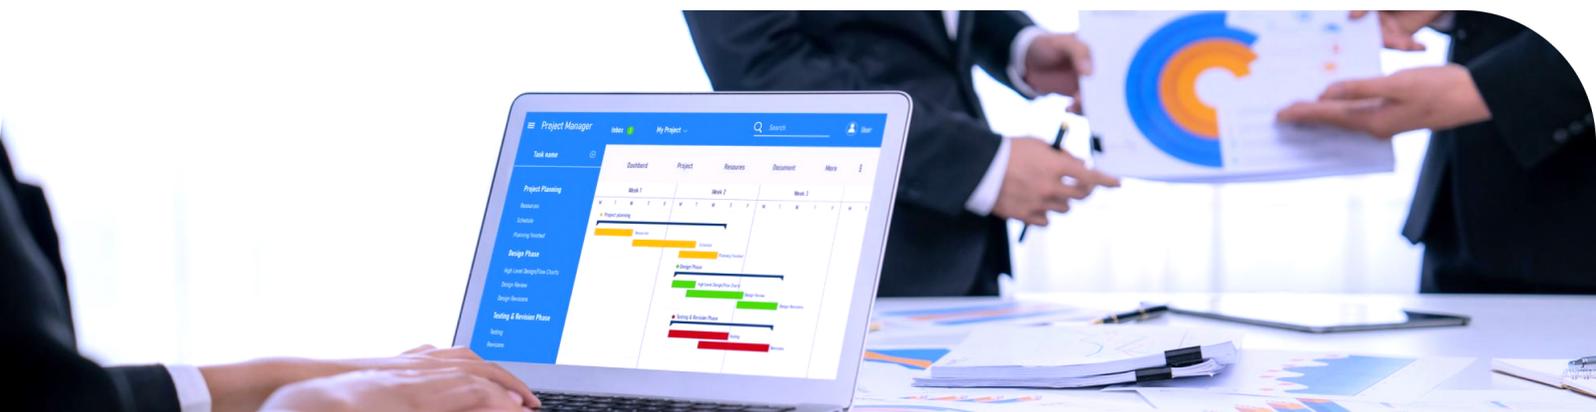
Reflection

What Worked

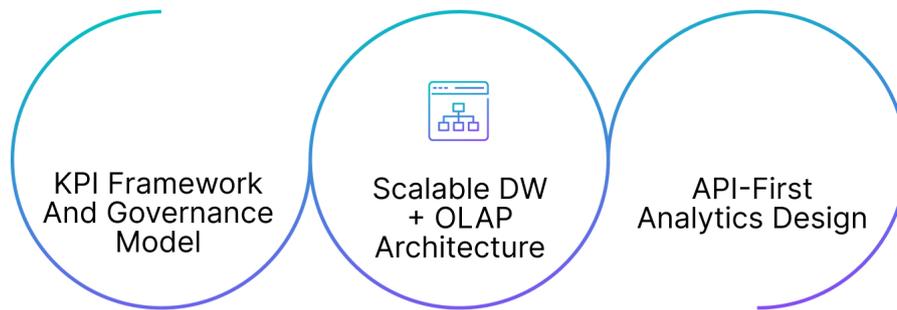
- 01 Strong KPI governance and stakeholder alignment
- 02 OLAP-driven performance architecture
- 03 Role-based reporting design

What Did Not Work Initially

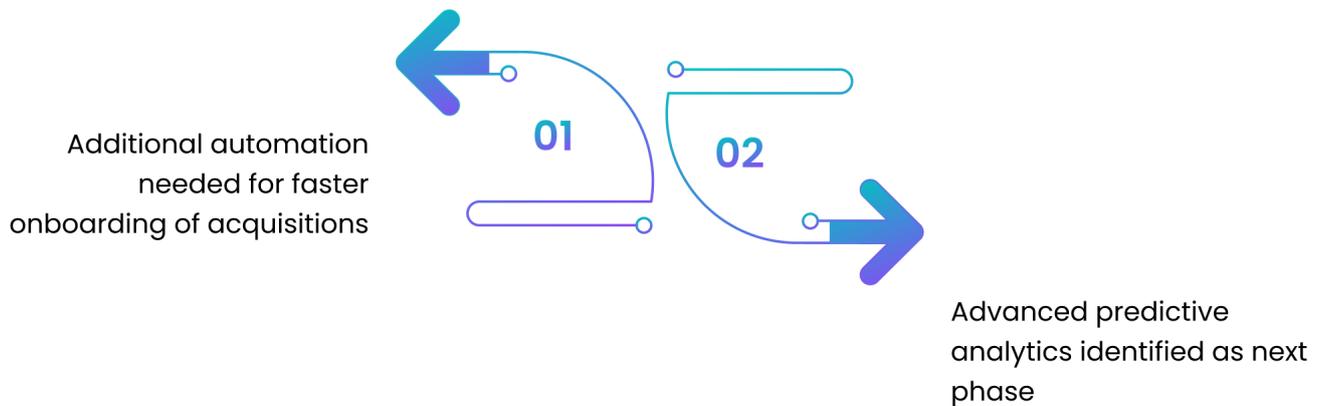
- 01 Overly granular early KPI drafts
- 02 Underestimated effort in legacy data normalization



What Is Reusable



Limits And Adaptation



Takeaway

This engagement transformed fragmented, unreliable reporting into a scalable enterprise analytics platform supporting hundreds of practices and leadership stakeholders. By combining a robust Data Warehouse, high-performance OLAP, standardized KPIs, and intuitive dashboards, the organization moved from reactive reporting to proactive, data-driven decision-making.

The key lesson: At enterprise scale, success depends less on tools and more on disciplined data modeling, KPI governance, and performance-first architecture.