

# GEMMA ANALYTICS

## PROPOSAL

Modern Data Platform & BI Migration

---

Prepared for

**Gregory's Food Service Group S.A.**

March 2026

Valid for 60 days

## EXECUTIVE SUMMARY

Gemma Analytics proposes to partner with Gregory's on the migration from the current on-premise SQL Server and Targit BI environment to a modern, cloud-native data platform – building the foundation for scalable reporting, richer analytics, and future AI capabilities.

Gregory's has built a strong BI culture over the past seven years with solid user adoption and well-understood data models. The goal is not to start over but to **carry that knowledge forward** into a platform that removes the current limitations: restricted data source connectivity, limited visualization options, descriptive-only analytics, and a tightly coupled MOLAP cube that constrains flexibility.

We propose a **phased engagement** that starts with a focused discovery and architecture design phase before committing to full implementation. This ensures the platform decisions (cloud provider, warehouse, orchestration, migration strategy) are grounded in Gregory's actual data landscape rather than assumptions.

**Lightdash** serves as the BI and semantic layer – a modern, open-source alternative to Targit with rich visualizations, metric trees, and a code-first approach that integrates tightly with the data platform.

## GREGORY'S TODAY

### CURRENT LANDSCAPE

Gregory's operates a franchise network of 367+ stores across Greece and 9 European countries, serving over 250,000 daily customers. The data landscape spans multiple systems:

Source	Data	Current Use
<b>ERP (Business Central)</b>	Sales, Orders, Fixed Assets, GL, Financials	Reporting & Financial Control
<b>Retail POS (Profit)</b>	Transactions, Products, Ticket Size, Payments	Store / Network Reporting
<b>Delivery Channels</b>	Online orders, customer data, web analytics	Order & Marketing KPI Reporting
<b>Aggregators</b>	eFood, Wolt, BOX orders and delivery metrics	Consolidated with online orders
<b>B2B CRM</b>	Tickets, store evaluations, training, complaints	Reporting & Compliance
<b>Power Apps (Dataverse)</b>	Various business area data	Early stage
<b>Asana</b>	Cross-department projects	Low penetration
<b>Excel / CSV</b>	IT, Legal, HR, supplier KPIs	Ad-hoc, siloed

The current data warehouse runs on **SQL Server on-premise** with **Targit** as the BI tool. The Targit MOLAP cube provides fast reporting and strong user adoption, but imposes significant constraints: limited data connectors, a narrow visualization library, descriptive statistics only, infrequent product updates, and a binding effect that makes it difficult to evolve the architecture independently.

The warehouse holds approximately **764 GB** across 916 tables and 828 views – though a migration will likely focus on a critical subset rather than a full lift-and-shift.

## WHERE GREGORY'S WANTS TO GO

1. **Cloud-native data platform** – scalable, managed, with broad connector coverage
2. **Modern BI layer** – rich visualizations, decomposition metric trees, custom visuals
3. **Beyond descriptive analytics** – AI capabilities, predictive statistics, scenario modeling
4. **Real-time where it matters** – faster refresh cycles for operational decisions
5. **External data integration** – combining warehouse data with public and online sources
6. **A platform with a future** – frequent updates, active roadmap, growing ecosystem

## OUR APPROACH

This is a **migration**, not a greenfield build. Gregory's has a mature data culture and well-understood business logic embedded in the current Targit cubes and SQL Server models. Our approach respects that:

- **Discovery first** – before committing to tooling or timelines, we audit the current environment, identify business-critical data flows, and design the target architecture together with Gregory's BI team
- **Preserve what works** – existing KPI definitions, business logic, and data models are migrated and validated, not reinvented
- **Incremental migration** – data sources and reports move to the new platform in priority order, with the legacy system running in parallel until stakeholders are confident in the new environment
- **Code-first, AI-assisted** – all pipelines, transformations, and configurations are version-controlled. We use AI-assisted development tooling as standard practice for accelerated delivery
- **Lightdash as the BI layer** – Lightdash provides metric definitions in code (dbt semantic layer), interactive dashboards, and self-service exploration – addressing Gregory's need for richer visuals, metric trees, and a modern analytics experience

## PROPOSED WORK STREAMS

The engagement is structured in three phases. Phase 1 (Discovery) is a fixed-scope, fixed-price deliverable. Phases 2 and 3 are scoped based on Phase 1 findings.

### Phase 1: Discovery & Architecture Design

A focused audit of Gregory's current data landscape, source systems, existing Targit models, and reporting requirements. Produces a concrete architecture blueprint and migration roadmap that Gregory's can evaluate before committing to implementation.

**Duration:** 2–3 weeks

**Model:** Fixed price

**Output:** Architecture blueprint, migration roadmap, source prioritization, cost projection

## Phase 2: Data Platform Build

Cloud infrastructure setup: data warehouse provisioning, ingestion pipelines for priority data sources, orchestration, CI/CD, monitoring, and the dbt transformation layer. Connects Gregory's core systems to the new platform.

**Start:** After Phase 1 sign-off

**Duration:** 6–10 weeks (scoped in Phase 1)

**Model:** Time & Materials or Fixed Price (to be defined)

## Phase 3: Analytics Engineering & BI Migration

Data modeling, KPI migration, and dashboard build-out in Lightdash. Existing Targit reports are rebuilt and validated against the current environment. New analytics capabilities (metric trees, advanced visuals) are developed iteratively with stakeholders.

**Start:** Overlaps with late Phase 2

**Duration:** 4–8 weeks (scoped in Phase 1)

**Model:** Time & Materials

## PHASE 1: DISCOVERY & ARCHITECTURE DESIGN

Phase 1 is a standalone, fixed-scope engagement. It produces a detailed architecture blueprint and migration plan that Gregory's can use to make an informed go/no-go decision on the full implementation.

### DELIVERABLES

#### Data Source Audit

- Inventory of all active data sources, refresh frequencies, and dependencies
- Assessment of current Targit cube logic: which KPIs, calculations, and business rules are embedded
- Identification of data quality issues, gaps, and undocumented transformations
- Source prioritization: what moves first, what can wait, what is retired

#### Architecture Blueprint

- Recommended cloud data warehouse (Snowflake, BigQuery, or other – based on Gregory's existing cloud presence, data residency requirements, and cost profile)
- Ingestion architecture: which sources use managed connectors, which require custom pipelines
- Transformation layer design (dbt) – mapping current cube logic to a modern dimensional model
- Orchestration strategy (Airflow, Prefect, or managed alternative)
- Lightdash configuration and semantic layer design
- Security, access control, and data governance model

#### Migration Roadmap

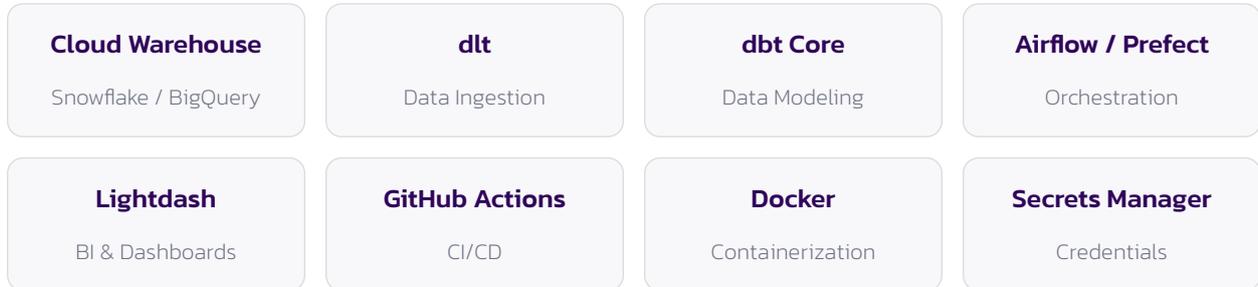
- Phased migration plan with clear milestones and dependencies
- Parallel-run strategy: how long the legacy system stays active
- Resource and cost projection for Phases 2 and 3
- Risk assessment and mitigation plan
- Recommended team structure (Gemma + Gregory's internal resources)

#### Stakeholder Sessions

- 2–3 working sessions with Gregory's BI team and key business stakeholders
- KPI prioritization workshop: which metrics matter most and which reports are critical path
- Architecture review session: presenting and refining the blueprint together

## RECOMMENDED TECHNOLOGY STACK

The specific tools will be confirmed during Phase 1 discovery. The following represents Gemma's standard modern data stack, which we expect to form the basis of the architecture:



- Cloud warehouse** – the choice between Snowflake and BigQuery depends on Gregory's existing cloud footprint, data residency needs (EU), and cost considerations. Both are strong options; Phase 1 will make a concrete recommendation.
- dlt (data load tool)** – Python-based ingestion framework with native connectors for Business Central, CRM systems, and APIs. Custom pipelines for sources without standard connectors (POS, aggregators).
- dbt Core** – industry-standard transformation layer. Gregory's existing star schema (Fact\_Sales, Dim\_Dates, Dim\_Products, Dim\_Stores) maps naturally to dbt's dimensional modeling approach. Business logic currently embedded in Target cubes is extracted and codified as version-controlled SQL.
- Lightdash** – open-source BI tool that reads metric definitions directly from the dbt semantic layer. Supports metric trees, rich visualizations, and self-service exploration. Actively developed with a clear product roadmap.

## INDICATIVE TIMELINE

- Phase 1: Discovery** Q2 2026  
 Data source audit, architecture blueprint, migration roadmap, stakeholder sessions. Fixed scope, 2--3 weeks.
- Phase 2: Platform Build** Q2--Q3 2026  
 Cloud warehouse setup, ingestion pipelines, orchestration, CI/CD, dbt transformation layer. Scoped in Phase 1.
- Phase 3: Analytics & BI** Q3 2026  
 KPI migration, Lightdash dashboards, stakeholder validation, parallel-run with legacy. Scoped in Phase 1.
- Advanced Capabilities** Q4 2026+  
 AI/ML capabilities, predictive analytics, external data integration, loyalty scheme analytics. Defined based on Phase 3 outcomes.

## INVESTMENT

Our daily rate is **EUR 1,210** (1 day = 8 hours).

### PHASE 1: DISCOVERY & ARCHITECTURE DESIGN (FIXED PRICE)

Deliverable	Est. Days
Data source audit & current-state analysis	3
Architecture blueprint & technology selection	2
Migration roadmap & cost projection	2
Stakeholder sessions (2–3 sessions)	2
Documentation & final deliverable	1
<b>Total</b>	<b>10 days</b>

**Phase 1 -- Discovery & Architecture Design**

**EUR 12,100**

Phase 1 is a standalone deliverable. Gregory's can evaluate the blueprint and decide on next steps with no further commitment. Travel to Athens for on-site sessions, if required, is billed at cost.

### PHASES 2 & 3: IMPLEMENTATION

Phases 2 and 3 are scoped and priced based on Phase 1 findings. To provide an order-of-magnitude estimate:

Phase	Indicative Range
Phase 2: Data Platform Build	EUR 25,000 – 45,000
Phase 3: Analytics Engineering & BI Migration	EUR 20,000 – 35,000
<b>Indicative total (Phases 2 + 3)</b>	<b>EUR 45,000 – 80,000</b>

Ranges depend on the number of data sources, complexity of existing business logic, and depth of the BI migration. Phase 1 will produce a precise scope and cost for each phase. Delivery model (fixed price or time & materials) is agreed per phase.

All prices net (excl. VAT). Fixed-price items are billed upon completion. Time & materials items are billed monthly based on actual days worked. Travel expenses, if any, are billed at cost.

## PREREQUISITES

- **Source system access** – read access to SQL Server warehouse, Targit environment, and key source systems for the audit
- **Stakeholder availability** – 2–3 sessions with BI team and business stakeholders during Phase 1
- **Existing documentation** – any available data dictionaries, Targit cube definitions, or report catalogs
- **Cloud account** – for Phases 2+: cloud provider account (GCP, AWS, or Azure) with appropriate permissions

## WAYS OF WORKING

Area	Approach
<b>Communication</b>	Shared Slack channel or Microsoft Teams for day-to-day coordination
<b>Sync Cadence</b>	Weekly call to review progress and align priorities
<b>Tooling</b>	AI-assisted development (Cursor, Claude Code) as standard practice
<b>Code Review</b>	All work via pull requests with review before merge
<b>Documentation</b>	All deliverables documented to handoff-ready standard
<b>Stakeholder</b>	Petros Tsongas as primary point of contact

## BEYOND THE MIGRATION

Once the core platform is in place and the BI migration is complete, the modern data stack unlocks capabilities that were not feasible on the legacy infrastructure:

### AI & PREDICTIVE ANALYTICS

- Demand forecasting per store and product category
- Anomaly detection on sales patterns and inventory
- Customer segmentation and lifetime value modeling (once loyalty data is available)

### LOYALTY SCHEME ANALYTICS

- Gregory's consumer loyalty programme (currently under development) can be natively integrated into the data platform from day one
- Customer profiles, purchase history, campaign responses, and engagement metrics flowing into the same warehouse as sales and operational data
- Enables cross-analysis: loyalty impact on ticket size, repeat visit frequency, campaign ROI

### REAL-TIME & OPERATIONAL ANALYTICS

- Near-real-time dashboards for store operations and delivery performance
- Alerting on KPI thresholds (e.g. store underperformance, delivery SLA breaches)

### EXTERNAL DATA INTEGRATION

- Combining warehouse data with public and online sources (weather, events, foot traffic)
- Competitive intelligence and market benchmarking

### AGENTIC AI & SELF-SERVICE

- AI agents that answer business questions directly from the warehouse in natural language
- Automated trend summaries and proactive alerts
- Dashboard generation from natural-language requirements

## ABOUT GEMMA ANALYTICS

Gemma Analytics is a Berlin-based data consultancy specializing in modern data platform implementation. We work with growth-stage companies to build production-grade data infrastructure – from warehouse setup and pipeline engineering through to KPI dashboards and team enablement.

Our team brings deep experience across the modern data stack, with 70+ completed data platform projects. We operate code-first and leverage AI-assisted tooling to deliver at speed without compromising on quality or maintainability.

Over these projects we have built extensive internal documentation, reusable patterns, and engineering best practices that carry over to every new engagement – accelerating development significantly from day one.

# **GEMMA ANALYTICS**

Gemma Analytics GmbH – Chausseestraße 17, 10115 Berlin  
[gemmaanalytics.com](https://gemmaanalytics.com)