## From Fragmented Reports to Financial Intelligence

# Designing the Site Financial Health Engine

Financial organizations rarely suffer from a lack of data. They suffer from fragmented data, inconsistent reporting processes, and an inability to transform financial information into actionable insights quickly enough to support operational decisions.

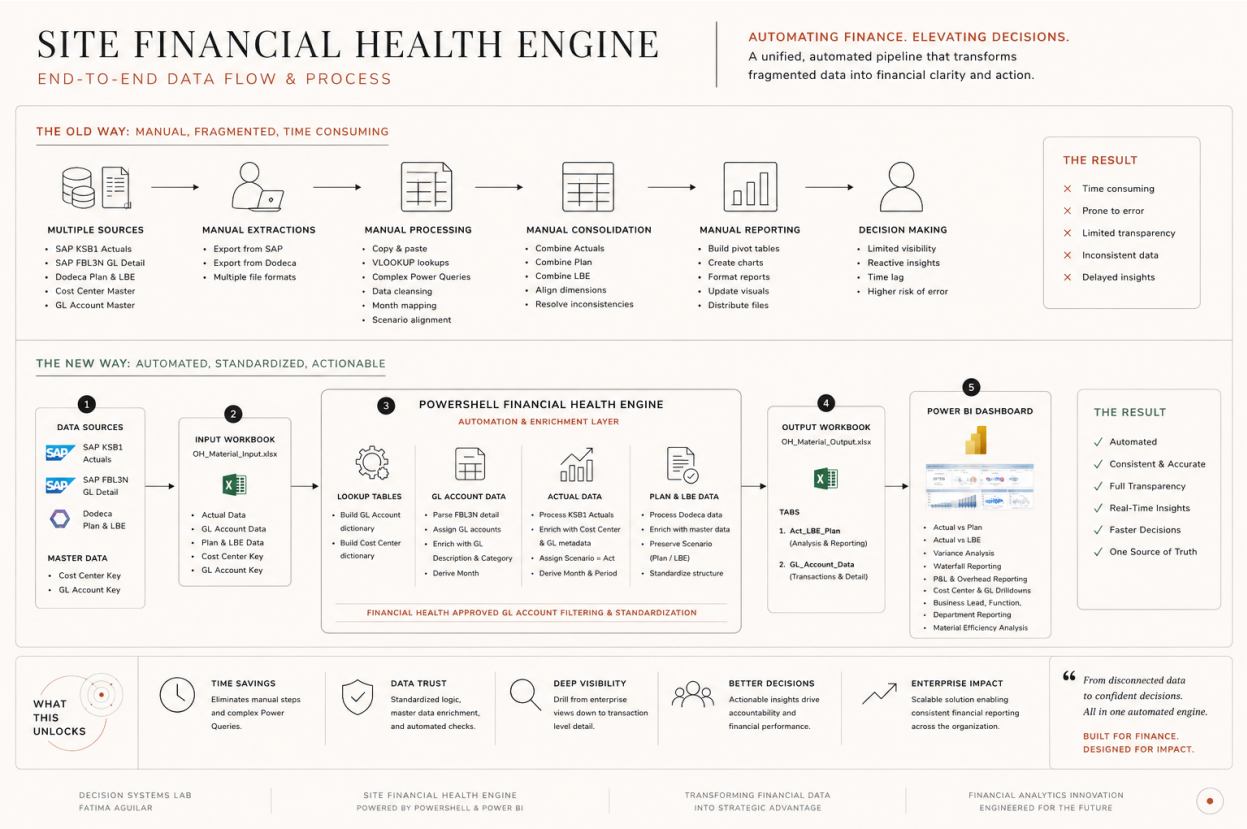

# 1. Introduction

Financial organizations rarely struggle with a lack of data. Instead, they struggle with fragmented information, inconsistent reporting processes, and the considerable effort required to transform raw financial records into meaningful insight. Month end reporting often depends on analysts manually collecting data from multiple enterprise systems, reconciling inconsistent formats, rebuilding business logic, and producing dashboards that explain results only after significant preparation.

At the site level, this challenge was especially evident. Financial reporting relied on data extracted from multiple independent sources, including SAP KSB1 Actuals, SAP FBL3N General Ledger detail, Dodeca Plan and Latest Business Estimate (LBE) reports, and multiple master data files, that each served a different purpose and followed different structures. Before any meaningful analysis could begin, these datasets had to be manually combined, enriched, validated, and standardized into a common format.

While each individual task was manageable, the reporting process as a whole had evolved into a fragmented workflow. Business logic was duplicated across spreadsheets, lookup tables were recreated repeatedly, Power Query transformations became increasingly complex to maintain, and visualizations were often built separately from the data preparation process. As reporting requirements grew, the amount of effort spent preparing data began to exceed the time available for interpreting it.

The Site Financial Health Engine was developed to redesign this workflow from the ground up. Rather than automating isolated spreadsheet tasks, the project established a standardized financial data pipeline that transforms multiple enterprise data sources into a single, validated analytical model. Built using PowerShell, Microsoft Excel, and Power BI, the engine automatically enriches, standardizes, and consolidates Actual, Plan, and Latest Business Estimate (LBE) data into a unified dataset that serves as the single source of truth for financial reporting and analysis.

The objective of the system extends beyond automation. Its purpose is to reduce the cognitive burden of financial reporting by ensuring that analysts spend less time locating, cleaning, and validating data, and more time understanding financial performance, investigating variances, and supporting strategic decisions. By treating financial reporting as an engineering problem rather than a

collection of spreadsheets, the Site Financial Health Engine establishes a scalable foundation for transparent, repeatable, and trustworthy financial decision support.

# 2. The Legacy Financial Reporting Process

Before the development of the Site Financial Health Engine, financial reporting relied on a sequence of independent manual processes performed across multiple applications and spreadsheets. While each step fulfilled a necessary business function, the workflow lacked a standardized architecture capable of consistently transforming enterprise financial data into reliable analytical outputs.

The reporting process began by extracting information from several independent systems. Actual financial performance was retrieved from SAP transaction KSB1, General Ledger activity from SAP FBL3N, planning information from Dodeca, and supporting metadata from multiple master data files. Because each source was designed independently, differences in structure, naming conventions, account hierarchies, and reporting dimensions required substantial manual reconciliation before analysis could begin.

Once extracted, these datasets were imported into Microsoft Excel where analysts manually performed lookups, mapped cost centers, aligned General Ledger accounts, standardized dimensions, cleaned inconsistent records, and reconciled reporting periods. These transformations were implemented through a combination of Excel formulas, Power Query workflows, and repeated manual validation.

Although Power Query significantly reduced repetitive spreadsheet work, the growing number of transformation steps created increasingly large and interconnected query structures. As additional reporting requirements were introduced, maintaining these workflows became progressively more difficult. Small changes to business rules frequently required modifications across multiple queries, increasing maintenance effort and reducing transparency.

After the data preparation stage, Actual, Plan, and Latest Business Estimate (LBE) datasets were manually consolidated into reporting workbooks. Pivot tables and charts were then created to support financial reviews, variance analysis, and management presentations. Because visualization existed as a separate step from data preparation, any changes to the underlying data often required rebuilding reports or manually refreshing multiple dependent components.

This fragmented workflow introduced several operational limitations. Data preparation consumed a significant portion of the reporting cycle, reporting logic became distributed across numerous spreadsheets and Power Query transformations, and visibility into the complete data pipeline was limited. Rather than providing a repeatable analytical system, the process depended heavily on analyst expertise and manual intervention to produce accurate financial reports.

These limitations highlighted that the primary challenge was not the availability of financial data, but the absence of a unified processing architecture capable of transforming disparate enterprise information into standardized, decision-ready intelligence.

# 3. Designing the Site Financial Health Engine

The Site Financial Health Engine was designed to address a fundamental limitation of traditional financial reporting workflows: the absence of a standardized processing layer between enterprise financial systems and business intelligence tools.

Rather than improving isolated reporting tasks, the objective was to redesign the entire financial data pipeline into a repeatable decision support system. Every stage of data preparation, enrichment, validation, and reporting was consolidated into a single automated architecture capable of producing consistent analytical outputs from multiple enterprise data sources.

Several design principles guided the development of the system.

**Automation First.** Every transformation that had previously required manual intervention, including data imports, account mapping, enrichment, filtering, and consolidation, was redesigned to execute automatically through code. This reduced repetitive analyst work while improving consistency across reporting cycles.

**Single Source of Truth.** Instead of distributing business logic across numerous Excel formulas and Power Query transformations, all financial rules were centralized within a single processing engine. This ensured that every dashboard, workbook, and analysis referenced the same standardized data model.

**Modular Processing.** Each financial dataset, General Ledger transactions, Actuals, Planning data, and Master Data, was processed independently before

being integrated into a unified output. This modular architecture simplified maintenance, improved scalability, and allowed individual components to evolve without requiring redesign of the entire system.

**Business Intelligence by Design.** Rather than treating dashboards as a separate reporting activity, the engine was designed to produce outputs already optimized for Power BI. Data preparation and visualization became components of the same architecture, allowing dashboards to function as a direct extension of the processing pipeline instead of an independent reporting layer.

**Transparency and Traceability.** Every transformation applied to the data follows a deterministic sequence that can be reviewed, validated, and reproduced. This increases confidence in financial reporting while making the processing logic significantly easier to understand than large interconnected Power Query workflows.

These principles transformed the reporting process from a collection of spreadsheet operations into an engineered financial intelligence platform. Instead of simply automating existing work, the Site Financial Health Engine introduced a standardized architecture that continuously converts fragmented financial information into reliable, decision-ready data.

# 4. System Architecture

The Site Financial Health Engine is organized as a modular financial processing pipeline that transforms raw enterprise financial data into standardized analytical outputs. Rather than relying on a collection of independent spreadsheets and disconnected reporting files, every stage of the workflow is executed within a single, deterministic architecture.

The system consists of five cooperating layers, each responsible for a distinct stage of financial data transformation.

## 4.1 Data Acquisition

The pipeline begins by collecting financial information from multiple enterprise systems. Actual financial activity is extracted from SAP, planning data is imported from Dodeca, and organizational metadata is retrieved from master data repositories. Each source contributes a different component of the organization's financial picture.

Instead of attempting to reconcile these datasets during reporting, the engine first preserves each source independently before applying standardized processing rules.

## 4.2 Data Standardization

Once imported, each dataset is normalized into a consistent structure.

General Ledger accounts are categorized and enriched with reporting metadata. Cost centers are matched to organizational hierarchies. Planning data is standardized to align with Actual financial activity. Common reporting dimensions, including fiscal periods, scenarios, departments, and business units, are generated automatically.

This stage eliminates the repetitive lookup tables, manual mappings, and spreadsheet transformations that previously consumed much of the reporting cycle.

## 4.3 Financial Processing Engine

The standardized datasets enter the core PowerShell processing engine.

Here, business rules are executed to filter approved General Ledger accounts, combine Actuals with Plan and Latest Business Estimate (LBE) data, generate reporting ready financial structures, and validate the integrity of every dataset before publication.

Unlike traditional Power Query workflows, every transformation is executed as code within a controlled sequence, making the processing logic easier to maintain, review, and extend as reporting requirements evolve.

## 4.4 Output Generation

After processing is complete, the engine automatically produces standardized output workbooks designed specifically for downstream reporting.

Rather than creating numerous independent Excel reports, the engine generates a small number of structured datasets that serve as the organization's financial reporting foundation.

These outputs become the single source of truth consumed by business intelligence tools and financial analysts alike.

### 4.5 Business Intelligence Layer

The final layer presents the processed financial information through interactive Power BI dashboards.

Because the data has already been standardized and validated upstream, dashboard development focuses entirely on analysis rather than data preparation. Financial leaders can immediately explore Actual versus Plan performance, variance analysis, departmental spending, cost center performance, General Ledger activity, and executive KPIs without rebuilding calculations for each report.

The dashboard is no longer another spreadsheet, it becomes the visualization layer of an engineered financial information system.

# 5. End-to-End Data Processing Pipeline

The Site Financial Health Engine transforms multiple independent financial data sources into a single analytical model through a deterministic processing pipeline. Rather than combining datasets manually during each reporting cycle, every transformation is executed automatically in a predefined sequence that standardizes, enriches, validates, and consolidates enterprise financial information.

The process begins with data acquisition from three primary enterprise sources. Actual financial activity is extracted from SAP KSB1, providing cost center expenditures by posting period. Detailed transaction-level information is imported from SAP FBL3N, supplying General Ledger activity used for transaction investigations and root-cause analysis. Planning and Latest Business Estimate (LBE) data are retrieved from Dodeca, providing forecast information that supports variance analysis and financial planning. These datasets are complemented by Cost Center and General Ledger master data used throughout the enrichment process.

Once imported into the Financial Health Input Workbook, the PowerShell engine initializes a series of lookup dictionaries using the Cost Center and General Ledger reference tables. These lookup structures provide the organizational context required to enrich every financial record with standardized business metadata, including Business Lead, Function, Cost Center Description, General Ledger Description, and reporting Category. By constructing these reference tables once at the beginning of execution, the engine eliminates thousands of repetitive lookup operations that would otherwise occur throughout the workflow.

Each financial dataset is then processed independently according to its business purpose. General Ledger transaction data is parsed to identify account headers, associate transaction records with the correct General Ledger accounts, enrich each record with descriptive metadata, derive reporting months, and produce a standardized transaction dataset. Actual financial data is simultaneously enriched with Cost Center and General Ledger information, assigned the "Actual" reporting scenario, filtered to approved Financial Health General Ledger accounts, and prepared for consolidation. Planning and LBE datasets undergo a parallel transformation process that standardizes their structure to match the Actual dataset while preserving their original reporting scenarios.
After each dataset has been independently validated and standardized, the engine consolidates Actual, Plan, and Latest Business Estimate records into a unified financial model. Every record now shares a common schema containing standardized dimensions such as Cost Center, Business Lead, Function, General Ledger Account, Category, Scenario, Period, Month, and Amount. This unified structure becomes the analytical foundation for all downstream reporting and ensures that every financial comparison is performed using consistent business logic.

The final stage of the pipeline automatically exports two reporting-ready datasets. The Act_LBE_Plan dataset supports executive financial reporting, variance analysis, Actual versus Plan comparisons, cost center reporting, and organizational performance monitoring. The GL_Account_Data dataset provides transaction-level visibility for detailed investigations, material efficiency analysis, spending trends, and root cause analysis. These outputs serve as the direct data source for the Site Financial Health Power BI dashboard, allowing every visualization to refresh from a single standardized source of truth without additional manual preparation.

By separating data acquisition, enrichment, transformation, validation, consolidation, and visualization into clearly defined processing stages, the Site Financial Health Engine converts what was once a fragmented reporting workflow into a repeatable financial information pipeline. The result is a scalable architecture that minimizes manual intervention while improving transparency, consistency, and confidence in financial decision support.

# 6. Business Intelligence Layer

The final layer of the Site Financial Health Engine is the Business Intelligence environment, where standardized financial data is transformed into interactive analytical dashboards that support operational and strategic decision making.

Unlike traditional reporting workflows, where dashboard development requires significant data preparation and manual calculations, the Financial Health Engine separates data engineering from visualization. All enrichment, validation, filtering, and consolidation occur upstream within the processing pipeline, allowing Power BI to function solely as the presentation layer of the system.

This architectural separation significantly simplifies dashboard development. Rather than recreating business logic inside reports, Power BI consumes a standardized dataset that has already been validated and structured by the Financial Health Engine. Every visualization, metric, and drill through analysis therefore references the same underlying financial model, ensuring consistency across all reporting views.

The dashboard provides multiple analytical perspectives tailored to different levels of financial decision making. Executive users can monitor overall financial performance through Actual versus Plan and Actual versus Latest Business Estimate (LBE) comparisons, while department managers can investigate performance at the Cost Center, Business Lead, Function, and General Ledger levels. Transaction level drill downs enable analysts to move seamlessly from high-level financial summaries to individual accounting entries without leaving the reporting environment.

One of the primary design objectives of the dashboard was to improve financial storytelling. Instead of presenting static pivot tables or isolated charts, the dashboard organizes financial information into a cohesive analytical experience that guides users from organizational performance indicators toward the underlying operational drivers of financial variance. Interactive filters, drill through capabilities, waterfall analyses, and trend visualizations allow financial questions to be answered dynamically rather than requiring multiple independent reports.

This shift fundamentally changes the role of the financial analyst. Rather than spending reporting cycles preparing data and rebuilding visualizations, analysts can focus on interpreting results, investigating root causes, communicating financial insights, and supporting business leaders with evidence based recommendations.

The Business Intelligence layer therefore represents more than a reporting interface. It serves as the decision support layer of the Site Financial Health Engine, transforming standardized financial information into actionable organizational intelligence while preserving a single source of truth throughout the entire reporting ecosystem.

# 7. Results and Business Impact

The implementation of the Site Financial Health Engine fundamentally changed the financial reporting workflow by replacing a fragmented collection of manual processes with a standardized, automated data pipeline. Rather than requiring analysts to repeatedly prepare, reconcile, and restructure financial information before analysis could begin, the system delivers reporting-ready datasets that support immediate exploration within Power BI.

The most immediate improvement was the standardization of financial reporting. Actual, Plan, and Latest Business Estimate (LBE) data are now processed using a single set of business rules, ensuring that every dashboard, variance analysis, and financial review is based on the same validated information. By centralizing transformation logic within the PowerShell engine, the system also eliminates inconsistencies that can arise when business rules are distributed across multiple spreadsheets and Power Query workflows.

The architecture also significantly improves maintainability. Because the processing logic is organized into modular components, modifications to General Ledger mappings, Cost Center hierarchies, reporting categories, or business rules can be implemented within the processing engine rather than manually replicated throughout multiple reporting files. This reduces maintenance effort while improving transparency and auditability.

From an analytical perspective, the Financial Health Engine shifts the role of financial reporting from data preparation to decision support. Financial analysts can spend more time investigating variances, identifying operational drivers of performance, and communicating actionable insights instead of performing repetitive data transformations. Executive stakeholders benefit from a consistent reporting environment that provides immediate visibility into organizational performance across Cost Centers, Business Leads, Functions, General Ledger Accounts, and reporting scenarios.

Perhaps the most significant outcome of the project is not the automation itself, but the establishment of a reusable financial reporting infrastructure. Future dashboards, analyses, and reporting requirements can now be built upon an existing standardized data model rather than requiring independent reporting solutions. As financial reporting needs continue to evolve, the Site Financial Health Engine provides a scalable foundation capable of supporting increasingly sophisticated analytical capabilities while preserving consistency, reliability, and trust in the underlying financial information.

The project also demonstrated the organizational value of treating financial reporting as an engineering discipline. By redesigning the reporting process as an integrated information system rather than a collection of spreadsheets, the Financial Health Engine establishes a framework that enables faster reporting cycles, greater confidence in financial data, and more informed operational decision making across the organization.

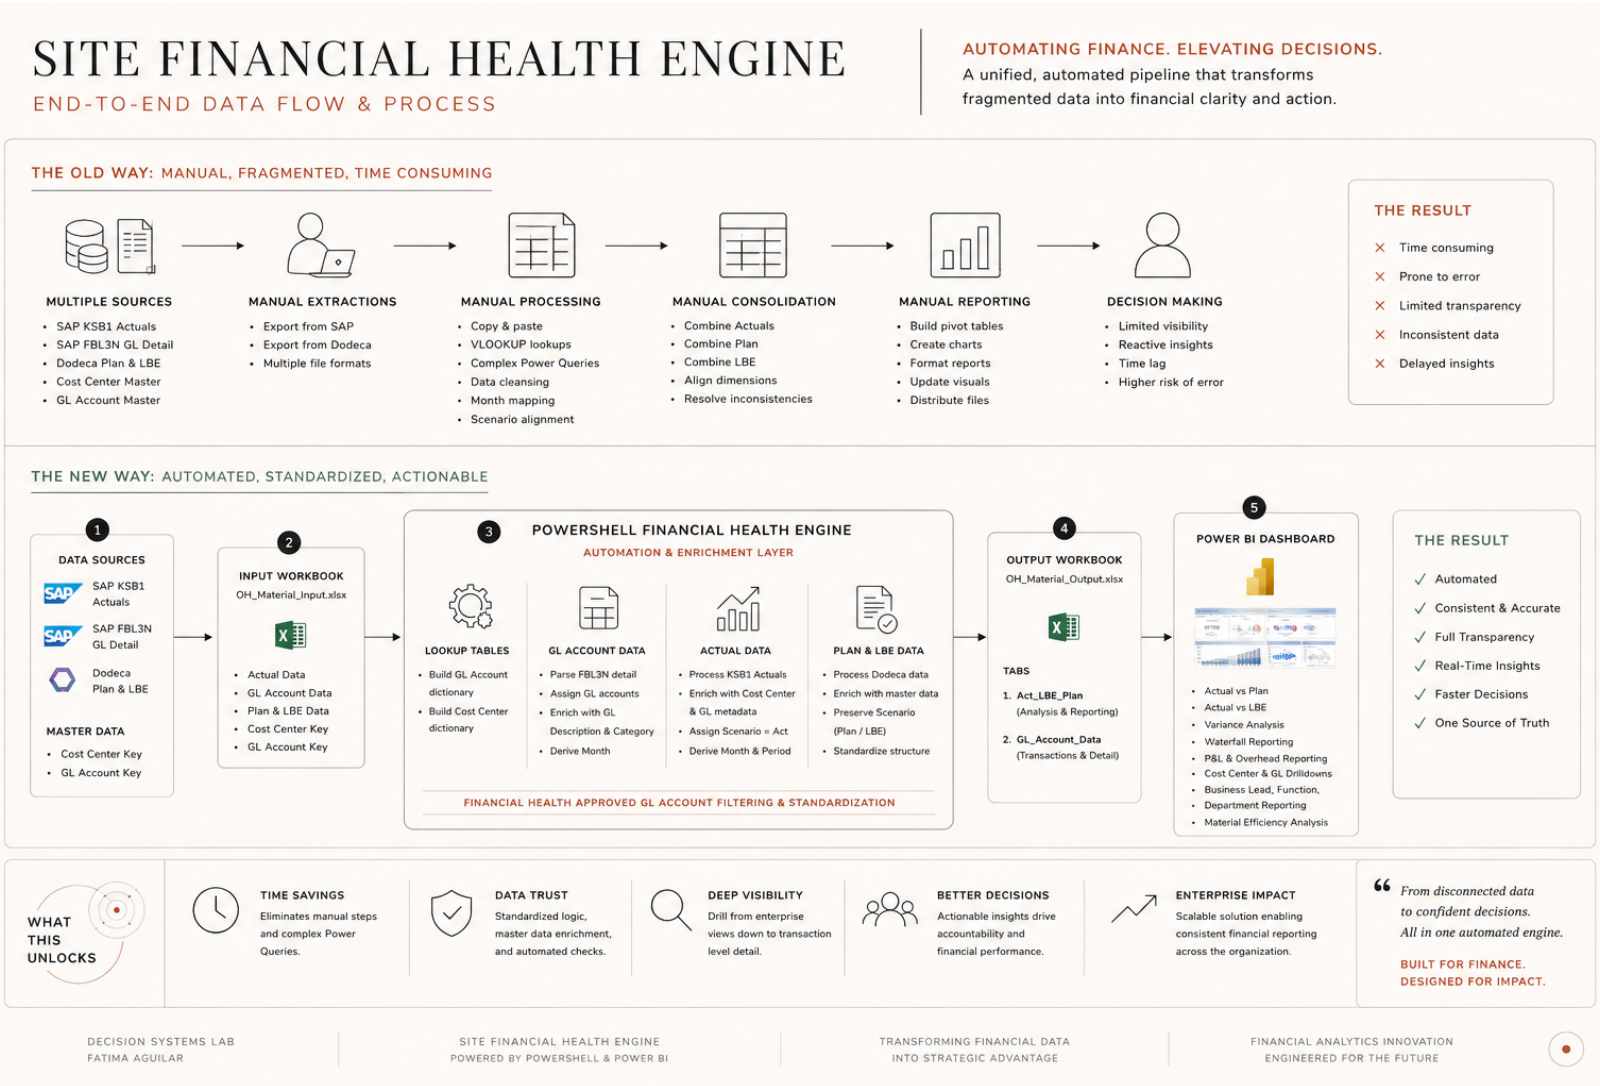

# 8. Lessons Learned

The development of the Site Financial Health Engine reinforced several engineering principles that extend beyond financial reporting. Although the project originated as an effort to improve month-end reporting, the process ultimately demonstrated that many reporting challenges are architectural problems rather than analytical ones.

## Build the Data Pipeline Before the Dashboard

One of the most important lessons learned was that high-quality dashboards are the product of high-quality data engineering. Interactive visualizations alone cannot compensate for inconsistent business rules, fragmented datasets, or repetitive manual transformations. By investing in a standardized processing pipeline first, dashboard development became significantly simpler, more maintainable, and more reliable.

## Standardization Creates Scalability

As reporting requirements increased, it became clear that maintaining numerous independent Power Query transformations and spreadsheets would not scale effectively. Centralizing business logic within a single processing engine allowed new reporting requirements to be incorporated without redesigning the entire reporting workflow. Standardization transformed financial reporting from a collection of isolated reports into a reusable analytical platform.

## Financial Reporting is an Engineering Problem

Traditional financial reporting often evolves organically through spreadsheets, formulas, and manual processes. While effective for small-scale analyses, these approaches become increasingly difficult to maintain as organizational complexity grows. Applying software engineering principles—including modular architecture, deterministic processing, reusable components, and separation of responsibilities—produced a reporting system that is significantly more transparent, maintainable, and adaptable than the legacy workflow.

## Automation Should Elevate Human Work

The objective of the Financial Health Engine was never to replace financial analysts. Instead, automation was used to eliminate repetitive preparation tasks so analysts could focus on higher-value activities such as variance investigations,

financial storytelling, operational decision support, and strategic planning. The greatest value of automation lies not in removing people from the process, but in allowing them to spend more time exercising professional judgment.

### Decision Support Begins with Trusted Data

Perhaps the most important lesson from this project is that effective decision-making depends on confidence in the underlying information. Standardized processing, consistent business logic, and transparent transformation pipelines create trust in financial data. Once that trust exists, dashboards become more than reporting tools—they become decision-support systems capable of guiding meaningful organizational action.

## 9. Future Work

The Site Financial Health Engine establishes a standardized financial reporting architecture, but its greatest potential lies in serving as the foundation for future intelligent decision support systems.

By consolidating enterprise financial information into a single, validated analytical model, the engine creates a structured data environment capable of supporting increasingly sophisticated analytical capabilities. Because the underlying processing pipeline already standardizes Actual, Plan, and Latest Business Estimate (LBE) data, future enhancements can focus on generating insights rather than preparing information.

Several opportunities for future development have been identified.

Predictive analytics could extend the system beyond descriptive reporting by forecasting cost center performance, identifying emerging financial trends, and proactively highlighting areas requiring management attention. Rather than waiting for month-end reviews, financial leaders could receive early indicators of operational risk based on historical spending patterns and forecast deviations.

Machine learning techniques could further improve financial monitoring by automatically detecting anomalous transactions, identifying unusual spending behavior, and recognizing patterns that may not be immediately visible through traditional variance analysis. These capabilities would augment existing financial reviews while preserving analyst oversight and judgment.

Large Language Models (LLMs) also present an opportunity to transform the interpretation of financial information. Instead of requiring analysts to manually explain variances and summarize financial performance, future versions of the

system could automatically generate executive-ready financial narratives, identify likely drivers of organizational performance, and recommend areas for further investigation. These capabilities would complement, rather than replace, the expertise of financial professionals by accelerating routine analysis and improving communication.

Looking further ahead, the architecture could evolve from a financial reporting platform into a comprehensive Financial Decision Support System. By integrating predictive modeling, intelligent recommendations, automated financial narratives, and interactive decision-support capabilities, the Site Financial Health Engine has the potential to become an active participant in organizational decision-making rather than a passive reporting tool.

Ultimately, the long-term vision extends beyond financial automation. It is to engineer systems that reduce the effort required to transform enterprise data into actionable organizational intelligence. As organizations continue to generate increasing volumes of financial information, the ability to build transparent, explainable, and trustworthy decision systems will become an essential capability for modern finance.

The Site Financial Health Engine represents one step toward that future. Rather than viewing financial reporting as the final objective, it demonstrates how engineering principles can establish the infrastructure upon which the next generation of intelligent decision systems can be built.