SEQUOIA

rise above

WHITEPAPER

Building Your Data Foundation

Why Fundraising Organisations Need a Data Warehouse

Executive Summary

Modern fundraising organisations are increasingly turning to business intelligence tools like Power BI, Tableau, and Looker to gain insight from their data. However, many organisations make a critical mistake: connecting their BI tools directly to their CRM or source systems without an intermediate data layer.

DID YOU KNOW?

The global data warehousing market was valued at approximately $34.5 billion in 2024, yet charities are still primarily focused on basic data management and face resource challenges in leveraging sophisticated data warehousing solutions.

This whitepaper explores the increasing accessibility of data warehouse technology and why investing in a well-designed data warehouse between your CRM and your reporting tools is not merely a technical nicety but a strategic imperative.

Faster Performance
Reports run in seconds, not minutes

Cleaner Data
Trustworthy, consistent information

Reduced Complexity
Simplified BI tool configuration

AI-Ready Foundation
Support for advanced analytics

01

The Challenge

Understanding the reporting problem facing fundraising organisations

Introduction: The Reporting Challenge

Every fundraising organisation faces the same fundamental challenge: turning raw data into actionable insight. Your CRM holds a wealth of information about donors, gifts, campaigns and relationships. You may have separate systems that capture email or social data, but raw data alone does not drive decisions.

To be truly data-driven, organisations need to transform data into reports, dashboards, and analysis that can be used to find the insights that matter.

Business intelligence tools have revolutionised what is possible. Power BI, Tableau and their competitors offer powerful visualisation capabilities, interactive dashboards, and the ability to share insight across the organisation. But the tool itself is only part of the equation. The quality of your reports depends fundamentally on the quality and structure of the data feeding them.

This is where many organisations struggle. Connecting a BI tool directly to a CRM might seem like the simplest approach, but it creates problems that compound over time. Reports become slow. Data quality issues proliferate. Analysts spend more time wrestling with data than generating insight and when the charity wants to do something more sophisticated, like tracking active file size over time or building predictive models, the limitations become even more apparent.

KEY INSIGHT

The solution is a dedicated data warehouse: an intermediate layer specifically designed to support reporting and analytics.

Roche's Maxim: A Guiding Principle

“Data should be transformed as far upstream as possible and as far downstream as necessary."

When thinking about where data transformation should happen, we refer back to Matthew Roche's maxim. This principle captures a crucial insight about data architecture.

> Every transformation you perform in your BI tool is a transformation that runs every time someone opens a report. In Power BI it adds complexity to your DAX. It makes reports harder to maintain and debug and critically, it means the same logic might be implemented differently across different reports.

By contrast, transformations performed upstream, in a data warehouse, are done once. They create a clean, consistent dataset that all reports and analysis can use. The logic is documented in one place. Changes propagate automatically to all downstream reports. And your BI tool can focus on what it does best: visualisation and interactivity.

The second part of the maxim acknowledges that some transformations must happen downstream.

“some transformations must happen downstream…
But these should be the exception, not the rule.”

Context-dependent filtering, user-driven parameters, and real-time calculations that depend on the specific view being requested belong in the BI layer. But these should be the exception, not the rule.

02

The Architecture

Understanding the four layers of a well-designed data foundation

The Four Layers of Data

A well-designed data architecture typically comprises four distinct layers, each with a specific purpose. Over the years these layers have been given numerous marketing labels, like a medallion of bronze, silver, gold. However, they boil down to a similar conceptual purpose.

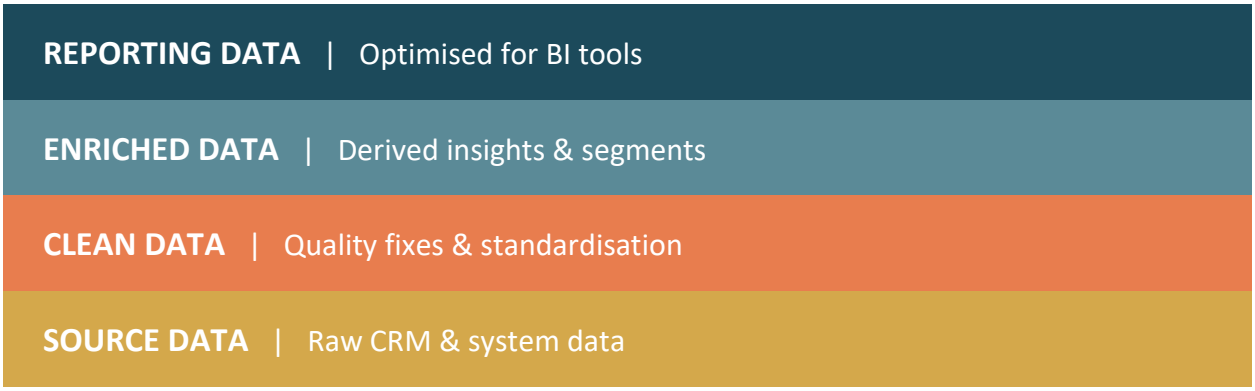

01 Source Data

This is your raw data as it exists in CRMs, finance systems, digital platforms, and other operational systems. Source data is extracted and loaded into your data warehouse, typically with minimal or no transformation. This preserves a record of what the source systems contained at the point of extraction, allowing for future reprocessing and auditing.

External Data and a Data Warehouse: The Perfect Match

A data warehouse is your best friend when dealing with external data sources because it pulls everything like APIs, and other systems into one clean, organised place. Instead of cramming external data into your CRM (where it might not even fit properly or eat up expensive storage), you get a purpose-built environment that handles the messy work of standardising and storing it all. This saves you from burning dev time on custom CRM integrations that weren't designed for bulk external data anyway. Your team gets reliable, historical records of everything without the headache of managing multiple connections or worrying about data disappearing from third-party sources.

Bottom line: it turns external data chaos into something actually useful, often at a fraction of the cost.

02 Clean Data

Source systems are rarely as clean as we would like. Regular giving cancellation dates might be incorrect or missing. Product categories might be inconsistent. Data entry errors create anomalies. This clean layer addresses these issues systematically. Here, you fix data quality problems, standardise formats, handle missing values, and ensure consistency.

> KEY INSIGHT
>
> **When you discover a new data quality issue, you fix it in one place, and the fix flows through to every report.**

03 Enriched Data

This is where the real value creation happens. The enriched layer derives new information from your clean data: recruitment source attribution, RFV (Recency, Frequency, Value) scores, donor segments, lifetime value calculations, churn risk indicators, and more.

04 Reporting Data

The final layer is specifically optimised for reporting. Here, data is structured into the star schema that BI tools work with most efficiently. Tables and columns are named in human-readable language. The result is a clean, performant, user-friendly dataset that makes building reports straightforward.

03

The Benefits

Why investing in a data warehouse pays dividends

The Case for a Data Warehouse

So why invest in this intermediate layer rather than building reports directly from your CRM?

1

Performance

CRMs are optimised for operational use, not analytical queries. A report that needs to aggregate millions of transactions while joining across multiple tables can bring a CRM to its knees. A data warehouse is purpose-built for these analytical workloads.

2

Consistency

When definitions and calculations live in the data warehouse, they are consistent across all reports. The definition of an "active regular giver" is the same whether you're looking at a board KPI dashboard or a detailed operational report.

3

Maintainability

Complex DAX formulas and calculated fields in BI tools are notoriously difficult to maintain. By contrast, logic in a data warehouse is centralised and documented. Changes are made once and propagate automatically.

4

Flexibility

A well-designed data warehouse enables analyses that would be impractical to build directly from source data. Tracking active file size over time, building cohort analyses, calculating rolling averages all become straightforward.

5

Future-Proofing

Your BI tool might change. Your CRM might be replaced. But a well-designed data warehouse provides continuity and the foundation for advanced analytics including machine learning and AI.

DID YOU KNOW?

31% of charities rate themselves as poor or not engaging with managing and using data. Using an agency to make sense of their data can be a quick way to boost data maturity.

04

Best Practices

Top tips for getting your warehouse right

Top Tips for Getting Your Warehouse Right

Based on our experience implementing data warehouses for fundraising organisations, here are our key recommendations:

1	**Build Shared Dimensions** Your warehouse should hold shared dimensions that can be reused: Date, Geodemographic, Geography, Product category tables and more. When built upstream, downstream tools will have common approaches for all to utilise.
2	**Define Business Logic Once** What makes a regular giver "active"? How do you calculate donor retention? These definitions should be implemented once, in your data warehouse, not repeated across multiple reports.
3	**Consistent Naming** Define a clear naming and schema convention at the start, and religiously adhere to it. Think about namings that guide users on how joins between tables occur.
4	**Document Everything** Maintain clear documentation of your data model: what each table contains, what each column means, how derived fields are calculated. This is invaluable when onboarding new team members.
5	**Invest in Data Quality** Build data quality checks into your warehouse processes. Monitor for anomalies. Alert when data falls outside expected ranges. Users quickly lose trust in reports that show obviously wrong numbers.
6	**Plan for Change** Business requirements evolve. Source systems change. Design your warehouse with change in mind. Use version control for your transformation logic.
7	**Governance and Access** Governance should provide freedom to allow data cultures to flourish whilst protecting sensitive data. Avoid being the bottleneck for data access.
8	**Think About Your Audience** Different users have different needs. Senior leaders want high-level dashboards. Analysts want flexibility to explore. Your data warehouse should support all these use cases.

05

Common Pitfalls

Mistakes to avoid when building your data warehouse

Common Pitfalls to Avoid

In our work with fundraising organisations, we see certain mistakes repeated frequently:

⚠ Building Reports Directly from the CRM

While technically possible, this approach leads to performance problems, inconsistent definitions, and overwhelming complexity in your BI tool. The short-term saving is quickly offset by ongoing pain.

⚠ Over-Engineering Too Early

It's tempting to build the perfect data warehouse before creating any reports. But perfect is the enemy of good. Start with a focused scope, build something useful, learn from it, and iterate.

⚠ Ignoring Data Quality

A data warehouse built on poor quality source data will produce poor quality reports. Invest time in understanding your data quality issues and addressing them systematically.

⚠ Neglecting Documentation

The person who builds your data warehouse will not always be available to answer questions. Documentation is essential for sustainability.

⚠ Duplicating Logic Across Layers

If you find yourself implementing the same calculation in both your data warehouse and your BI tool, something has gone wrong. This creates maintenance burden and risk of inconsistency.

06

Technology

Choosing the right platform for your organisation

Technology Considerations

The good news is that data warehouse technology has become increasingly accessible. Options include:

Platform	Best For
Azure Synapse, BigQuery, Redshift, Snowflake	Scalable, managed cloud solutions
Microsoft Fabric	Integrated Microsoft ecosystem
SQL Server	On-premises or existing MS licensing
dbt (data build tool)	Modular transformation management

Why Bring in an Agency to Help?

Building and maintaining a data warehouse isn't your charity's core mission. While the technology has become more accessible, getting it right still requires specialised skills in data architecture, ETL processes, and ongoing optimisation. Expertise that's expensive to hire and hard to retain in-house.

Your team is already stretched thin with imports, reporting and donor relationships, so adding building a data warehouse to someone's plate usually means it either doesn't get done properly or pulls critical resources away from mission-critical work.

Outsourcing to specialists means you get a professionally designed system from day one, with best practices baked in and someone else handling the technical headaches when things need updating or troubleshooting. You get to insights faster. Plus, working with experienced partners who've built warehouses for multiple organisations means you benefit from proven solutions rather than reinventing the wheel.

The right choice depends on your organisation's technical capabilities, existing infrastructure, budget, and scale. But the architectural principles outlined in this whitepaper apply regardless of which technology you choose.

07

Case Study

Sequoia Warehouse and Canopy in action

Case Study: Sequoia Warehouse and Canopy

The Challenge

Like many charities, this organisation was struggling to gain consistent, timely insight from its data. The CRM system was built for data entry, not analysis, leaving teams dependent on manual exports, disconnected spreadsheets, and slow reporting processes.

The Solution

The charity selected Sequoia's Owned Warehouse option combined with the Canopy Reporting Suite:

Microsoft Fabric
Modern data warehouse with full ownership

Automated Pipelines
Data flows from CRM automatically

Dimensional Model
Fundraising-optimised structure

Canopy Reports
Ready-made Power BI dashboards

The Results

8	1	100%
Weeks to Implement	Source of Truth	Ownership

The Canopy Reporting Suite provided immediate access to comprehensive dashboards covering income and fundraising KPIs, donor retention and acquisition metrics, channel and campaign performance, and supporter journey analysis including lifetime value and second gift rates.

> *The charity now has fast, consistent reporting capabilities that update automatically as new data flows in.*

Conclusion

Building your data foundation for the future

> A well-designed data warehouse is not a luxury; it is a strategic asset that underpins effective use of business intelligence tools.

Organisations that invest in this foundation find that their reporting becomes faster, more consistent, and more trustworthy. Analysts spend less time fighting with data and more time generating insight. And the platform supports sophisticated analyses, from cohort tracking to predictive modelling, that would be impractical without proper data architecture.

The principles are straightforward: transform data upstream, structure it for reporting, and keep your BI tool focused on visualisation and user interaction. The implementation requires investment, but the returns—in efficiency, consistency, and capability—justify that investment many times over.

KEY INSIGHT

Whether you are just beginning your BI journey or looking to improve an existing implementation, take the time to get your data foundation right. Your future self, and everyone who uses your reports, will thank you.

About Sequoia

Sequoia helps fundraising organisations unlock the value in their data. From data warehouse design and Power BI implementation to supporter insights and analysis, we bring deep sector expertise and technical excellence to every engagement.

Let's discuss how we can help your organisation build a better data foundation:

www.sequoia-insights.com

Tim McKenzie: tim.mckenzie@sequoia-insights.co.uk

Andrew Sargent: andrew.sargent@sequoia-insights.co.uk