

THOUGHT LEADERSHIP

Business continuity

How to replace core data systems
without critically impacting
business-as-usual



We are architects of systems.
We deliver an experience, not just technology.

Legacy system replacement can be a constant headache in large organisations. **Jason Michaelides** explains how building an enterprise data abstraction layer can remove much of the risk, protect business continuity and deliver other benefits besides.

As sure as night follows day, your important business systems will eventually need replacing. Whether you're upgrading to a newer software version, switching to a different vendor, or consolidating multiple systems into one, you'll need to ensure the transition is done in a way that protects business continuity.

This is especially challenging in larger organisations. Here, multiple business processes and applications will typically rely on – and be integrated with – the source system you're looking to replace. When you swap out the source system, each of these integrations will need to be painstakingly rebuilt to work with the new system. The number of integrations likely to be involved with an important business system means there's a high risk of something going wrong on cutover day.



But what if we said it was possible to eliminate much of that risk, by making the changeover essentially invisible to the consuming systems and business processes?



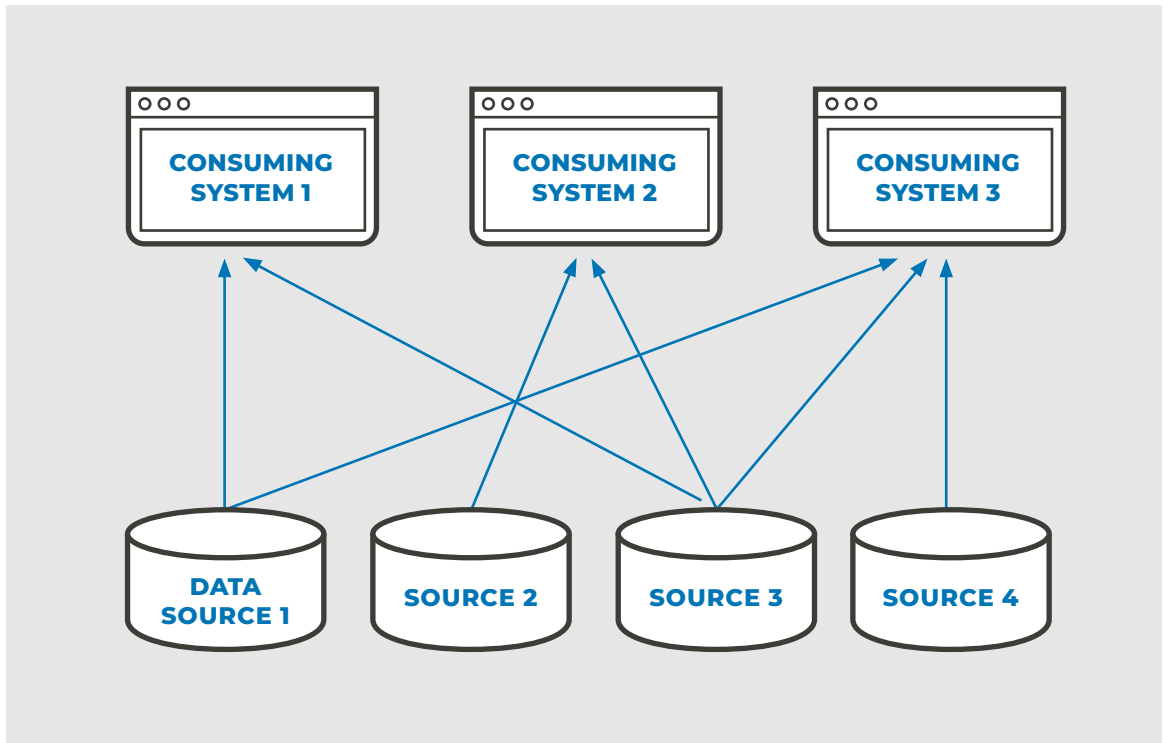


Figure 1: A typical current state in many organisations, with source data systems connected directly to one or more consuming systems.

When building an enterprise data layer makes sense

There are different ways of going about this, largely dependent on the scale of change you’re planning for. Below, we outline the approach we recommend when you’re looking to replace multiple source data systems, or very large systems.

In summary, the solution is to build an abstracted data layer that sits between the source data system(s) and the applications and processes that consume data from them. This decouples the consuming systems and processes from the source data systems, such that the data requests sent by the consuming systems are no longer sent direct to the source system(s), but to the abstraction layer, which then fetches the necessary information and passes it to the consuming application.

THE BENEFITS OF AN ABSTRACTED ENTERPRISE DATA LAYER

This has a number of advantages, not least that when you’re seeking to replace an end-of-life source data system, all your consuming applications and processes can continue running the same queries they do today, with the new data layer seamlessly handling the fact that the data is now coming from a different place.

In addition, having a layer between your source and consuming systems means there will typically only be one integration required

with each source data system, rather than many. This massively reduces risk and complexity when you come to replace any one of these systems. Equally, owners of the consuming systems only need to maintain a single integration (with the abstraction layer), regardless of where their source data is coming from.

The data layer can be either physical or virtual, and may be a bespoke build or use a commercial or open source product.

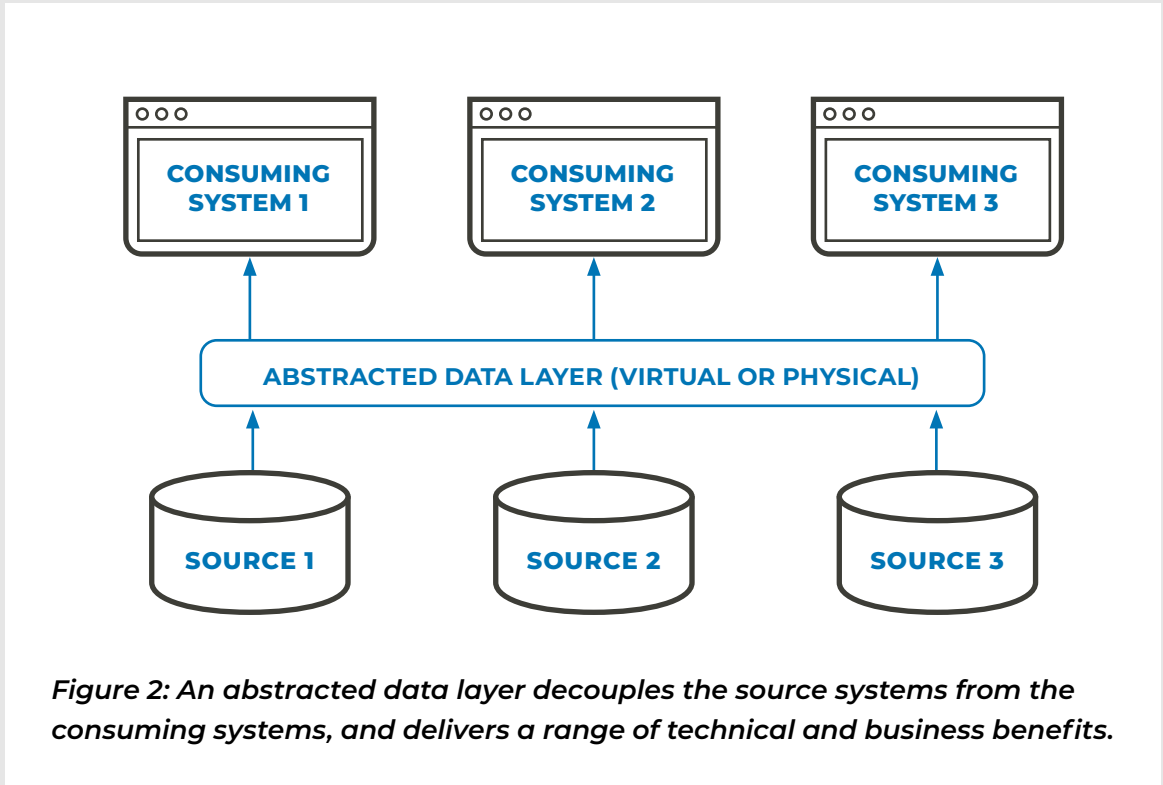


Figure 2: An abstracted data layer decouples the source systems from the consuming systems, and delivers a range of technical and business benefits.

How to build an abstracted enterprise data layer

To build an enterprise data layer to facilitate lower-risk and lower-complexity transitions from old to new data systems, your teams need to understand a number of things in detail.

Firstly, what data does each consuming system require, including its format, frequency and access method (such as a direct connection or an API)?

Secondly, what transformations need to take place to get data from the source systems into the necessary state? The data layer must recreate these, and they'll most probably need to be different for your old source data system(s) and your new one(s). This is essential if you're

“As well as assisting with legacy application replacement, an abstracted data layer can contribute to numerous other business benefits.”

Jason Michaelides, Calibrate Consulting

to ensure the output data is the same, and therefore that the changeover is invisible to the consuming systems.

Thirdly, you need a detailed understanding of permissions and regulatory requirements around the data in the source systems. Which consuming systems are allowed to access what? Does data need to be masked for certain users, but not others? Do certain datasets have data residency requirements?

Addressing these issues will require consistent, business-wide classification of all your data, which most organisations already have in place. The data layer needs to have the means of enforcing this correctly.

The breadth of domains we've just touched on demonstrates the importance of close collaboration between the various stakeholders, including source system owners, consuming application owners, regulatory compliance teams and the people building the data layer.



STEP-BY-STEP MIGRATION

Given we typically recommend this type of enterprise data layer to organisations that are seeking to migrate very large or multiple source data systems, we also usually recommend phasing the migration to further reduce risk.

One approach is the customer-by-customer technique, where you carefully stagger the move of your clients from the old source system to the new. You might want to start with your lowest-value or least strategic customers, for example, and iron out any teething troubles before transitioning your more strategic accounts. A customer-by-customer migration will require a control table in your data layer, to log which customers have been moved to the new source system. This will ensure that when data queries come in from a consuming system during the migration, the data layer always knows where to fetch the necessary data from.

THOUGHT LEADERSHIP

A LONG-TERM INVESTMENT



While we've focused primarily on the benefits of an enterprise data layer when it comes to legacy system replacement, the solution you put in place

can form the basis for numerous other advantages in the future. For example, when architected correctly, the data layer can become like a 'supermarket' for data, which helps people around the business find and use data to improve their products, services and efficiency.

It therefore represents an investment that has the power to repay your business many times over, if designed and built in the right way.



CONTACT US

✉ info@calibrateconsulting.com

🌐 www.calibrateconsulting.com

☎ +44(0) 3330 112844

🌐 [/calibrate-consulting](https://www.linkedin.com/company/calibrate-consulting)