WHITE PAPER BI ACCELERATION

How to get more from your BI tools

by optimizing your data stack

Exasol
The analytics database

Introduction

This guide is aimed at organizations that use one or multiple business intelligence (BI) tools. It's both relevant for end-users and those responsible for managing the tool's and configurations, as well as the architecture of the wider data stack. This guide illustrates how your BI tools can benefit from gradually adding a high-performance database to your technology environment.

Today's BI user has more control of their data stack than ever before.

First, the good news. Today's BI user has more control of their data stack than ever before. You can largely choose the data sources and services from the bottom up and rely on them offering a good deal of embedded integrations and automation. Make the right choices, and you end up with a faster, hassle-free and more economic data stack, powering supercharged BI.

But this can be a double-edged sword. In a busy marketplace, with technology vendors releasing all sorts of new services promising wonderful things, where can you find actual gains?

Understand the importance of the database layer

Let’s start with the database layer. It can be daunting if you’re just beginning to reappraise your current database layer and what’s needed to support your BI tool(s) going forward.

The database layer is likely to be responsible for more than just your BI. It could be the beating heart of your other data environments and data science platforms. So choosing a database and configuring it to maximum effect has far-reaching implications.

A new database can influence:

- The data sources you can pull in to your BI tool(s)
- The efficiency and speed with which your system completes Extract, Transform and Load (ETL) tasks
- The quantity and complexity of the queries you run.

Just as a suboptimal database can drain power and productivity from your data stack and leave any BI tool gasping, a high-performance database can lift your system to a new level.

Choice is great if you’re starting from scratch. But that’s unlikely for most organizations. The architecture of most enterprise-sized organizations is likely to be made up of years of tech procurement, start-stop adoption, inherited systems and complicated integrations.

Just mapping your current data stack might be challenging and that’s before you can even contemplate swapping out the fundamental database layer. Yet doing nothing presents more of a risk – it might result in your BI tool(s) falling short. And this means someone, somewhere, is making better and faster decisions than you.

So rather than regarding it as an enormously complicated upgrade, let’s look at it as a gradual journey and start small.

The architecture of most enterprise-sized organizations is likely to be made up of years of tech procurement, start-stop adoption, inherited systems and complicated integrations.

Step 1

The real-time, deep analytics offload – a quick win

You're probably already using one or multiple BI tools to run comprehensive, irregular, ad-hoc analyses, or it's on your radar to do so soon. For example, you might measure the effectiveness of a new sales campaign, analyze data from Internet Of Things devices to adjust parameters in your manufacturing, or run a credit check on a customer.

Here's where a live connection from any BI tool to a new high-performance database can help. Just move these tasks to your legacy data warehouse and continue to run the more predictable and standard analytics with the BI tools.

A new database should easily integrate with your current data warehouse, or use your existing ETL tools to bring in data straight from the source. The rest of your environment remains unchanged, while you test the combination of the BI platforms and the high-performance database.

The live connection to your new high-performance database is likely to deliver fast, independent, flexible and deep analyses to BI end-users. Instead of relying on pre-aggregated data, as a BI user you're now able to:

- **Combine multiple data sets**
- **Define the parameters important to you**
- **Discover new insights specific to your remit**

The live connection also leaves any BI platform to do what it does best. Data caches local to the BI tool are no longer stretched beyond the regular, standard analysis and reporting they've been designed for.

Your Database Administrators (DBAs) will have more time and inclination to focus on tasks that maximize value. Indeed, add in other time savings you achieve from managing a highly autonomous analytics database, and as a DBA you might expect to free up as much as 80% of your time.

This near real-time, deep analytics offload strategy introduces you to the benefits of adding a high-performance database to your data stack without requiring your existing, legacy data warehouse to be replaced.

The real-time, deep analytics offload – a quick win

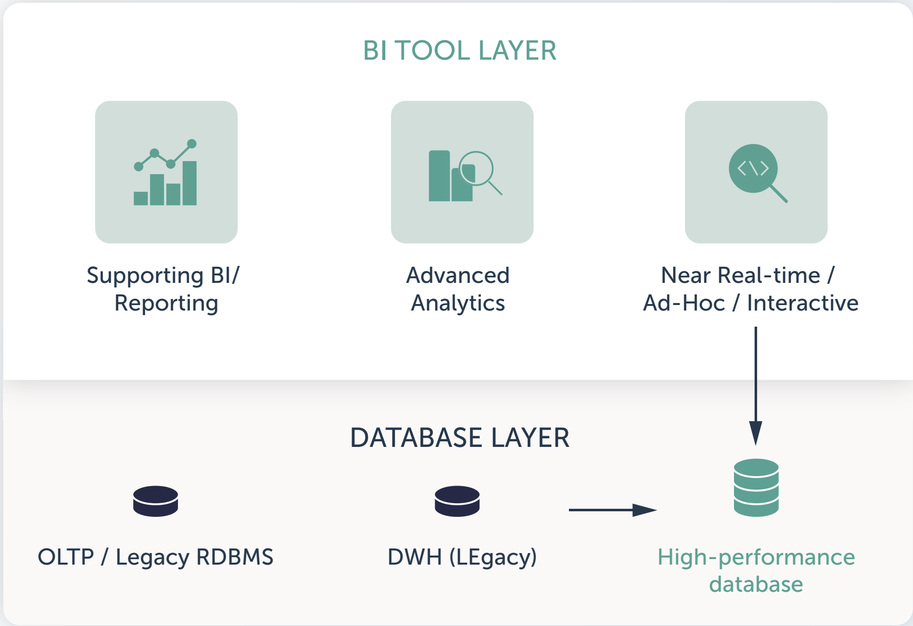

Exasol
The analytics database

Step 2

Extend the new database layer to benefit more from your BI tools

Once the new database is in your stack, it's up to you how fast you want to leverage further benefits. Advanced analytics databases are easily integrated with most other data software tools. It can be tempting to replace the databases and data warehouses in your legacy data layer in one go with your new high-performance database.

However, most enterprise environments are complex and sensible pragmatism means it's better to incrementally move workloads off your legacy data warehouse and onto the new database. Your new high-performance database will still power your BI platforms – from data caches and standard reporting through to the fast and furious world of discovery data science via live connections.

Advanced analytics databases are easily integrated with most other data software tools.

The real-time, deep analytics offload – a quick win

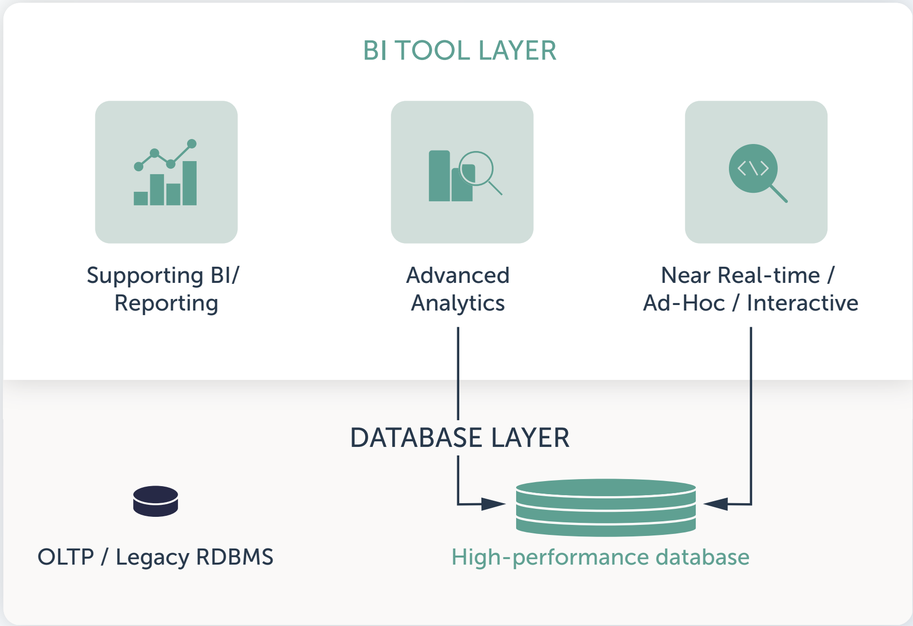

Exasol
The analytics database

Step 3

Create a less complex, more efficient stack from the bottom up

With a high-performance database layer at the center of your BI stack, there are many choices for implementation. You're most likely to base the stack on one of the leading cloud providers such as Microsoft Azure, Google Cloud or Amazon Web Services (AWS).

Alternatively you could opt for a multi-vendor approach, or a combination of private, public and even on-premises if you prefer (or simply require) a hybrid approach. More and more organizations understand how different hosting models impact cost, performance and security, so having a choice of infrastructures to deploy your BI tool(s) and the data stack on, has become an essential consideration.

You could also set your high-performance database between your BI platform and your data sources to leverage all the data your organization generates.

This includes:

- **Structured data of your enterprise resource planning (ERP), enterprise performance management (EPM) and customer relationship management (CRM) applications**
- **Data coming out of email, social media, sensors and log files**
- **Structured data generated by scripting languages such as Python via preprocessing inside the database**

A cleaner and more efficient data stack for your business

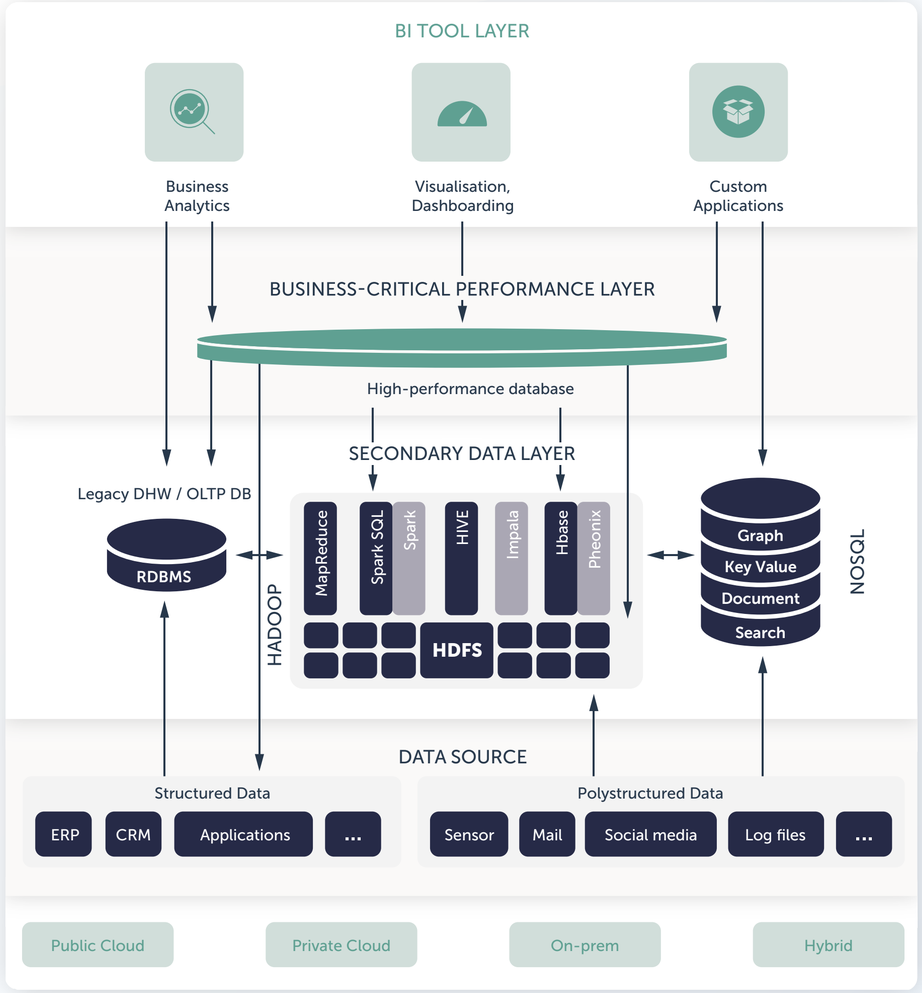

Exasol
The analytics database

Why is this important?

Because today's most valuable insights are achieved by throwing everything together and letting machine learning work its magic. If you decide against the option to add a particular set of data to your BI tool(s), you might end up missing a critical piece of insight to make that killer decision.

Equally important is the ability to plug in your chosen ETL tool. Again, the integrability of your high-performance database is crucial. If you rely on multiple data pipelines, not only will the stack be more complex and less efficient, it'll drain your DBA's time by managing them.

You might have retained some of your legacy data systems such as data warehouse solutions, or a Hadoop or NoSQL environment. That's no problem – all the data from these partitioned environments can flow up into your new high-performance layer.

This speeds up data cleaning, transformation and integration. Your BI platform runs queries faster – via a data cache or live connection set-up – as there's only one data source. By transforming your stack into a single database system, your BI platform works with one single source of the truth for all your data sources.

More than superior performance – the commercial benefits

Adding a high-performance database isn't just about better performance. A less complex environment delivers other cost benefits.

A mature data layer might include solutions from multiple vendors and that means managing a lot of contracts. And where there's administration, there's cost. By comparison, a simplified and automated database layer has lower overheads for tuning and hardware thus reducing the total cost of ownership (TCO).

Let's look at how a cleaner database layer benefits the lower stack layers.

You already know that a high-performance database's appeal is the level of built-in integration and automation with ETL tools and individual data sources. This reduces TCO even more. It also creates a flexible playground to pick and choose the specific tools you want. You move from a rigid structure where choice is restricted and improvements slow to execute to a fluid and adaptable stack.

A mature data layer might include solutions from multiple vendors and that means managing a lot of contracts.

Conclusion: Your Improved BI tools

Let's remember how introducing our new high-performance database layer to your stack improves BI:

It's faster, doing more for the same cost

It's performing deeper real-time analysis empowering your organization to make more intelligent and impactful decisions

It's sitting on a simpler, more automated stack freeing up resources to be deployed to higher-value data science projects

It's analyzing more data sources while working with one single source of truth

It's cloud-enabled – fully or partially – and reduces your TCO while enhancing security

The Result?

A more optimized (set of) BI tool(s) driving faster and more evidence-based decisions, helping you build a better business.

BI Acceleration with Exasol

Many organizations already use our analytics database with their BI tools to realize the benefits explained in this guide. Recognized as one of the best-performing databases available, we continue to push the boundaries of integration and automation, enabling BI tool customers and users to build their data stack with a wide choice of tools, strong capabilities and unrivalled power.

Experience the Exasol difference for yourself – register for a free trial at **exasol.com/test-it-now/cloud/**

Exasol
The analytics database

Exasol

The analytics database

About Exasol

The Exasol high-performance analytics database is built to run faster than any other database, delivering next-level performance, scale and ease of use. Analyze billions of rows in seconds; run high-performance analytics securely in the cloud or on-premises; deliver frictionless analytics with self-indexing that automatically tunes performance; and scale out analytics for one transparent price.

To learn more about Exasol, please visit: https://www.exasol.com/

GERMANY

Exasol AG
Neumeyerstr. 22–26
90411 Nuremberg
Germany

T: +49 911 23991-0
F: +49 911 239 91 24
E: info@exasol.com

UNITED KINGDOM

Exasol UK Limited
Parkshot House
5 Kew Road
Richmond, London
Greater London
TW9 2PR
United Kingdom

+44 20 3813 8310
info@exasol.com

USA

Exasol, INC.
c/o WeWork
1372 Peachtree Street
Atlanta, GA, 30309
USA

T: +1 415 363 5500
E: info@exasol.com