

# The database deployment playbook

Your guide to the what, where and how  
of deploying data analytics

**Exasol**  
The analytics database

**Contents:**

**Select the play you want to focus on  
and jump straight through to that section.**

<a href="#">Introduction</a>	<a href="#">03</a>
<a href="#">Play 1: On-prem and beyond</a>	<a href="#">04</a>
<a href="#">Play 2: Your cloud, their cloud, multi-cloud</a>	<a href="#">08</a>
<a href="#">Play 3: Hybrid – the agility of choice</a>	<a href="#">13</a>
<a href="#">How you can benefit from the Exasol advantage</a>	<a href="#">16</a>



# How this playbook will help

You use data every day to make better decisions, faster, more often. But you could significantly limit your ability to achieve this by getting one big decision wrong – where you deploy your analytics database.

It's a question that many in your organisation will have a viewpoint on. Many in the industry will think they have the one true answer for you.

## **But the truth is simple – you need choice.**

Perhaps you aren't sure where all your data should be, and you want to test a few things out. Perhaps you want the confidence to be able to change your mind if cloud doesn't work for you. Perhaps you have a roadmap for gradually migrating over several years and you want a partner that will walk the journey with you and deliver unmatched performance throughout. We could go on. Whatever the scenario is at your organisation, you need access to trustworthy data to make faster decisions, more accurately. You need to be able to choose to work with all of your data, wherever it lives.

So, whether you're in a highly regulated industry like financial services and need to keep sensitive data on-premises for security and compliance reasons; you're fully focused on gaining the agility provided by the cloud; or you need the best of both worlds, choice is paramount.

**This leaves the big questions – how do you know what option to choose for your organisation? And will this decision give you the flexibility to adapt to the changing needs of your business and marketplace, in one or two years' time?**

## **How can this playbook help you?**

We'll take a big, complex decision and lay it out in its clearest terms. The following pages will break down each deployment option at your disposal – on-premises, cloud or hybrid – and map out the benefits and drawbacks of each. And we'll provide you with practical next steps, identifying the most important questions you need to ask internally before making the optimal deployment decision for your analytics database.

So, if you're directly using an analytics database in your role as a data analyst or data scientist, for example, or if you're responsible for data infrastructure at your organisation, we're here for you, wherever you are on your analytics journey.

Read on to find the right deployment choice for your organisation.



# Play 1 On-premises and beyond

Working with your data on-premises is still an important option to have. In some cases it can be essential to meet compliance regulations such as HIPAA in industries like healthcare, or the wider reaching GDPR (General Data Protection Regulation). There may also be other performance reasons specific to your organisation as to why you look to this deployment choice. But keeping certain workloads on-premises doesn't mean you have to shut down all other deployment options. Many organisations' journeys to the cloud will begin on-premises – this is a starting point, not a dead end.

## Why might you be considering on-premises?

### Compliance and security

Some jurisdictions mandate that certain data is not allowed to leave its borders, and so on-premises is the only option where there isn't adequate local cloud data provision. Security is another consideration - because an on-premises network does not use the internet to transfer data, some would argue it is inherently safer than cloud.

### To make the most of legacy investments

For some organisations, running their databases on-premises ensures that legacy investments, notably hardware, see out their lifespan. And there is also the cultural consideration; the perception that systems that can be seen and touched give you more control over the databases running on them, be it tuning, patching, or integrating them with other data sources and applications.

### Save money for certain data workloads

For some large-scale and complex database workloads, on-premises can be the least expensive deployment option, both for set-up and ongoing cost predictability. 24x7 workloads are a good example of this given they don't benefit from the dynamic usage of clouds.

# Customer snapshot - delivering results on-premises



Germany's largest healthcare insurer, SBK, is proof of what can be achieved when deploying an analytics database on-premises. It's driven the delivery of typical data queries down from 6.5 hours, to 21 minutes – accelerating its time to data significantly while meeting data privacy requirements. Beyond the performance, the automation of processes and self-tuning mean that SBK employees can now spend more time with customers, rather than on maintenance.

# 2000

*data processes migrated to Exasol*

# 21 minutes

*down from 6.5 hours*

*“New and faster data analyses should improve internal processes. This saves money and lets our advisers present innovative offers to customers.”*

**Joachim Zaspel, Technical architect data  
and digitalisation mangement at SBK**

## Why would you reassess on-premises?

For all the benefits of an on-premises database deployment, there are also potential drawbacks to consider.

### Lack of flexibility

On-premises data warehouses are fixed entities and can't flex to accommodate bursts of activity that require more compute or memory, such as Black Friday for retailers.

### Infrastructure can drain budgets

Buying and operating your own infrastructure can also be expensive. You need the hardware; and the latest models don't stay superior for long. Constantly upgrading to the latest kit can be a drain on budgets.

### Security challenges

Then there's the flip side to the 'safer' argument. If an organisation is operating its IT from a single site (as is often the case with on-premises models), then it and its databases are only as strong as that location. And as well as physically securing the hardware, there's the network to secure too. Maintaining security in an ever changing threat landscape is very costly and requires expensive (and scarce) specialist security skills. Meanwhile, hackers are increasingly well-resourced and constantly finding new ways to break into corporate networks. Any weaknesses can be exploited by hackers to steal or compromise data.

Conversely, cloud hyperscalers (such as AWS or Google Cloud) are constantly evolving their security capabilities. Their scale and diversity of customers makes that a necessity. It also provides them with the knowledge and resources to innovate and enhance security at a pace that in-house IT teams are unlikely to be able to match.



## Why on-premises is the first step in bridging to the cloud

Even if you decide on-premises is the right fit for your analytics database today, that doesn't mean you have to completely shun the idea of working with your data in the cloud. The benefits of shifting from CAPEX to OPEX spend; the ability to redeploy resources from managing hardware to higher value tasks; and the scalability to increase compute power at will with pay-as-you-go (PAYG) consumption, will be attractive to many organisations. But you don't have to go all in straight away, nor do you have to compromise on performance if you start gradually migrating certain workloads to the cloud.

### Match your ambition with the right deployment option

Migrating databases from on-premises to a cloud deployment is not without risks. If your organisation is only performing basic analysis on largely static data, with little need for departmental collaboration or real-time insights, then the gains from moving to cloud may not justify the effort to get there. But as soon as analytics ambitions grow, and capabilities such as data warehousing, Business Intelligence (BI) and AI/ML become important, the case for migrating to cloud becomes more compelling.

Where the case for a cloud deployment has been successfully made, migrating databases from on-premises should not be considered a one-hit exercise. A steadier journey can be safer but then that raises the question of which workloads to move first and what infrastructure you need to maintain what's left on-premises. For instance, you may consider moving more isolated workloads to the cloud first as low hanging fruit, but you'll need to be careful this doesn't lead to transferring large amounts of data between on-premises and the cloud, which would make the process slow and expensive.

**In section three, we'll look at how a hybrid model can bridge the move from on-premises to cloud, as organisations gradually migrate the most suitable databases first.** We'll also cover how many organisations conclude that a permanent state of co-existence between on-premises and cloud database deployments is the most effective model for them.

### Working with your data on-premises – what questions do you need to ask?

- Do you have significant compliance obligations and operate in a highly regulated sector that dictates that data must stay within certain locations or jurisdictions?
- Do your organisation's security protocols require that some data should remain on-premises?
- Have you made significant investments in on-premises hardware that you need to make the most of?
- Do you need departmental collaboration or real-time insights? Or are your current data analytics needs currently more basic?
- Do you have heavy lifting data warehouse deployments, using substantial data and analytics use cases? This could see costs escalate quickly in the cloud – have you checked whether you can budget for this?
- On-premises may be right for you today, but what about tomorrow? Are there data workloads that could benefit from gradually moving to the cloud through a hybrid approach? If so, which workloads would you prioritise for migration?

### Take action

Ask these questions internally and you'll be far more likely to make the deployment decision that is right for your organisation.

And if you're facing resistance from budget holders or influencers, use these questions to make your case.

# Play 2 Your cloud, their cloud, multi-cloud

There's every chance that, if you're running analytics in your organisation, regardless of size or maturity, there will be some sort of plan to move databases to the cloud – either fully or for limited workloads. The economic and operational benefits are just too compelling to ignore; not least the shift from CAPEX to OPEX investments, and the scalability and agility that brings.

As we covered in 'Play 1', one of the big questions you'll be faced with when it comes to migrating databases to the cloud is how fast should you do it. But there are other big considerations that stem from the different deployment options you have when it comes to the cloud.

**Here's what you need to consider for each option.**

## Your cloud

### Watch out for CAPEX spend

For a private cloud run internally, your organisation will still be on the hook for equipping its data center with the right hardware, and ongoing maintenance, and so will continue to incur CAPEX spend.

Managed-service private cloud – allows your organisation to refocus its attention away from matters of technological infrastructure (to a degree). But as the ultimate owners of the data centre and the hardware, it can still be an expensive option.





## Public cloud

Option two is to run your databases on the vendor's cloud infrastructure. It gets over all the costs and stresses of a private cloud. It's also immediate.

There are some variations:

### SaaS

This provides greater scalability and flexibility. You should expect to be able to provision a database in minutes, pay for only what you use, and scale at will. Or, when a certain level of consumption is reached, move to a more cost-efficient committed spend contract.

## Self-managed

The databases are run on a dedicated virtual machine (accessed via the cloud) that is specific to an organisation. This provides more control, as the virtual machine is operated by you, but internal expertise is also required to operate it.

There is also the option of running an account on public cloud infrastructure, with similarly fast deployment and elasticity as a SaaS model, but with more control.

These variations present comparatively little immediate risk to an organisation. However, as an organisation's analytical ambitions grow, so will the capabilities they need. If an organisation's analytics database falls short of those capabilities, that could mean migrating to another more suitable provider. That takes time – and more so the more workloads you have running. This can lead to a form of vendor-lock in, where the organisation is unable to innovate, losing pace with the market.

# Customer snapshot - opening up new business opportunities with cloud flexibility



Revolut, a leading FinTech organisation with over 18 million users in more than 200 countries, is a perfect example of how deployment choice can open up new business opportunities. Running Exasol's high-performance analytics database on Google Cloud Platform, Revolut has made query times 100x faster, and crucially it isn't locked into the platform.

As a result, the organisation's data strategy is more flexible and it's been able to expand the business into countries where Google Cloud has not been directly available.

100,000

Queries run on a daily basis

100x

Faster query times

*"We are an extremely data-driven company and Exasol was a game changer for us. Queries that took hours now complete in seconds. People gain more trust in data using it on a daily basis. Today, almost every department is relying on Exasol."*

Demeter Sztanko, Head of Data Engineering at Revolut

## Multi-cloud

The multi-cloud option is about deploying your databases with a range of public hyperscalers or incorporating a SaaS model from your database provider. A multi-cloud strategy allows your organisation to select different cloud services from different providers, giving you the freedom to use the best combination of capabilities for each workload. In turn, this helps to address the risk of vendor lock-in.

### Benefits of scale

The key advantage here is – as the name suggests – scale. Hyperscalers have hundreds of public cloud data centres around the world, serving thousands of businesses. They can therefore keep their costs competitive, while their customers benefit from multi-location deployments (should national data sovereignty laws mandate the residency of certain data) and unlimited scalability.

### Check compatibility and integrations

Be aware that not every database runs on every cloud infrastructure. Moreover, the third-party tools that your organisation may prefer to use alongside its database – be it to imbue data or perform sophisticated analytics and reporting – may not be workable.

### Watch out for data siloes

A multi-cloud strategy can create other challenges, especially if the organisation has not thought long-term. When siloed workloads increase, expand or become more mission-critical, the challenge of integrating disparate databases running with different public cloud providers becomes relevant. This can present an issue. Even similar databases running on different clouds may have been configured in such a way that they are no longer compatible.

## Cloud considerations

Within the variations of cloud offerings you could opt for, there are many ways in which it can be the optimal environment for your organisation. The elasticity and scalability that helps you cater for peaks in demand across your business; the consumption-based pricing that sees you pay only for what you use; the agile development capabilities and removal of capital expenditure (CAPEX). **But any shift to the cloud has to be on your terms and meet the needs of your organisation.**

The change from CAPEX to OPEX is a case in point. It's a big positive that can help you increase agility and save money but that's not always going to be the case. If your organisation is used to an annual CAPEX cycle where you purchase equipment and amortise it over five years for a tax write-off, suddenly all of that vanishes when you move to the cloud. Is your organisation prepared for that?

You may also have to cater for heavy-lifting data warehouse deployments, with substantial data usage and analytics use cases – in this case, cloud can become very costly. Public cloud's PAYG model is better suited and more cost-efficient for variable workloads or bursts of data analysis.

Similarly, there may be parts of your organisation where people have built careers around data centre operations. They could feel threatened by or adverse to the idea of moving to the cloud, so you need to consider this if you need to earn buy-in internally.

The concept of 'data gravity' is also important to think about. Your organisation should consider its databases in the scope of the wider technology stack, such as the various sources of data, data storage solutions, and where BI tools are running. For example, if you run your BI tool in the cloud, it makes sense to run the database in the cloud as well, to minimise data transfer and latency.

### Cloud – what questions do you need to ask?

- Are you ready to migrate all of your data workloads to the cloud or do you need to take a gradual approach?
- Is your organisation ready to shift from a CAPEX to an OPEX financial model?
- Have you assessed the data workloads that you want to move to the cloud? Is the cloud going to be the most cost-effective option for you – for instance, are there resource-intensive data warehouse workloads that could become expensive?
- Are there departments within your organisation that could feel threatened by a cloud migration? Have you planned how to sell the idea internally – could you recommend retraining or redeployment of resources, for example?
- If you take a multi-cloud approach, will your databases on different clouds be compatible with each other or even with on-premises apps?
- When you select a cloud infrastructure, are you aware of integrations and general compatibility of any third party apps you use in relation to your database?
- What are your ambitions regarding data analytics? Are your use cases too basic to necessitate a shift to the cloud?
- How do data sovereignty requirements impact your management of data across different locations?

Cloud can be a great option for your organisation, but remember to interrogate each option thoroughly before making your decision. Your business and its requirements are unique – choose your own path.

# Play 3 Hybrid - the agility of choice

The journey to the cloud is not always a linear one but it does often start on-premises. This is why the hybrid approach to deployment is in play – as a consequence of cloud becoming a key component of data management architectures.

For some of the smartest organisations, a mix of different deployment models is not a temporary fix, but rather a long-term solution. But hybrid architectures are complex and nuanced.

**So, what should you consider when thinking about taking a hybrid approach to deployment?**

## Benefits of hybrid

### Flexibility

Some workloads (think highly confidential data) are more suited (or mandated) for on-premises; while databases being used for high-performance analytics benefit from the scalability of cloud. This hybrid approach promotes control, giving an organisation the flexibility to scale specific environments up or down depending on their needs at any given time.

### A journey to the cloud at your own pace

We've covered how costs, legacy investments, regulation, attitudes to risk and security, system interoperability and technological know-how shape a deployment strategy. These constraints may leave you or your colleagues frustrated that you can't move as fast as you'd like.

On the other hand, you may want to purposefully take your time, gradually migrating the less important workloads first, before turning to mission critical databases, thereby mitigating risk, as well as spreading out costs.

### Future-proof your organisation

A hybrid approach allows you to meet today's needs as well as keeping your options open should requirements change in future. Change need not mean a complete overhaul of your data management infrastructure.



# Customer snapshot - making hybrid work

# OTTO

Otto, the German home and lifestyle e-commerce giant, has significantly accelerated its data analysis and reporting since working with Exasol. Now thousands of users can access the data they need, at the same time, without compromising on performance.

And that's not all, thanks to Exasol's architecture flexibility, Otto can scale to match the changing demands placed on the business.

# 1000+

*of Otto users can access the data at the same time without sacrificing speed or performance*

# 50%

*savings from working with Exasol*

*"Why did our analytics database prove impressive?  
"Primarily due to an excellent price-performance ratio: low costs coupled with exceptional performance and integration with Hadoop."*

Andreas B., Product Owner at OTTO

# Hybrid - key considerations

Deciding whether or not to go with a hybrid approach will come down to balancing the risks and benefits of managing data across diverse and distributed environments. There are several key points to consider:

## Cost

Transferring data out of a public cloud can be expensive as Cloud Service Providers charge transfer fees, so you will need to be conscious of these rates before making your decision.

## Latency

It's essential to understand the characteristics regarding the flow of data when you plan your hybrid data management architecture. You need to know whether data is flowing into or out of the cloud for example, as well as the expected volumes of data – this will impact a number of things. Latency is one of these factors that this will impact while integration, metadata and governance practices will also require consideration.

## Standardisation

This makes migration of data from one deployment model to another less complex, because the databases have been designed to work in the same way, wherever they run.

## Integration

Integrability is important because databases operating in different environments must be able to connect, to feed advanced analytical projects. The more manual integrations you must perform between databases in different environments, the slower, more costly, and less reliable your analysis will be. This can very quickly negate many of the economic and operational benefits of a hybrid deployment strategy.

## Hybrid – what questions do you need to ask?

- Are you clear on which workloads you want to run on-premises versus a cloud environment?
- How quickly will you migrate certain workloads to the cloud?
- Do you have a comprehensive understanding of the flow of data within your organisation? Have you considered how this could impact latency?
- Have you considered how many and which integrations will be needed in order to connect databases operating in different environments?
- What impact will too many manual integrations have on performance?

## Take action

Take the lead on getting answers to these questions to make sure that a hybrid approach is the optimal option for your organisation.

Now you're considering the options, what's the next step to turn this into reality?

# How you can benefit from the Exasol advantage

## What next?

Perhaps you aren't sure where all your data should be, and you want to test a few deployment options out. Perhaps you want the confidence to be able to change your mind if cloud doesn't work for you. Perhaps you're a cloud native but need better performance for your analytics, or want to experiment with different cloud environments. Or maybe you have a roadmap for gradually migrating completely over several years and you want a partner that will walk the journey with you.

**In each scenario, Exasol is the answer.**

### Exasol deployment options at a glance

- **Exasol On-premises:** Deploy Exasol on-premises as either a single node or as an Enterprise Cluster, ensuring privacy, regulatory compliance and low latency.
- **Exasol SaaS:** Enjoy a fully managed SaaS cloud, with full service provisions including database installation, operation, administration, upgrades, 24x7 monitoring, and incident management.
- **Exasol on Public Cloud:** Deploy Exasol on any public cloud platform - Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform, with no vendor or platform lock-in. Seamlessly integrate with critical cloud services such as data ingestion, data lake, identity and access control, and downstream BI and analytics services of the big public hyperscalers.
- **Hybrid:** Get the best of all worlds by selecting the optimal deployment option for specific workloads, and use Exasol's Virtual Schemas to create read-only links between databases, allowing access to data wherever it sits without actually moving the data.

### Get started with Exasol SaaS

See for yourself what the performance of Exasol can do for your organisation. Sign up for a trial and get \$500 of complimentary usage, including:

- All the features, at full speed
- Fast deployment with fully managed infrastructure
- Elastic and infinite scale
- The separation of storage and compute
- Immediate insights, with SQL queries that run straight from your browser
- Unlimited users

[Test Exasol now →](#)



# 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: [exasol.com](https://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](mailto: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](mailto: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](mailto:info@exasol.com)