ModelOps
Buyer's Guide
Governance, Scale, & Control

ModelOps Buyer's Guide

For any business that builds risk and fraud models, one of the most significant challenges they face is model deployment and monitoring.

Model deployment processes can take more than a year, by which time model degradation may have occurred and new regulatory changes may need to be implemented.

This guide will help you understand why ModelOps is crucial to lending agility and what key capabilities to look for in a vendor.
The intention is to provide a clear overview of ModelOps with criteria for selecting a ModelOps platform.

The ability to fast-track these processes has become increasingly important as pressure from economic instability and changing consumer behaviour drives lenders to adopt new models in faster cycles, while also updating existing models more regularly.

According to our research, more than three-quarters (76%) of senior decision-makers believe that organisations that can develop and deploy models in faster cycles will increasingly have a competitive advantage.

Why?

Because it allows them to reduce time-to-market for model deployment, accelerating ROI for models (55% of lenders build models that never make it into production) and enabling them to adapt to market and regulatory changes more quickly.

As for monitoring, many businesses struggle to track drift quickly enough. Versioning is often managed on spreadsheets, making auditing prone to error, and performance monitoring is based on metrics aggregated across versions, features, and tests, with model artefacts housed in different places. Dedicated ModelOps software makes model deployment and monitoring much more effective and efficient.

55%

of lenders build models that never make it into production

What is ModelOps?

Model Operations, otherwise known as ModelOps, generally describes the systematic approach of managing the deployment and governance of a model.

The result of following this systematic approach is that organisations can standardise and scale their model requirements across different environments in a consistent way.

ModelOps makes sure that:

- Models are deployed correctly into live production within the decisioning software
- Model versioning is documented and clear
- Models are monitored to ensure they still make good predictions
- Models are updated or retrained when data or business conditions change
- The process is auditable and compliant, so regulators and managers can trust it

ModelOps is similar to DevOps in that it aims to break down legacy silos between different teams. In the case of ModelOps, those teams are the data scientists responsible for developing models and the IT software engineers responsible for deploying models into live decisioning. By speeding up or automating away the friction associated with this interface, ModelOps can significantly reduce the time required to deploy models from months to days.

When this is primarily focused on Machine Learning (ML) models, it is referred to as MLOps. However, there is another important distinction between MLOps and ModelOps. MLOps often refers to the full model lifecycle, including data preparation and model development, whereas these early stages in the model lifecycle are typically excluded from ModelOps.

Essentially, ModelOps is a collection of tools, technologies, and best practices to deploy, monitor and manage models once they have been developed. **When done right, it streamlines the process between data science and IT teams, allowing for a consistent and reliable flow of models into production.**

How long does model deployment take – and why is this important?

71%

of businesses take longer than six months to deploy a model

A new or updated model is developed to improve performance in a specific area.

Therefore, model deployment speed is a critical capability for progressive lenders, as it enables them to start benefiting from any improvement faster. However, our research with Forrester Consulting shows that **less than a third** of businesses (29%) can deploy a model into decisioning in under six months. With close to three-quarters (71%) taking more than six months.

LESS THAN A THIRD OF BUSINESSES CAN DEPLOY A MODEL IN UNDER 6 MONTHS

How long, on average, does it take to complete model deployment related activities?

29%	35%	36%
0-6 MONTHS	6 M - 1 YEAR	1 YEAR OR MORE

71% of businesses take **more than 6 months** for model development and deployment

The number 1 pain point when it comes to developing and updating models:

"The time it takes to develop, update, and push new models into production"

Base: 1320 EMEA & APAC decision-makers at Financial Services & Telco providers. Source: Experian research conducted by Forrester Consulting, August 2024

This time lag between developing the model and successfully deploying it into production has become the litmus test for lending agility.

Optimising model deployment timeframes is more important than ever, as in some cases, the underlying conditions (and therefore the dataset that was used to create the model) have changed during this delay.

The result of delayed deployment is that the model's performance may start to decline, meaning it needs to be redeveloped – potentially delaying ROI for years and increasing the risk of poor credit decisions.

Our research shows that around half (48%) of Financial Services and Telcos are already updating their models more frequently than ever before, and 54% agree that ModelOps will play a crucial role in shaping the industry over the next five years.

The ability to consistently update models through a reliable and scalable process is likely to be the main differentiator of winning organisations in the decade ahead.

MODELS HAVE A SHORT SHELF LIFE

Model degradation can happen before you know it

Financial Models degrade only 1 year after training. The longer it takes to deploy, the more variability in model outputs and the more risk to your lending practices.

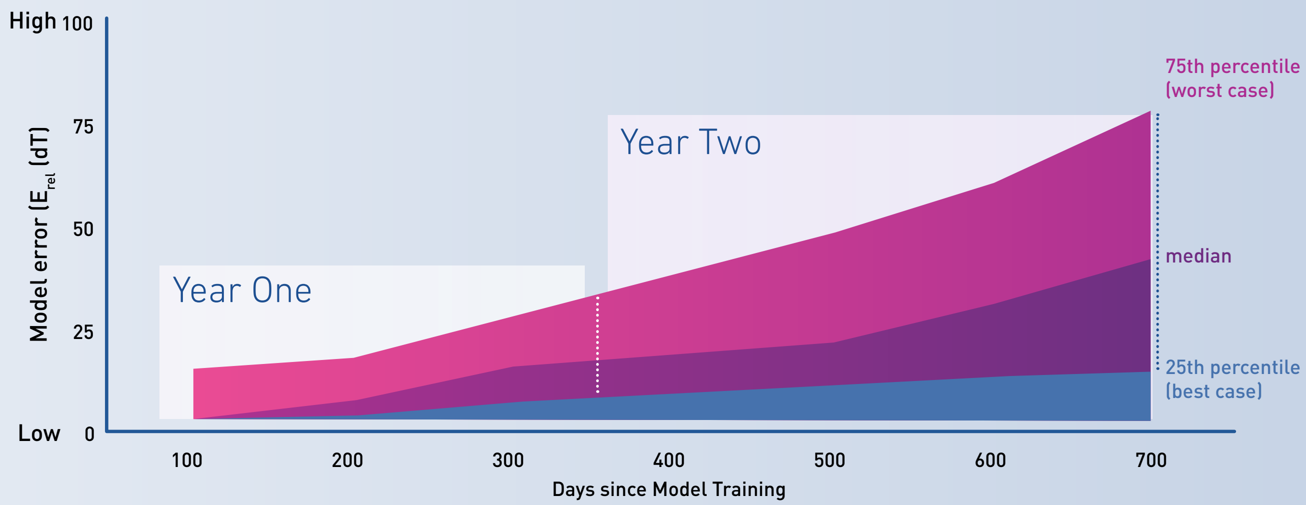

**Source: Vela, D., Sharp, A., Zhang, R., Nguyen, T., Hoang, A. and Pianykh, O.S., 2022. Temporal quality degradation in AI models. Scientific Reports, 12(1), p.11654. Note: Erel (dT) is the relative model error on the two test sets (Mean Squared Error at evaluation over Mean Squared Error at time of training)*

What are the top model deployment and monitoring challenges? And how can they be overcome?

Our research indicates that the primary challenge in updating models is the time-consuming process of updating and deploying new models into production. As we've discussed, this can have a direct impact on profitability and time to ROI for new models. But why does model deployment take so long, and how can this process be streamlined?

Common model deployment and monitoring challenges

CHALLENGE	IMPACT
Long deployment cycles 6 – 18 months	Degraded model accuracy Delayed ROI
Recoding models for production	Increased cost Increased risk of errors
Lack of monitoring and governance	Regulatory non-compliance Model drift
Siloed data and tooling	Inefficient workflows Poor collaboration between teams

In this guide, we've identified five model deployment challenges and their solutions:

- Data integration
- Containerisation
- Monitoring
- Regulatory reporting
- Recoding

CHALLENGE: **CONNECTING AND INTEGRATING DIVERSE DATA SOURCES**

ModelOps solutions effectively act as a calculation engine for scores.

Once a model is deployed, it requires data from different sources to be fed into it for real-time calculation. The challenge businesses face is integrating and connecting all these different data sources.

This process typically involves IT building out those data connections and ensuring they are compliant. And given the variety of data sources used to calculate modern scores, each comes with its own challenges and nuances for integration.

Solution

Scores need to be calculated in real time for each submitted application. The right ModelOps solution can make this process much easier. With pre-built connections into a range of diverse data sets, businesses don't need their IT team to go through the painful process of integrating each new data source.

Instead, data sources are pre-integrated, and any data refreshes are taken care of – an important point for real-time score calculation. As new data feeds emerge, the ModelOps software can scale effortlessly, enabling businesses to maintain fast, consistent, and highly accurate scoring whilst reducing operational overheads.

CHALLENGE: **MODEL AND METADATA CONNECTIVITY THROUGH CONTAINERISATION**

Models that are ready for production should be stored with their corresponding metadata in a centralised location that is accessible to multiple business units.

Each of these different teams should be able to easily access this repository with versioning and tracking available to manage audits.

Bundling each model with all of its related resources, configuration settings, libraries and dependencies is known as containerisation. This self-contained unit of software packages all the code and its associated dependencies into a single 'container'.

Containerisation allows multiple teams to access and work with data, files and software associated with each model. This improves collaboration and simplifies model auditing by tracking model runtime details, code libraries and version requirements in a single place.

Without this containerisation, the process of moving a model into different computing environments, such as testing and production, is a significant challenge.

Solution

Containerisation is becoming an important part of any effective ModelOps solution.
The containerisation process removes the need for model recoding for decisioning software because it separates the model execution from the decisioning platform's native language and runtime. In simple terms, containerisation turns models into callable production services, so decisioning platforms don't need to re-implement them – instead the decisioning software just calls them.

CHALLENGE: **MODEL RECODING**

One of the biggest causes for delaying the progress of models from the development stage into production is the need to recode them. Models are generally developed in data science environments using coding languages such as Python and R. Before these models can be deployed into live decisioning environments, they must first be recoded into common production-level coding languages, such as Java and C++. This process typically requires IT engineering or operations to complete this recoding exercise.

The recoding process is time-consuming and painful for the teams involved.

The first challenge is to ensure that there are no small errors introduced during the translation process, as these can affect the model's logic – potentially leading to different outputs in the live environment compared with the development environment. The consequences of any errors are significant, as they could impact model accuracy. Therefore, thorough testing of the recoded models is required to detect and debug them. This extensive testing is both time and resource intensive.

Recoding must precisely replicate the feature engineering and data processing rules of the original model, which can be complicated, especially when using alternative data sources that may have inconsistent formats.

Finally, recoding requires extensive collaboration between data science teams and IT engineers. A lack of clear documentation describing the model's structure or data flow can lead to misunderstandings and errors between the two teams, resulting in additional delays and frustration during the recoding process. The need for data scientists to retest and revalidate recoded models prevents them from performing other important tasks.

Solution

The good news is that advanced ModelOps solutions can eliminate the need for recoding. This enables data scientists to own the process end-to-end, removing the dependency on IT teams. New models can move from development into registration and production, with just a few clicks required to drag and drop them into the relevant decisioning workflow. This represents a step change in model deployment processes – speeding up testing and lowering risk.

TYPICAL DEPLOYMENT PROCESS

ANALYTICS TASKS

- Development of the model
- Schedule resource with IT team
- Pass model to IT team for deployment into decisioning software

IT TASKS

- Recode model for decisioning platform
- Code API e.g. for executing model via REST call
- Technical testing for correct implementation and bug fixing

WHAT OFTEN GOES WRONG

- Logic drift between original and recoded model
- Long hand-offs and release cycles
- Revalidation after every translation

ANALYTICS TASKS

- Validate the performance of the recoded model
- Amend model (if required)

STREAMLINED MODEL DEPLOYMENT PROCESS

Data scientist develops the model and then registers the model in the Model Ops software. They can then deploy to the decisioning software without IT support.

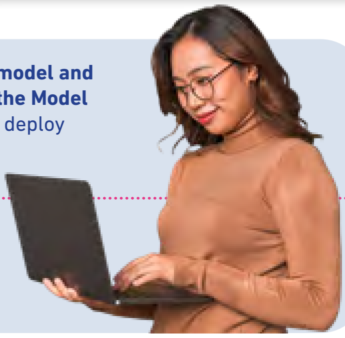

NO RECODING OR MANUAL INTEGRATION

The model's API and metadata flow automatically from the Model Ops software to the decisioning software. This ensures faster deployment.

HOW IT WORKS

The original model code (e.g. Python) runs unchanged. All dependencies are packaged with it using containerisation. The decisioning engine sends input data, receives a score/decision and applies policy and orchestration.

The decisioning software no longer needs to understand the model - only how to call it.

CHALLENGE: **INEFFICIENT MONITORING**

Where model versioning is about identity and traceability, model monitoring is about tracking behaviour and performance over time. Versioning is still fundamentally important to monitoring, as it enables monitoring to be accurate, explainable and actionable.

Model monitoring helps answer questions like:

- Is predictive power degrading?
- Is performance worsening for certain segments?
- Is bias increasing post-deployment?
- Is the model still fit for purpose with new data?

Increasing regulatory pressure means that being able to answer these questions has never been more important.

Many businesses find the monitoring process manual and time-consuming, with 59% stating that they spend an excessive amount of time preparing reports for management. These types of tasks can also place a negative drag on the model-development team.

Businesses want to automate much of this process, utilising dashboards that are automatically updated with the latest performance data. Automated monitoring enables businesses to track lifecycle lineage to support model governance processes while allowing them to store artefacts all in one place – development code, documentation, test data, results, and approvals.

Model drift is also a common challenge in model monitoring. By closely monitoring model outputs, lenders can quickly react and update models to keep them as accurate as possible. As a result, models will always require continuous monitoring and rapid retraining cycles that traditional processes may struggle with.

Solution: Automated monitoring processes

Model monitoring should involve a combination of automatic processes and dashboards to show how well models are performing. ModelOps software helps simplify processes by automatically tracking model performance metrics and detecting data drift. Integrated dashboards provide visibility into how models are performing in the real world.

The software should allow users to set up automatic alerts to notify data scientists and ML engineers when performance deviates from defined thresholds, allowing for timely intervention and retraining. When monitoring capabilities are integrated with model governance processes, they help institutions meet compliance and audit requirements by tracking the model lifecycle and performance history.

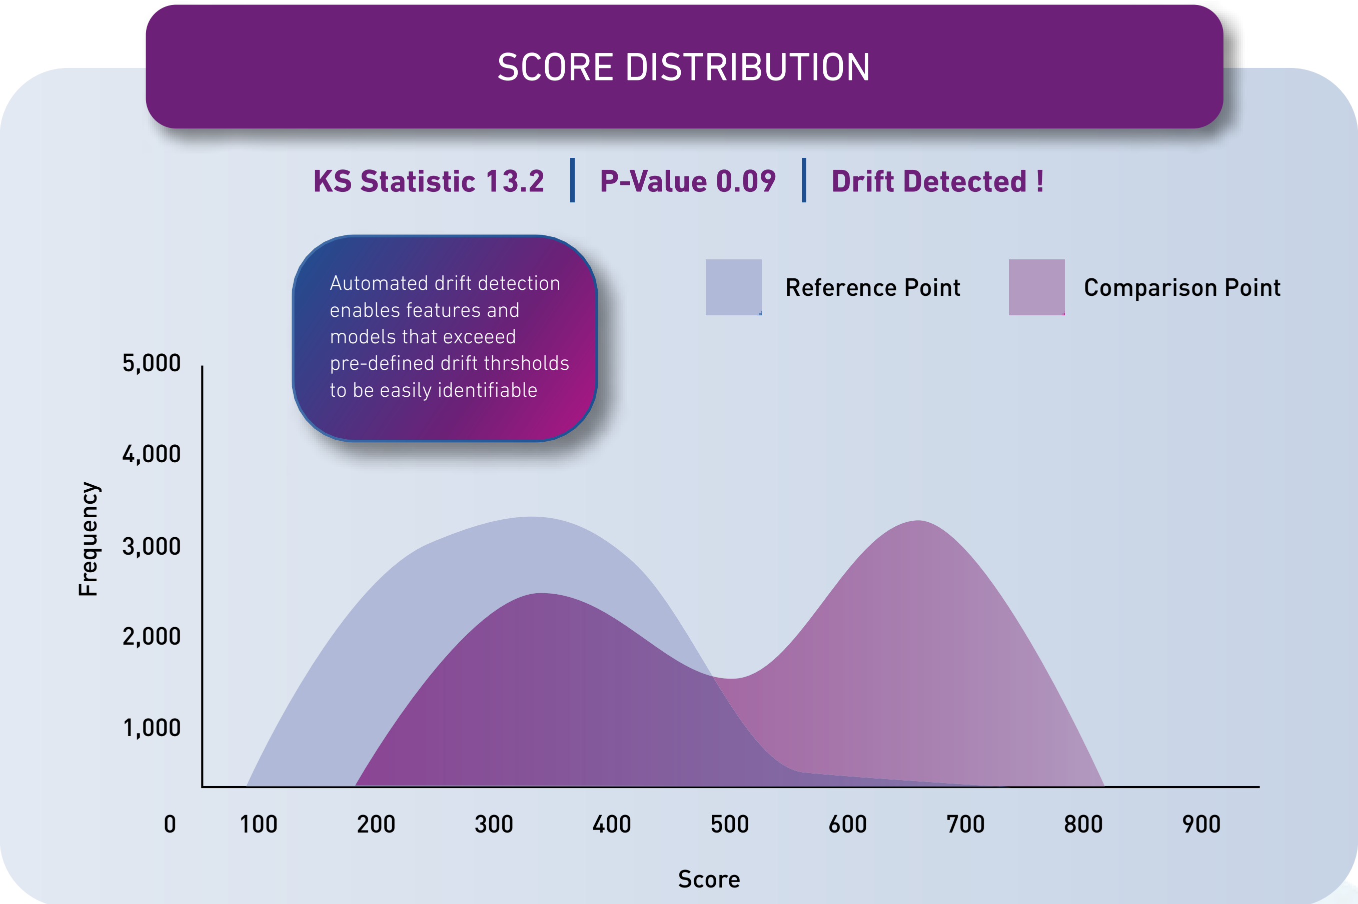

CHALLENGE: **REGULATORY REPORTING OF MODELS**

The governance, auditing, and regulatory reporting aspects of ModelOps are extremely important, but they often involve manual production of regulatory documentation.

Regulatory reporting draws on a lot of information from ModelOps-managed artefacts, such as model inventory and classification, version history, validation results, etc. With various regulatory documentation to complete, teams often spend days or weeks gathering information from different systems and applications to fulfil the necessary reporting requirements. This is a drawn-out, complex process.

More than
two-thirds
of businesses in our survey agree that the biggest advantage of a GenAI assistant is to reduce the time required to produce regulatory documentation.

Solution: Generative AI model monitoring assistants

The advent of GenAI has enabled the integration of AI assistants to support regulatory reporting.

Using templated regulatory documents – which are based on local requirements and can be fully customised for each business – information from the model development process, such as model code and diagnostics, can be auto-populated using GenAI, which fetches the information and places it in the correct fields within the template. GenAI is effectively augmenting the completion – helping to summarise information. This technology can therefore play an important role in the validation process to support compliance with local legislation and/or client policies.

Human oversight remains an integral part of the process, but the use of automation significantly improves efficiency, thereby accelerating the completion of regulatory reporting. Therefore, some solutions now include GenAI functionality that can accelerate the translation between data science and regulatory reporting within the model risk management process.

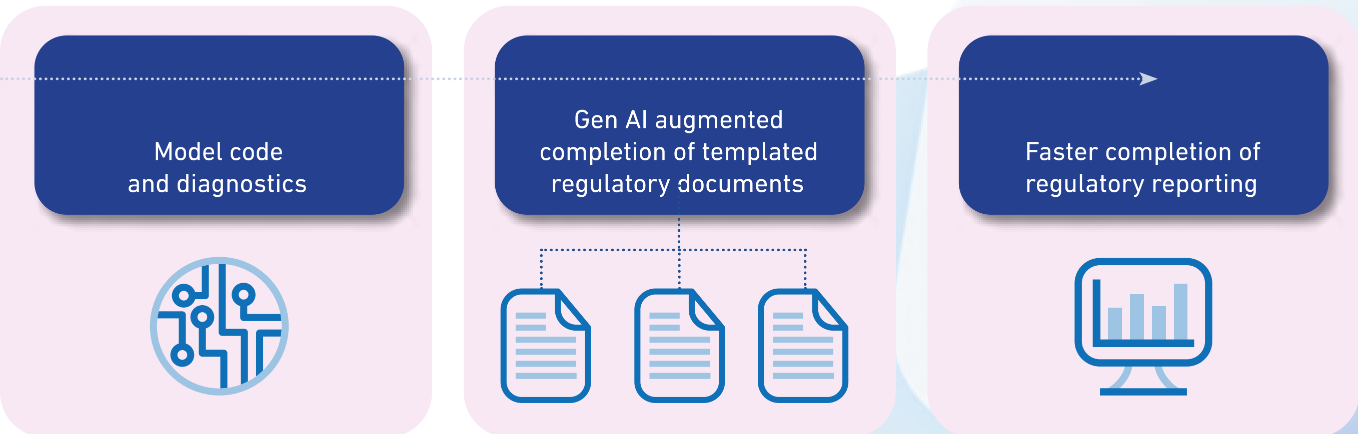

Key criteria for choosing a ModelOps solution

The following factors should all be considered when evaluating potential ModelOps vendors:

Define the tools you need: There is a range of Model Ops solutions on the market, each with different capabilities and features. Understanding your jobs-to-be-done will help you identify the right solution for your business.

User experience – one of the biggest challenges of adopting new technology is getting the buy-in of the teams that must use it on a daily basis. A simple and intuitive interface is essential to simplify adoption.

Compatibility with existing tech stack – any new ModelOps solution should integrate seamlessly with your current data and IT infrastructure.

Scalability – the solution must be capable of industrialising the deployment of large-scale AI/ML projects and managing a growing number of models (predictive, statistical, and ML) efficiently.

Vendor support and expertise – it is important that the vendor has experts available, not only to support implementation but to provide ongoing support and onboarding of teams to ensure they are trained on the solution.

Containerisation – packaging a model with all of its associated files into a single container is essential to ensure the model runs consistently across different computing environments, such as testing and production or from on-prem to cloud. This capability is crucial for scaling up model deployment.

Versioning – containerisation allows different teams to share the exact same container, ensuring that everyone is working in an identical, version-controlled environment. This simplifies versioning and makes it easy to trace and roll back to previous model versions if needed for auditing or debugging.

Continuous model monitoring – models must be automatically monitored, and when any drift is detected, an alert should be triggered with remedial steps to take.

Security and encryption – Data stored within ModelOps should follow strict security and data protection policies, including model artefacts and collaterals. Data should be encrypted in transit and at rest.

Introducing Experian's ModelOps Solution: Ascend Ops

Ascend Ops is Experian's cloud-based ModelOps platform.

It enables rapid deployment of models without re-coding, supports real-time and batch environments, and integrates seamlessly with Experian's Ascend Platform.

It provides full visibility into model performance, automated monitoring for data drift, and audit-ready governance, all of which are critical elements for meeting regulatory requirements.

Key capabilities

- **No-Code deployment:** Deploy models without recoding.
- **Real-time and batch execution:** Supports both decisioning modes for underwriting and portfolio management.
- **Integrated monitoring:** Automated drift detection, model performance tracking, and alerting.
- **Governance and auditability:** Full transparency and traceability for regulatory compliance.

By consolidating deployment processes and automating monitoring, Ascend Ops empowers data science teams to focus on innovation rather than operational bottlenecks.

Integration with Experian's modelling ecosystem

Ascend Ops is part of a broader ecosystem that includes data, analytics, decisioning and compliance solutions, which are trusted by leading lending organisations worldwide. Each component can be used individually or as part of the unified platform. This significantly reduces integration complexity, enabling a streamlined experience across different teams. Ascend Ops is fully integrated with Experian's decisioning software on the Ascend Platform to simplify the deployment of models into live production.

Visit Ascend Ops →

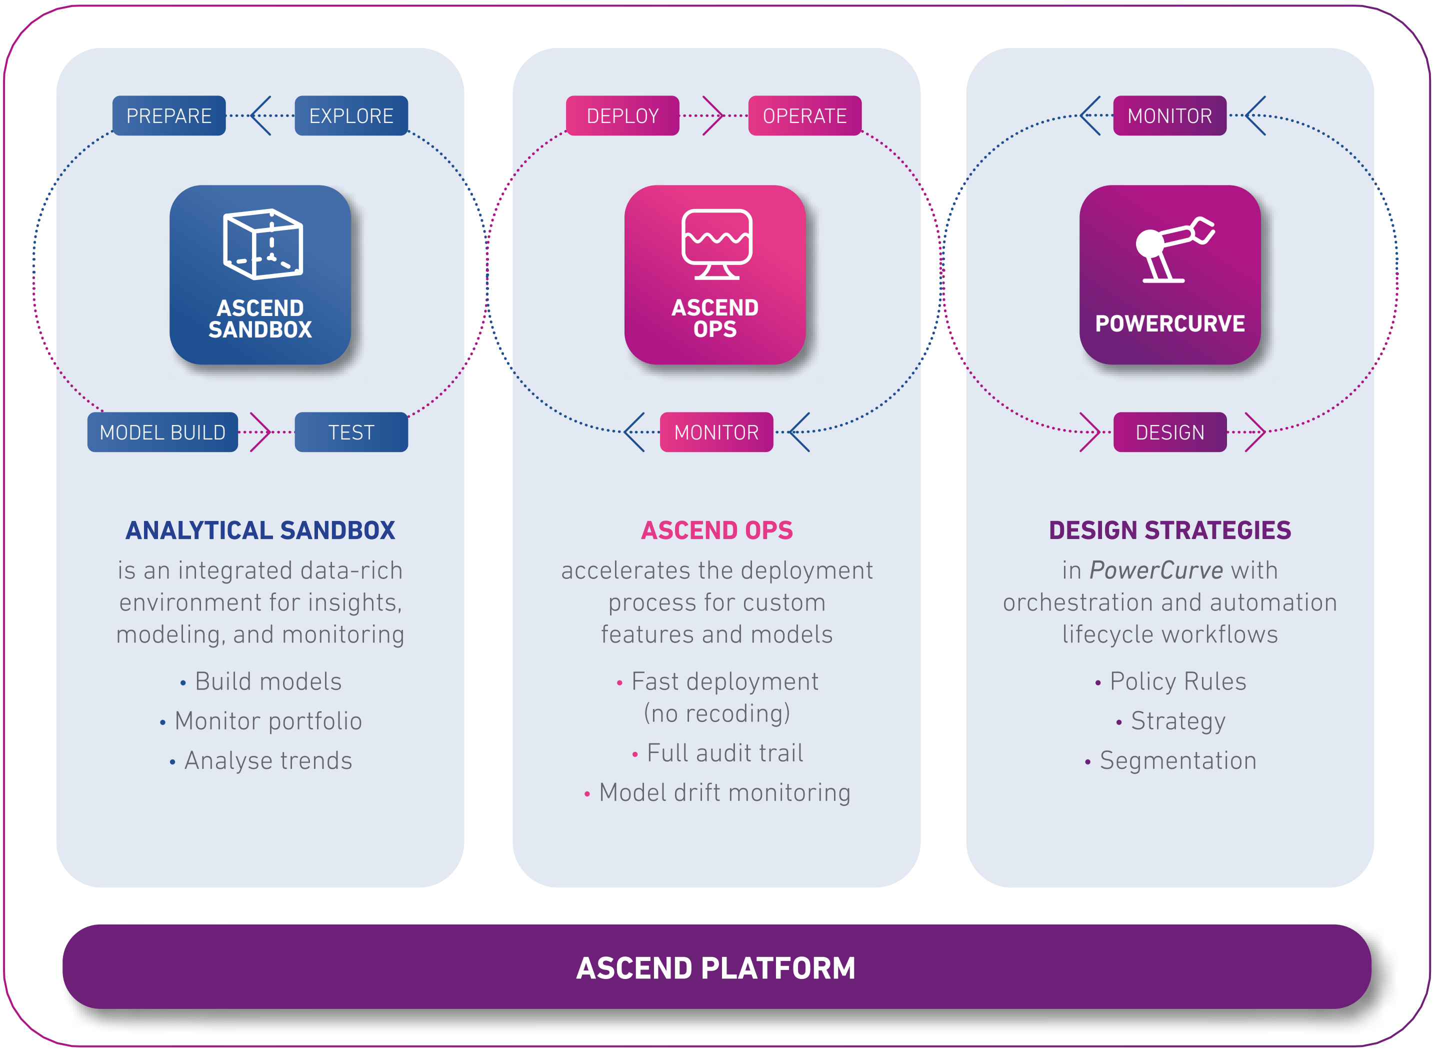

Conclusion

Like any high-impact technology, implementing a ModelOps solution doesn't happen overnight.

Every organisation will have a slightly different journey depending on its size, existing model systems and collaborative culture.

However, the potential returns are significant. Businesses that use an advanced ModelOps solution can deploy models in much shorter timeframes, which helps them to adapt their decisioning as the world changes. Ultimately, this results in greater profitability and faster ROI for modelling projects.

Experian has the data, software and expertise to help your organisation through this process.

Interested in ModelOps?

Contact your Experian Account Director to schedule a demo.