



Manage AI, with trust and confidence in business outcomes

IBM Watson OpenScale is an
open solution that helps remove
barriers to enterprise-scale AI

IBM Watson OpenScale is the open solution for businesses to operationalize trusted AI and extend their deployments enterprise-wide.

The solution enables you to

- Measure performance of production AI and its impact on business goals
- Track actionable metrics and alerts in a single console
- Enable the business user or project manager to understand AI outcomes
- Apply business results to create a feedback loop that sustains AI outcomes
- Govern and explain AI to maintain regulatory compliance
- Automatically detect and mitigate harmful bias to improve outcomes
- Accelerate the integration of AI into existing business applications

Harnessing AI to drive business value

Roadblocks on the AI journey

Businesses today are increasingly certain that AI will be a driving force in the evolution of their industries over the next few years. Many are already taking their first steps on this journey, whether by building AI-powered chatbots to augment their call centers, or by automating back-office tasks by using AI to process documents or recognize images.

Yet for every successful AI project, there are many that fail to reach widespread adoption in the business and achieve their expected outcomes. Even the most expert data science teams will typically only deploy a handful of models into production every year. This is partly because the mechanics of AI deployment can be complex, and there are still gaps in skills and tooling that can make it difficult for data science, IT operations, and business teams to work in lockstep. But beyond the operational challenges, there are also much more profound issues of trust and transparency that businesses need to address before they can truly turn AI into a business advantage.

Building trust

Knowledge workers must be able to trust AI and explain the decisions it helps make before they will incorporate it in their business processes. If AI is a black box that simply takes in data and produces obscure, unexplainable outcomes, then there is no way for the business to judge whether these systems are producing fair, accurate outcomes, or have confidence in AI's ability to augment decision-making. Equally, the business will not be able to explain outcomes to customers, auditors, or compliance teams.

Today, many promising models never make it into production because businesses cannot afford to trust AI decisions they do not fully understand. A business exposes itself to significant risk if it delegates responsibilities to an AI system that does not fully align with its aims and policies. For example, severe financial or reputational damage could result if a model unfairly discriminates against a particular group of customers because its training data did not represent a large enough sample of that population, or if model performance drifts and no longer meets the intended business goals.

Making decisions explainable

In many industries, adherence to GDPR and other comprehensive regulations presents a significant barrier to AI

adoption. Even if a company is satisfied its models are fair and trusts the results, regulators may demand a more rigorous lineage of AI platforms and components.

For this reason, it is critical that AI outcomes are fully explainable by keeping a complete track of inputs and outputs of any AI-powered application. It should be possible to audit the lifecycle of every AI asset from initial design, to training and deployment, through to operation and retirement. For a given model, it should be possible to identify the team who built it and the data sets they used to train it, as well as the inputs it received in production and the outputs it produced.

Ensuring fairness

Almost any AI model, no matter how carefully designed, is likely to exhibit a certain amount of bias. A model is only as good as the data with which it is trained, and since training data sets can never be 100 percent representative of real-world data, there is always risk that a newly trained model may not perform well in production. Moreover, since most data domains are continuously evolving, model accuracy tends to drift over time.

The key is real-time visibility at runtime, where decisions are made. If you can monitor the accuracy, performance, and fairness of your AI models throughout their operational lifecycle and provide analytics to help line-of-business users understand the reasoning behind the results, then you can overcome one of the most significant roadblocks on the AI journey.

Introducing IBM Watson OpenScale

IBM Watson OpenScale makes it easier for data scientists, application developers, IT and AI operations teams, and business-process owners to collaborate in building, running, and managing production AI. This empowers businesses to confidently integrate machine learning capabilities into their applications and scale seamlessly as the demand for AI grows.

How it works

IBM Watson OpenScale includes a powerful operations console that makes it easier for business users to track and measure AI outcomes. This allows business users to correlate outcomes to their organization's KPIs and improve models to account for changing business situations. These analytics capabilities can also be easily integrated with many common

business reporting tools to provide insights to a wider audience. The solution augments the AI environment with instrumentation, payload logging, and monitoring services that provide deep insights, end-to-end auditability, and fine-grained control.

For example, with Watson OpenScale, a business user can define and monitor for custom metrics, or apply out-of-the-box metrics, in order to track model performance. Then, with drift detection capabilities in payload, a user will better understand how to manually improve the model.

Watson OpenScale also runs a sophisticated set of diagnostic services to assess the accuracy of the model. State-of-the-art anomaly and bias detection capabilities, developed by IBM Research, identify and automatically mitigate harmful biases in both the data and the model, without requiring retraining from the data science team. Bias checks can be performed both at build time and during runtime, to help ensure that any issues are caught as early as possible.

Watson OpenScale also features model explainability for visibility into model decisions. A business user can understand which factors contributed to an AI outcome to meet regulatory demands and customer expectations around transparency. Watson OpenScale also features a cutting edge capability from IBM Research - contrastive explanations - that show the minimum changes to inputs that would need to occur for an outcome to change. Contrastive explanations provide even deeper understanding into the factors behind a given output, which will be increasingly important as AI is embedded in more workflows and businesses demand greater visibility into the inner workings of models.

An open solution for AI

While IBM Watson OpenScale integrates seamlessly with IBM tools for building and running AI models, such as IBM Watson® Studio and IBM Watson Machine Learning, it has been designed as an open solution that will easily operate with model development environments from other vendors and open source tools, including TensorFlow, Keras, SparkML, Seldon, AWS SageMaker, AzureML and more. Regardless of your existing investments in model design, training, and evaluation tools, IBM Watson OpenScale offers value by closing the gaps between your data science team, IT team, and business process owners. Above all, it provides a unique set of monitoring and management tools that help you build trust and implement control and governance structures around your AI investments.

The solution enables organizations to keep their options open in terms of AI deployment, with a wide range of connectors that allow users to build and deploy their models anywhere – in the IBM Cloud™, in IBM Cloud Pak™ for Data, or on a variety of other cloud platforms.

Monitor credit risk for bias, explainability

IBM Watson OpenScale makes it possible for organizations to take their AI assets out of development and into the real world – helping to solve business problems and deliver value, while significantly mitigating risk.

For example, as lenders face increasing pressure from the market to give more loans to a wider variety of customers, they still need to limit their risk, so they are relying on new types of data to help them make lending decisions. In order to analyze more varied, complex data, they are turning to more complex and advanced machine learning techniques, which can be black boxes.

For both regulatory and business reasons, lenders need to

be able to understand why their credit risk models make specific decisions for individual applications, and they need to ensure that their models are not operating with bias against certain groups.

Watson OpenScale can help lenders avoid black box AI and opaque predictions, provide insight into the performance of a credit risk model in production, limit risk exposure from regulations, and ultimately create more fair and explainable outcomes for customers.

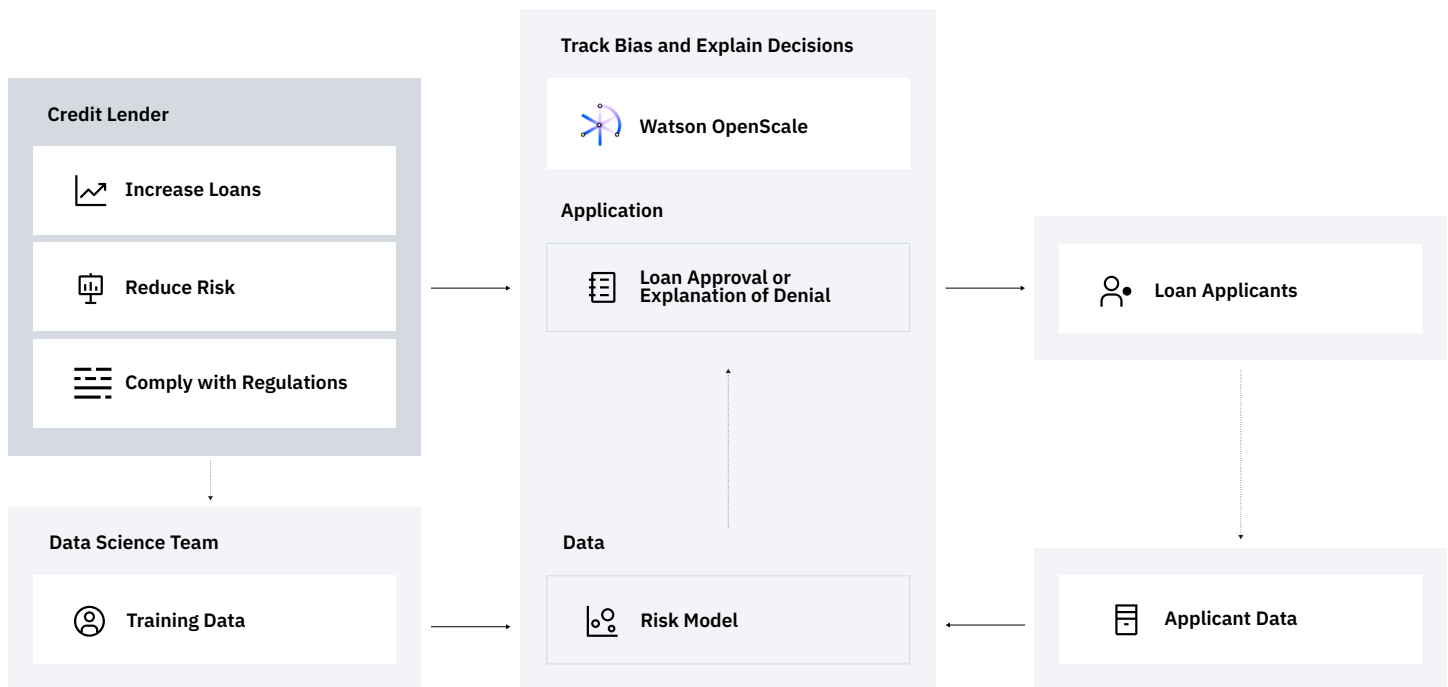
AI you can rely on

IBM Watson OpenScale provides the final piece of the puzzle to help organizations get AI projects out of development and into production. By providing full explainability and monitoring, the solution helps businesses ensure fair outcomes, remain compliant with regulations, and increase confidence in the value of AI.

For more information

To learn more about IBM Watson OpenScale Contact your IBM representative or IBM Business Partner Visit www.ibm.com/cloud/watson-openscale

Use Case Deep Dive Credit Risk Modeling



Monitor risk models for performance, bias, and explainability to limit regulatory exposure and produce better outcomes.