Build and scale trusted AI on any cloud

Use IBM Watson Studio for IBM Cloud Pak for Data to automate the AI lifecycle for ModelOps, and to implement explainable AI through model monitoring.

Highlights

— Build, run and manage AI models; optimize decisions at scale across any cloud.

— Operationalize AI anywhere as part of IBM Cloud Pak for Data, the IBM data and AI platform.

— Unite teams, simplify AI lifecycle management and accelerate time to value with an open, flexible multicloud architecture.

Realize results faster

— Automate AI lifecycles with ModelOps pipelines.

— Speed data science development with AutoAI.

— Prepare and build models visually and programmatically.

— Deploy and run models through one-click integration.

— Promote fair, explainable AI with model monitoring.

— Drive better business outcomes by optimizing decisions.

Businesses are using AI to predict outcomes, streamline operations, improve efficiency, and protect against cyber threats and fraud. These capabilities can help leaders stay ahead of competitors and market fluctuations while managing costs and risks. However, organizations also face pressure to demonstrate the ethical, responsible use of AI and meet business expectations even while they change their processes, technologies and organizations to implement AI.

The key to meeting all these demands and operationalizing AI across the enterprise is a unified data and AI platform based on a containerized, modern architecture. With this, you can implement data science and AI capabilities together with data management, data governance, data warehousing and other AI tools in an intuitive, integrated environment. Data scientists, developers, business analysts and line-of-business experts gain access to the structured and unstructured data they need to train models and run analyses. IBM Watson® Studio empowers you to operationalize AI anywhere as part of IBM Cloud Pak® for Data, the IBM® data and AI platform. You can unite teams, simplify AI lifecycle management and accelerate time to value with an open, flexible multicloud architecture.

With IBM Watson Studio, data scientists, developers and analysts get the tools to do their jobs—drawn from open source, third-party and IBM technologies—integrated into a single platform with a consistent experience. Organizations can automate AI lifecycles by using IBM AutoAI to help beginning and expert data scientists conduct experiments faster. Drag-and-drop visual data science tools also help people without coding backgrounds. Watson Studio helps you inject decision intelligence into your applications with the combined power of predictive and prescriptive analytics.

IBM

This seamless collaboration in a unified environment leads to substantial productivity gains that save both time and money in building, deploying and managing AI models. By using model monitoring and explainable AI capabilities, you can track the impact your models are having on your business, while detecting model bias, drift and risk as part of the end-to-end AI lifecycle. This is why IBM Watson Studio on IBM Cloud Pak for Data is the ideal environment to build your model operations, or ModelOps.

The ideal environment for operationalizing AI

IBM Watson Studio is ideally suited to help you build ModelOps—a principled approach to operationalizing a model in applications. ModelOps synchronizes cadences between the application and model pipelines. With ModelOps you can accelerate your AI deployment time-to-value using data, models and resources from edge to core to cloud. In the operationalization scenario, you need the data and AI platform that simplifies the process of deploying, optimizing and updating these deployed models. IBM Watson technology empowers you to save time and costs for managing and monitoring the releases, tracking model performance and its impact on key performance indicators while ensuring model governance, security and resiliency.

Work with a wide variety of open source data science tools

With Watson Studio you can use open source frameworks such as PyTorch, TensorFlow and scikit-learn. This helps you bring together development tools including popular IDEs, Jupyter notebooks, JupyterLab and CLIs—or languages such as Python, R and Scala. Data scientists can use Notebooks or RStudio® for analysis within a project. Watson Studio environments make it easier to work with open source data science tools without suffering vendor lock-in.

Federated learning for secure, improved training

Organizations want to improve their models without causing security or privacy issues. Federated learning:

– Enables sites with large volumes of data to be collected, cleaned, and trained on an enterprise scale without migration
– Accommodates for the differences in data format, quality, and constraints
– Complies with data privacy and security needs, even while training models with different data sources

Available through IBM Watson Studio, federated learning also benefits organizations where different entities from different countries and cloud providers may want to use their data to train a model without sharing the data.

Faster training and inference for deep learning workloads

Deep learning can interpret text, images, audio and video at scale, generating patterns for recommendation engines, sentiment analysis, financial risk modeling and anomaly detection. IBM Watson Studio on IBM Cloud Pak for Data helps a business:

– Scale compute, people and apps dynamically across any cloud.
– Manage and unify large data sets and deep learning models with transparency and visibility.
– Adapt models continuously with real-time data from edge to hybrid clouds.
– Optimize cloud and AI investments with faster training and inference.

With Watson Studio you can build your deep learning models from initial prototype to enterprise-wide quicker. In this environment you can push deep learning models for apps in a containerized, hybrid cloud foundation while uniting data and model deployment anywhere. Watson Studio also helps you share and optimize GPU and CPU allocations that are tuned to workload demands.

Explainable AI and model monitoring

With model monitoring and management on a data and AI platform, you can:

- – Monitor model fairness, drift and risk.
- – Visualize and track AI models in production.
- – Validate and test models to mitigate regulatory, reputational and operational risks.
- – Increase end-to-end visibility of AI lifecycles.

Furthermore, you can prepare and configure model monitors with model input, including training data. This helps you track and visualize model insight in a single environment. With endpoint monitors, you can log scoring requests or evaluate models for fairness and drift, using a payload logging endpoint. Then you can generate code snippets for the payload and feedback endpoints and for debiased transactions so that you can integrate them into your applications.

Automate data prep, feature engineering and hyperparameter optimization

The AutoAI graphical tool in Watson Studio automatically analyzes your data and generates candidate model pipelines customized for your predictive modeling problem. These model pipelines are created iteratively as AutoAI analyzes your dataset and discovers the data transformations, algorithms, and parameter settings that work best for your problem setting. Results are displayed on a leaderboard, showing the automatically generated model pipelines ranked according to your problem optimization objective. You can also explore top algorithms, pipelines, and feature transformation on a relationship map, as shown in *Figure 1*.

Predict and optimize outcomes

With Watson Studio, you can take predictive results and apply prescriptive capabilities to optimize decisions. You can import or create decision optimization models in Python, Optimization Programming Language (OPL), or natural language using the Modeling Assistant. With CPLEX® optimization engines, your organization can evaluate millions of possibilities to find the most appropriate prescriptive solutions and make optimal business decisions. This helps improve operational efficiency by combining data science capabilities, machine-learning techniques, model management, and deployment.

Give analysts a drag-and-drop model development environment without coding

IBM SPSS® Modeler, available through Watson Studio, offers business analysts and data scientists an easy-to-use, interactive way to develop predictive models without the need for programming—as shown in *Figure 2*. It provides automated modeling with out-of-the-box, industry-leading algorithms as well as a range of advanced analytics, including text analytics, geospatial analysis and optimization.

Built-in data preparation and profiling with Data Refinery

With Watson Studio you can increase productivity with easy-to-use data prep so that data is cleansed, refined and enriched. Watson Studio provides a powerful set of self-service capabilities that can be used to explore datasets, iteratively cleanse and refine them, and visualize the results. The solution also automatically tracks and documents every step in the refinement process to provide end-to-end data flow and makes it easy to save and share the output with other users.

Data Refinery integrates with a wide range of cloud and on-premises data stores through a secure gateway, enabling users to load data from source systems in minutes or even seconds. The availability of so many out-of-the-box connectors means that data scientists and other knowledge workers are no longer reliant on support from the IT team. If their access credentials are valid, and a suitable connector exists, they can set up a new connection for themselves.
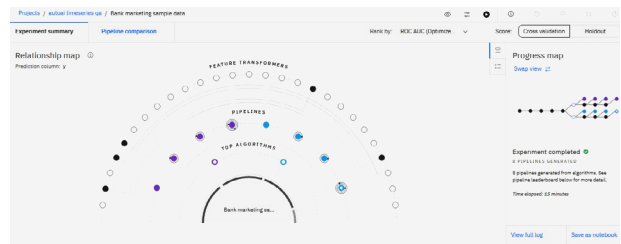

Figure 1. Automate data prep, feature engineering and parameter optimization
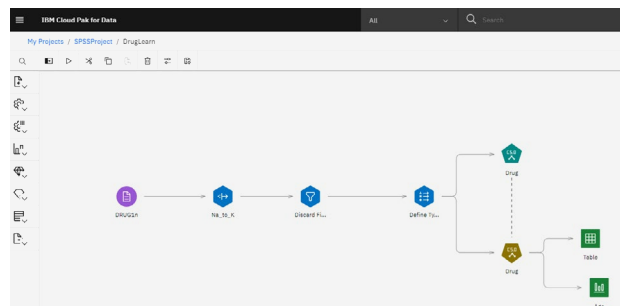

Figure 2. Build predictive models using a drag-and-drop interface

Visualize insights with integrated dashboards

Integrated analytic dashboards can enable analysts to discover insights from data and turn the analytical results into user-friendly visualizations for sharing with a broad audience of business users. Dashboards can be created from within a project using drag-and-drop tools that access available data assets. The system provides automatic recommendations for effective visualizations based on the type of data selected, while built-in templates and styles make it fast and easy to format reports as needed.

Deliver self-service access to data and other knowledge assets

Through integration with Watson Studio on IBM Cloud Pak for Data, you can automate and simplify data discovery, facilitate curation and provide active governance. Intelligent AI-powered search capabilities help users find the structured and unstructured data, notebooks, and other knowledge assets they need, while metadata such as tags, comments, and quality metrics help them decide whether a data set will be useful to them and how best to extract value from it. Lineage of assets, including models, is automatically captured to give users the ability to understand where an asset came from, where it was used and what the inputs were. Integrated active governance capabilities give users confidence that they are permitted to use a given data set while automatically masking sensitive data so they are not able to see it. This helps ensure that the assets in the catalog are used responsibly by others in the organization.

Accelerate analytics development with IBM Analytics Engine

With IBM Analytics Engine, you can run Jupyter notebooks and jobs from tools in Watson Studio by selecting IBM Analytics Engine as your runtime environment. It enables data scientists to rapidly provision, manage, run and retire Apache Hadoop and Apache Spark clusters. It increases flexibility by keeping compute and storage infrastructure separate, so each can scale independently to prevent loss of data if a compute cluster fails. It also simplifies the analytics infrastructure and streamlines workflow.

Enrich apps with integrated AI services

Available as part of IBM Cloud Pak for Data, Watson Studio is designed for seamless integration with IBM AI services, enabling users to create a service instance and associate projects with services to enable collaboration and extended use. Services such as speech-to-text, text-to-speech, tone analyzer and natural language understanding are available in IBM Cloud Pak for Data and provide APIs that you can run in notebooks in Watson Studio.

Choose the right deployment for your organization

With Watson Studio, you can use the cloud of your choice—whether IBM Cloud®, Amazon Web Services, Microsoft Azure, Google Cloud, your own private cloud—or run on-premises. An open, hybrid cloud strategy empowers you to build and manage AI workloads from anywhere, without vendor lock-in.

For more information

To learn more about how IBM Watson Studio can help you develop AI-powered solutions, check out the 2021 Gartner Critical Capabilities for Data Science and Machine Learning Platforms report, or visit: ibm.com/cloud/watson-studio.

To get started right away, review the Watson Studio on IBM Cloud Pak for Data documentation and start your no-cost Watson Studio trial.