

IBM Technology



G-Cloud 15 IBM - Watsonx.ai Service Definition Document

Email: ukcat@uk.ibm.com

Contact: Anne-Marie Wheeler

Service Description

IBM watsonx.ai as a Service Multicloud

This Service Description describes the Cloud Service. The applicable order documents provide pricing and additional details about Client's order. Upon acceptance of Client's order, this Service Description and the Cloud Services Agreement found at <http://www.ibm.com/terms/csa> (or equivalent agreement for Cloud Services between the parties) applies to Client's use of the Cloud Services.

1. Cloud Service

The functionality of the watsonx.ai platform is comprised of watsonx.ai Runtime (formerly known as Watson Machine Learning) as well as watsonx.ai Studio (formerly known as Watson Studio Paygo). If Client wishes to deploy the watsonx.ai platform, entitlements must be acquired to both watsonx.ai Runtime and watsonx.ai Studio.

1.1 Offerings

IBM watsonx.ai as a Service Multicloud is comprised of two distinct cloud services. Combined they offer a wide range of optional features and capabilities for developing AI applications and solutions.

The Client may select from the following available offerings.

1.1.1 IBM watsonx.ai Runtime as a Service Multicloud

Enables training and deployment of machine learning models and integrating them into applications. The Cloud Service provides predictive analytics at scale for applications needing model training, batch, or online scoring. The Cloud Service allows developers to access machine learning technology using simple APIs without having to manage any infrastructure.

1.1.2 IBM watsonx.ai Studio as a Service Multicloud

Enables an integrated development environment for a suite of tools and capabilities that are designed to assist Clients with a wide range of data science and machine learning tasks to enhance their productivity. The following optional IBM offerings are embedded within this Cloud Service:

- SPSS Modeler which provides the ability to create visual data flow executions for data science and machine learning (ML) tasks.
- Decision Optimization which provides an environment to perform prescriptive analytics by leveraging CPLEX and CP Optimizer engines.

2. Data Processing and Protection Data Sheets

IBM's Data Processing Addendum at <http://ibm.com/dpa> (DPA) and the Data Processing and Protection Data Sheet(s) (referred to as data sheet(s) or DPA Exhibit(s)) in the links below applies to IBM's processing of personal data on behalf of Client.

<https://www.ibm.com/terms/?id=DCDEDC1A3B3840CA8D2E82FF64849F52>

To view the compliance certification available by region for watsonx.ai platform services, please consult the documentation: <https://datapatform.cloud.ibm.com/docs/content/wsj/getting-started/regional-datactr.html?context=wx&locale=en&audience=wdp>

3. Service Levels and Technical Support

3.1 Service Level Agreement

IBM provides the Client with the following availability service level agreement (SLA). IBM will apply the highest applicable compensation based on the cumulative availability of the Cloud Service as shown in the table below. The availability percentage is calculated as the total number of minutes in a contracted month, minus the total number of minutes of Service Down in the contracted month, divided by the total number of minutes in the contracted month. The Service Down definition, the claim process and how to contact IBM regarding service availability issues are in IBM's Cloud Service support handbook at <https://www.ibm.com/support/pages/node/731179>.

Availability	Credit (% of monthly subscription fee*)
--------------	--

Less than 99.9%	2%
Less than 99.0%	5%
Less than 95.0%	10%

* The subscription fee is the contracted price for the month which is subject to the claim.

3.2 Technical Support

Technical support for the Cloud Service, including support contact details, severity levels, support hours of availability, response times, and other support information and processes, is found by selecting the Cloud Service in the IBM support guide available at <https://www.ibm.com/support/home/pages/support-guide/>.

3.2.1 Optional Support Offerings

Additional support options are available for this Cloud Service. If Client selects to order an optional support offering, the terms defined in the Optional Support Offerings Service Description at <https://www.ibm.com/terms/?id=i126-9789> will apply to the support offering specified in Client's TD.

4. Charges

4.1 Charge Metrics

The charge metric(s) for the Cloud Service are specified in the Transaction Document.

The following charge metrics apply to this Cloud Service:

- Resource Unit (RU) is an independent measure of a resource managed by, processed by, or related to the use of the Cloud Service.

The resources counted in the Cloud Service are described in the tables below.

- Capacity Unit Hour (CUH)** is an hour of use of the selected feature within a specified compute environment that will be multiplied by the specified number of required capacity units.
 - For the purposes of tooling within the IBM watsonx.ai Runtime service, 2 CUH are equivalent to 1 RU
 - For the purposes of tooling within the IBM watsonx.ai Studio service, 1 CUH is equivalent to 1 RU
- Hour** is an hour of use of the Cloud Service for model hosting. For the purposes of the IBM watsonx.ai Runtime service, 1 Hour is equivalent to 1 RU.
- Content Tokens, Data Points or Characters** are the resources counted for the IBM watsonx.ai Runtime service. The amount of Content Tokens, Data Points or Characters processed by a model is first divided by 1000 and factored by a model specific multiplier to determine the total quantity. The total quantity, when divided by 10,000, is equivalent to the corresponding number of RUs.
 - A Content Token is a unit of input and output content such as individual words, and subwords of a sentence.
 - For foundation models, the total number of input Content Tokens and output Content Tokens are counted.
 - For embedding models and for synthetic data generation, only output Content Tokens are counted.
 - Character is any encoded symbol (such as a letter, number, icon, or space) that represents information. The number of input and/or output Characters are counted.
 - Data Point is a unit of input and output content expressed as one or more numbers. The total number of input Data Points and output Data Points are counted.
- Page** is a collection of content, contained on one side of a sheet of paper when printed, whether produced in physical or digital format, managed or processed by the Cloud Service. For the purposes of the IBM watsonx.ai Runtime service, 33 pages is equivalent to 1 RU. Character is any encoded symbol (such as a letter, number, icon, or space) that represents information. For the purpose of the text extraction feature, the following file formats will be billed as follows:

- PDF – 1 page in the pdf is 1 Page
- HTML - up to 3,000 Characters = 1 Page
- Plain Text - up to 3,000 Characters = 1 Page
- Markdown - up to 3,000 Characters = 1 Page
- Word Doc - up to 3,000 Characters = 1 Page
- PPT - 1 slide = 1 Page
- Image - 1 Page
- In the case of file formats that do not have delineated pages, a Page will be considered as consisting of 3,000 Characters.

5. Additional Terms

For Cloud Service Agreements (or equivalent base cloud agreements) executed prior to January 1, 2019, the terms available at <https://www.ibm.com/acs> apply.

5.1 AI Technology Terms

a. Foundation and Embedding Models

The Cloud Service may include generative AI foundation models and/or embedding models as described in the documentation for the Cloud Service. Each foundation model and embedding model included with the Cloud Service will contain a model information document that provides information about that model ("Model Information"), including details on the model provider and applicable license terms. IBM-developed foundation and embedding models made available with the Cloud Service are considered part of the Cloud Service and cannot be used separately from the Cloud Service. Some models included with this Cloud Service are a Non-IBM Product and are separately licensed to the Client under the terms of the applicable third party license agreements(s) and not this agreement. By using a model that is a Non-IBM Product, Client accepts the applicable license terms, including any associated obligations and use restrictions specified in the third party license agreement and Model Information.

Client is prohibited from extracting the Model Weights of any model. "Model Weights" means the numerical values associated with the parameters of the training, optimization and performance of a model.

All Client inputs to a model, including by inferencing or prompting the model, and any content provided by the Client for further training of a model, including fine-tuning or prompt-tuning, are Client Content. Except as set forth in (c) (Output Indemnity) below, Client is solely responsible for Outputs generated by any model. IBM does not claim to have ownership rights to the Output. Unless separately authorized by Client through a voluntary feedback submission, IBM will not use the Client Content or model Outputs to train the foundation models or embedding models.

b. AI Use Restrictions

Client agrees not to, and not to direct or allow third parties, use foundation models in connection with the use of this Cloud Service : (i) for mass surveillance, racial profiling, or any use that violates or encourages the violation of basic human rights or other applicable laws and regulations; (ii) to distribute false, misleading, disparaging or obscene information or content; (iii) to provide fully automated decision making in connection use cases involving critical processes or the risk of loss of life, property or impact on an individual's legal rights; (iv) in a manner that impersonates another for deceptive purposes or conceals the fact a user is interacting with AI; or (v) to distribute or intentionally generate malware or other harmful code.

This Cloud Service is not intended for use in a high risk context as defined under the EU AI Act or other applicable regulations. If you chose to use it in such, please contact ChiefPrivacyOffice@ca.ibm.com.

c. Output IP Indemnity

IBM's obligation to protect Client against third party claims as set forth in the applicable base agreement includes third party claims that the result generated by the Covered Model in response to a Client's prompt or other input (collectively, "Output") infringes that third party's patent or copyright provided that (in addition to the exclusions and Client obligations provided in the relevant base agreement):

- (1) Client has a paid subscription to the Cloud Service or the Covered Model at the time Client generates the Output.
- (2) Client enables all available filters and logging features recommended by IBM in the product documentation or delivered with the Covered Model or in the Cloud Service, and Client supplies information collected by such filters and features as requested by IBM in connection with a third party claim;
- (3) Client was using a supported version of the Covered Model;
- (4) the Output at issue was generated in connection with an Indemnified Use Case;
- (5) Client takes mitigating actions reasonably requested by IBM, which may include ceasing use of the allegedly infringing Output; and
- (6) Client had no reasonable basis to know that the Content or Output was infringing or likely to cause an infringement.

“Covered Model” means an IBM developed foundation model made available with the Cloud Service (“IBM Covered Model”) and a third-party foundation model provided with the Cloud Service or as a separately purchased entitlement to the Cloud Service and identified in the product documentation or in its separate transaction document as being an indemnified third party model (“Third Party Covered Model”).

“Indemnified Use Cases” means use within Client’s enterprise (a) of an IBM Covered Model in a manner consistent with such model’s intended use as described in its published Model Information or other Transaction Documents (such as license information or program specifications), and (b) of a Third Party Covered Model to analyze, summarize or classify Content or extract information or insights from Content provided that Client has accepted and complied with the applicable third party license.

All other terms of ‘(a) Foundation and Embedding Models’ in this section 5.1 and the Liability and Intellectual Property Protection terms in the base agreement remain in effect.

d. **Liability for Infringement Claims Against Third Party Covered Models**

IBM’s entire liability associated with its obligation to defend Client against claims that a Third Party Covered Model or its Output infringes a third party patent or copyright and pay amounts related to such claim as set forth in the applicable base agreement is limited to the amounts paid by Client that are attributable to the applicable Third Party Covered Model in the 12 months preceding the claim.

All other terms of the Liability and Intellectual Property Protection terms in the base agreement remain in effect.

5.2 HIPAA

The use of Health Information and Health data as Types of Personal Data and/or Special Categories of Personal Data (collectively, “Health Data”) is not permitted with this Cloud Service. In no event shall this Cloud Service be used for processing personal health information (“PHI”) as a health care clearinghouse within the meaning of Health Information Portability and Accountability Act of 1996 (“HIPAA”).

5.3 Anaconda Packages

The Cloud Service includes certain Anaconda-built data science and machine learning packages from the Anaconda Repositories which are made available from the Cloud Service for the convenience of Client. These Anaconda Packages are considered Non-IBM Products and are subject to their respective license agreement.

Client is also entitled to access the Anaconda Distribution (formerly known as Anaconda Individual Edition) (<https://www.anaconda.com/download>) and related services, and repositories collectively, the “Repositories”. Client’s use of the Anaconda Distribution and the Anaconda Packages are provided on an ‘as is’ basis as Non-IBM Products and are subject to their respective license agreement under the Anaconda End User License Agreement (<https://legal.anaconda.com/policies/en/>) and Anaconda Terms of Service (<https://legal.anaconda.com/policies/en/?name=terms-of-service#terms-of-service>), as applicable, without any restriction on (i) use for commercial purposes, and (ii) redistributing or embedding in a commercial product for redistribution or resale.

Access and use of the Anaconda Distribution and the Anaconda Packages is only for purposes of integration with IBM watsonx.ai Runtime, watsonx.ai Studio, the watsonx.ai Studio component in IBM Cloud Pak for Data, and IBM watsonx.ai offerings.

5.4 Technology Previews

IBM may make certain pre-release models or capabilities available to Client at no additional cost as a Preview Cloud Service for the purpose of evaluating its functionality and to provide feedback to IBM ("Technology Preview"). Use of Technology Previews in a production environment or for commercial purposes is not allowed. Technology Previews may have a limited set of features and may not be deployed in all data centers. Models available as Technology Preview are identified as such by the designation of "-rc" (release candidate) at the end of a model name; in the user interface next to the model name in the model gallery and in the Model Information; and in the API response.

Some Technology Preview models may be Non-IBM Products. By using a Technology Preview model that is a Non-IBM Product, Client accepts the applicable terms, including any associated obligations and use restrictions specified in the Model Information.

The Data Sheet as referenced in Section 2 does not apply to Technology Previews.

The Output IP Indemnity does not apply to Technology Previews.

Client is responsible for any Content input to the Technology Preview model and any decision to proceed with, or use, any output that are suggested by the Technology Preview model. IBM has no responsibility for any such output.

In exceptional circumstances it may be necessary for IBM to take steps to suspend or limit unreasonable use by Client of the Technology Preview.

5.5 Auto-generated Notebook

The auto-generated notebooks are subject to the International License Agreement for Non-Warranted Programs (or equivalent) and License Information document for Watson Studio Auto-generated Notebook ("License Terms"), such agreements located in the link below. Specifically, the Source Components and Sample Materials clause included in the License Information document for Watson Studio Auto-generated Notebook applies to the auto-generated notebooks. By downloading, copying, accessing, or otherwise using the materials, Client agrees to the License Terms found at: <http://www.ibm.com/terms/?id=L-AMCU-BYC7LF>.

5.6 Restricted Materials – Libraries and Models

The Cloud Service includes libraries and models ("IBM Materials") that may be used with Jupyter notebooks provided with the Cloud Service. IBM and its licensors are the sole owners of the IBM Materials. Client is permitted to use and modify the IBM Materials only in support of Client's internal use of the Cloud Service in accordance with the terms of the Service Description. Client may not reverse assemble, reverse compile, translate, or reverse engineer the IBM Materials, except as expressly permitted by law without the possibility of contractual waiver. The IBM Materials may not be used, copied, extracted from, or further distributed separately from the Cloud Service with which they were provided.

6. Overriding Terms

6.1 Synthetic Data Generator – Data Quality Terms

The following prevails over anything to the contrary in the base Cloud Service terms between the parties: IBM does not warrant the quality of synthetic data generated by the Cloud Service. Client acknowledges that the synthetic data is not real data and is produced from a model that requires the user to define the characteristics derived from a source data set selected by the Client, and that the monitoring, assessment, and curation of the data used as a source data set is an important aspect of the ethical and responsible use of artificial intelligence solutions.

6.2 OEM Commercial Models

The following takes precedence and prevails over any conflicting term in Section 5.1(a) Foundation and Embedding Models of this Service Description and the Liability and Intellectual Property Protection terms of the governing base agreement.

The Cloud Service includes access to Original Equipment Manufacturer ("OEM") Commercial Models which are considered part of the Cloud Service and are a Third Party Covered Model.

The third party retains ownership of the (i) the original, unmodified version of the OEM Commercial Model, and (ii) any customized and/or fine-tuned versions of such model (“**OEM Third Party**”), except for in the case of both (i) and (ii) the **Retained Rights** (as defined below).

“**Retained Rights**” means (i) Model Inputs, (ii) Generated Output, and (iii) any intellectual property rights developed by or on behalf of Client, including rights related to using or otherwise exploiting the OEM Commercial Models (as opposed to the Third Party Covered Models themselves).

“**Model Inputs**” are any data, content and software uploaded, exposed, made available, shared with, or input into an OEM Commercial Model (including for the avoidance of doubt a fine-tuned version of the OEM Commercial Model) for any purpose, including but not limited to, for fine-tuning, usage, training, improvement. Model Inputs are Client Content.

“**Generated Output**” means any text, images, and other outputs generated by an OEM Commercial model (including the fine-tuned version of the OEM Commercial Model) as a result of Client’s use of any version (including fine-tuned version) of the OEM Commercial Model.

Neither IBM nor OEM Third Party claims to have a license to or ownership of the Model Inputs or Generated Output.

Client has no obligation to provide IBM or OEM Third Party with any fine-tuned OEM Commercial Model.

Unless separately authorized by Client, neither IBM nor OEM Third Party will leverage, use (except to deliver the Cloud Service to the Client), or otherwise exploit the fine-tuned version of the OEM Commercial Model.

6.3 Custom Foundation Models

The following prevails over anything to the contrary in the Warranty and the Liability and Intellectual Property Protection sections of the base agreement between the parties:

Custom Foundation Models are considered Client Content. Client is solely responsible for all Custom Foundation Models, including without limitation, compliance with third party license terms and any associated obligations and use restrictions, assessing the suitability of the Custom Foundation Model for Client’s intended use, support, and for outputs generated by a Custom Foundation Model.

IBM Cloud Additional Service Description

IBM watsonx.ai Runtime as a Service

Except as noted below, the terms of the IBM Cloud Service Description apply.

1. Cloud Service

The functionality of the watsonx.ai platform is comprised of watsonx.ai Runtime (formerly known as Watson Machine Learning) as well as watsonx.ai Studio (formerly known as Watson Studio Paygo). If Client wishes to deploy the watsonx.ai platform, entitlements must be acquired to both watsonx.ai Runtime and watsonx.ai Studio.

1.1 Offerings

The Client may select from the following available offerings.

1.1.1 IBM watsonx.ai Runtime as a Service

IBM watsonx.ai Runtime is a managed service that enables training and deployment of machine learning models and integrating them in applications. The Cloud Service provides predictive analytics at scale for applications needing model training, batch, or online scoring. The Cloud Service allows developers to access machine learning technology using simple APIs without having to manage any infrastructure.

1.2 Optional Services

The following Cloud Services require IBM watsonx.ai Runtime as a prerequisite.

1.2.1 IBM watsonx.ai Foundation Models and Embedding Models

The Cloud Service enables access to IBM developed foundation models, and third-party foundation models as Non-IBM Product elements of watsonx.ai. The available foundation models are trained on a broad set of data and can be used for a variety of tasks including summarization, classification, content generation, and extraction.

The Cloud Service enables access to IBM developed embedding models, and third-party embedding models as Non-IBM Product elements of watsonx.ai. The available embedding models provide dense vector representations for text that can be used in various applications, most notably semantic search and retrieval augmented generation (RAG).

Client will be billed for the number of Resource Units consumed as described in Section 4.1.

1.2.2 IBM watsonx.ai Custom Foundation Model

The Cloud Service enables Clients to import Custom Foundation Models into a watsonx.ai deployment space for use within watsonx.ai. Custom Foundation Models are Client's own or third party foundation models that are trained and/or tuned externally to the Cloud Service. Not all model architectures are supported. Consult the Documentation for information on supported model architectures.

Client will be billed for the number of Hours the Custom Foundation Model is deployed.

1.2.3 IBM watsonx.ai Text Extraction

The Cloud Service enables access to text extraction methods to convert highly structured files and/or convey information in diagrams and tables into JSON or other file formats to make the information consumable by other features such as natural language interpretation.

Client will be billed for the number of Pages managed or processed by the Cloud Service.

2. Data Processing and Protection Data Sheets

The Data Sheet applicable for this service and the terms of this section provides the details and terms, including Client responsibilities, around use of this service. The following Data Sheet(s) apply to this service:

<https://www.ibm.com/terms/?id=850D9360405711E5B2E4A36A7B0C4479>

To view the compliance certification available by region, please consult the documentation:

<https://dataplatform.cloud.ibm.com/docs/content/wsj/getting-started/regional-datactr.html?context=wx&locale=en&audience=wdp>

3. Service Levels and Technical Support

3.1 Service Level Agreement

The service level agreement set forth in the base IBM Cloud Service Description applies to this service.

3.2 Technical Support

The support terms set forth in the base IBM Cloud Service Description apply to this service.

4. Charges

4.1 Charge Metrics

The charge metric(s) for the Cloud Service are specified in the Transaction Document.

The following charge metrics apply to this Cloud Service:

- Instance is each access to specific configuration of the Cloud Services.
- Resource Unit is an independent measure of a resource managed by, processed by, or related to the use of the Cloud Service.

For purposes of the foundation models and embedding models included with this Cloud Service, each 1000 Content Tokens, Data Points or Characters processed by a model is one (1) Resource Unit. Client will be billed the total number of Resource Units, multiplied by a model specific price per input or output inference.

 - A Content Token is a unit of input and output content such as individual words, and subwords of a sentence.
 - For foundation models, the total number of input Content Tokens and output Content Tokens processed during each month of a billing period are each rounded up to the nearest 1000.
 - For embedding models and for synthetic data generation, only output Content Tokens are counted and rounded up to the nearest 1000.
 - For alignment tuning, only input Content Tokens are counted and rounded up to the nearest 1000.
 - Character is any encoded symbol (such as a letter, number, icon, or space) that represents information. The number of input and/or output Characters processed by the Cloud Service during each month of a billing period are rounded up to the nearest 1000.
 - A Data Point is a unit of input and output content expressed as one or more numbers. The total number of input Data Points and output Data Points processed during each month of a billing period are each rounded up to the nearest 1000.
- Capacity Unit-Hour is an hour of use of the selected Cloud Services capacity type that will be multiplied by the specified number of required capacity units for such capacity type.

For the purpose of this Cloud Service, a minimum charge of 0.016 Capacity Unit-Hour will apply each time a Capacity Type is initiated.
- Hour is an hour of use of the Cloud Service.
- Page is a collection of content, converted to PDF if from a web page, contained on one side of a sheet of paper when printed, whether produced in physical or digital format, managed or processed by the Cloud Service.

For the purpose of the text extraction feature, the following file formats will be billed as follows:

 - PDF – 1 page in the pdf is 1 Page
 - HTML - up to 3,000 characters = 1 Page
 - Plain Text - up to 3,000 characters = 1 Page
 - Markdown - up to 3,000 characters = 1 Page
 - Word Doc - up to 3,000 characters = 1 Page
 - PPT - 1 slide = 1 Page
 - Image - 1 Page

In the case of file formats that do not have delineated pages, a Page will be considered as consisting of 3,000 characters.

4.2 Excess Use Charges

Client's entitlement to the Cloud Service may include specific limitations. If Client exceeds those limitations, Client will be charged for the excess use.

5. Additional Terms

For Cloud Service Agreements (or equivalent base cloud agreements) executed prior to January 1, 2019, the terms available at <https://www.ibm.com/acs> apply.

5.1 AI Technology Terms

a. Foundation and Embedding Models

The Cloud Service may include generative AI foundation models and/or embedding models as described in the documentation for the Cloud Service. Each foundation model and embedding model included with the Cloud Service will contain a model information document that provides information about that model ("Model Information"), including details on the model provider and applicable license terms. IBM-developed foundation and embedding models made available with the Cloud Service are considered part of the Cloud Service and cannot be used separately from the Cloud Service. Some models included with this Cloud Service are a Non-IBM Product and are separately licensed to the Client under the terms of the applicable third party license agreements(s) and not this agreement. By using a model that is a Non-IBM Product, Client accepts the applicable license terms, including any associated obligations and use restrictions specified in the third party license agreement and Model Information.

Client is prohibited from extracting the Model Weights of any model. "Model Weights" means the numerical values associated with the parameters of the training, optimization and performance of a model.

All Client inputs to a model, including by inferencing or prompting the model, and any content provided by the Client for further training of a model, including fine-tuning or prompt-tuning, are Client Content. Except as set forth in (c) (Output Indemnity) below, Client is solely responsible for Outputs generated by any model. IBM does not claim to have ownership rights to the Output. Unless separately authorized by Client through a voluntary feedback submission, IBM will not use the Client Content or model Outputs to train the foundation models or embedding models.

b. AI Use Restrictions

Client agrees not to, and not to direct or allow third parties, use foundation models in connection with the use of this Cloud Service : (i) for mass surveillance, racial profiling, or any use that violates or encourages the violation of basic human rights or other applicable laws and regulations; (ii) to distribute false, misleading, disparaging or obscene information or content; (iii) to provide fully automated decision making in connection use cases involving critical processes or the risk of loss of life, property or impact on an individual's legal rights; (iv) in a manner that impersonates another for deceptive purposes or conceals the fact a user is interacting with AI; or (v) to distribute or intentionally generate malware or other harmful code.

This Cloud Service is not intended for use in a high risk context as defined under the EU AI Act or other applicable regulations. If you chose to use it in such, please contact ChiefPrivacyOffice@ca.ibm.com.

c. Output IP Indemnity

IBM's obligation to protect Client against third party claims as set forth in the applicable base agreement includes third party claims that the result generated by the Covered Model in response to a Client's prompt or other input (collectively, "Output") infringes that third party's patent or copyright provided that (in addition to the exclusions and Client obligations provided in the relevant base agreement):

- (1) Client has a paid subscription to the Cloud Service or the Covered Model at the time Client generates the Output.
- (2) Client enables all available filters and logging features recommended by IBM in the product documentation or delivered with the Covered Model or in the Cloud Service, and Client

supplies information collected by such filters and features as requested by IBM in connection with a third party claim;

- (3) Client was using a supported version of the Covered Model;
- (4) the Output at issue was generated in connection with an Indemnified Use Case;
- (5) Client takes mitigating actions reasonably requested by IBM, which may include ceasing use of the allegedly infringing Output; and
- (6) Client had no reasonable basis to know that the Content or Output was infringing or likely to cause an infringement.

“Covered Model” means an IBM developed foundation model made available with the Cloud Service (“IBM Covered Model”) and a third-party foundation model provided with the Cloud Service or as a separately purchased entitlement to the Cloud Service and identified in the product documentation or in its separate transaction document as being an indemnified third party model (“Third Party Covered Model”).

“Indemnified Use Cases” means use within Client’s enterprise (a) of an IBM Covered Model in a manner consistent with such model’s intended use as described in its published Model Information or other Transaction Documents (such as license information or program specifications), and (b) of a Third Party Covered Model to analyze, summarize or classify Content or extract information or insights from Content provided that Client has accepted and complied with the applicable third party license.

All other terms of ‘(a) Foundation and Embedding Models’ in this section 5.1 and the Liability and Intellectual Property Protection terms in the base agreement remain in effect.

d. **Liability for Infringement Claims Against Third Party Covered Models**

IBM’s entire liability associated with its obligation to defend Client against claims that a Third Party Covered Model or its Output infringes a third party patent or copyright and pay amounts related to such claim as set forth in the applicable base agreement is limited to the amounts paid by Client that are attributable to the applicable Third Party Covered Model in the 12 months preceding the claim.

All other terms of the Liability and Intellectual Property Protection terms in the base agreement remain in effect.

5.2 HIPAA

Notwithstanding information in the applicable Data Sheet for this Cloud Service regarding the Health Information Portability and Accountability Act of 1996 (“HIPAA”) and the permitted use of Health Information and Health data as Types of Personal Data and/or Special Categories of Personal Data (collectively, “Health Data”), use of Health Data with this Cloud Service is subject to the following limitations and conditions:

When deployed in IBM Cloud Dallas region, only the following plan is eligible to be provisioned to implement the controls required under HIPAA Privacy and Security Rule for use with Health Data:

- IBM watsonx.ai Runtime HIPAA-Ready Plan

The Cloud Service may not be used for the transmission, storage, or other usage of any Health Data protected under HIPAA unless (i) Client notifies IBM in advance that Client will use Health Data with the Cloud Service; and (ii) IBM and Client have entered into an applicable Business Associate Agreement.

In no event shall the Cloud Service be used for processing PHI as a health care clearinghouse within the meaning of HIPAA.

5.3 Anaconda Packages

The Cloud Service includes certain Anaconda-built data science and machine learning packages from the Anaconda Repositories which are made available from the Cloud Service for the convenience of Client.

These Anaconda Packages are considered Non-IBM Products and are subject to their respective license agreement.

Client is also entitled to access the Anaconda Distribution (formerly known as Anaconda Individual Edition) (<https://www.anaconda.com/download>) and related services, and repositories collectively, the “Repositories”.

Client's use of the Anaconda Distribution and the Anaconda Packages are provided on an 'as is' basis as Non-IBM Products and are subject to their respective license agreement under the Anaconda End User License Agreement (<https://legal.anaconda.com/policies/en/>) and Anaconda Terms of Service (<https://legal.anaconda.com/policies/en/?name=terms-of-service#terms-of-service>), as applicable, without any restriction on (i) use for commercial purposes, and (ii) redistributing or embedding in a commercial product for redistribution or resale.

Access and use of the Anaconda Distribution and the Anaconda Packages is only for purposes of integration with IBM watsonx.ai Runtime, watsonx.ai Studio, the watsonx.ai Studio component in IBM Cloud Pak for Data, and IBM watsonx.ai offerings.

5.4 Federated Learning

Use of IBM Homomorphic Encryption Layers (Helayers) is limited to use in support of Federated Learning for aggregation of encrypted models.

5.5 Technology Previews

IBM may make certain pre-release models or capabilities available to Client at no additional cost as a Preview Cloud Service for the purpose of evaluating its functionality and to provide feedback to IBM ("Technology Preview"). Technology Previews may have a limited set of features.

Models available as Technology Preview are identified as such by the designation of "-rc" (release candidate) at the end of a model name; in the user interface next to the model name in the model gallery and in the Model Information; and in the API response.

Some Technology Preview models may be Non-IBM Products. By using a Technology Preview model that is a Non-IBM Product, Client accepts the applicable terms, including any associated obligations and use restrictions specified in the Model Information.

The Data Sheet as referenced in Section 2 does not apply to Technology Previews.

The Output IP Indemnity does not apply to Technology Previews.

Client is responsible for any Content input to the Technology Preview model and any decision to proceed with, or use, any output that are suggested by the Technology Preview model. IBM has no responsibility for any such output.

In exceptional circumstances it may be necessary for IBM to take steps to suspend or limit unreasonable use by Client of the Technology Preview.

6. Overriding Terms

6.1 Mistral Commercial Models

The following takes precedence and prevails over any conflicting term in Section 5.1(a) Foundation and Embedding Models of this Additional Service Description and the Liability and Intellectual Property Protection terms of the governing base agreement.

The Cloud Service includes access to Mistral commercial models which are considered part of the Cloud Service and are a Third Party Covered Model.

Mistral retains ownership of the (i) the original, unmodified version of the Mistral model, and (ii) any customized and/or fine-tuned versions of such model, except for in the case of both (i) and (ii) the **Retained Rights** (as defined below).

"Retained Rights" means (i) Model Inputs, (ii) Generated Output, and (iii) any intellectual property rights developed by or on behalf of Client, including rights related to using or otherwise exploiting the Mistral models (as opposed to the Mistral models themselves).

"Model Inputs" are any data, content and software uploaded, exposed, made available, shared with, or input into a Mistral model (including for the avoidance of doubt a fine-tuned Mistral model) for any purpose whatsoever, including but not limited to, for fine-tuning, usage, training, improvement. Model Inputs are Client Content.

"Generated Output" means any text, images, and other outputs generated by a Mistral model (including a fine-tuned Mistral Model) as a result of Client's use of the Mistral model (including a fine-tuned Mistral model).

Neither IBM nor Mistral claims to have a license to or ownership of the Model Inputs or Generated Output.

Client has no obligation to provide IBM or Mistral with any fine-tuned Mistral model.

Unless separately authorized by Client, neither IBM nor Mistral will leverage, use (except to deliver the Cloud Service to the Client), or otherwise exploit the fine-tuned Mistral model.

6.2 Custom Foundation Models

The following prevails over anything to the contrary in the Warranty and the Liability and Intellectual Property Protection sections of the base agreement between the parties:

Custom Foundation Models are considered Client Content. Client is solely responsible for all Custom Foundation Models, including without limitation, compliance with third party license terms and any associated obligations and use restrictions, assessing the suitability of the Custom Foundation Model for Client's intended use, support, and for outputs generated by a Custom Foundation Model.

IBM Cloud Additional Service Description

IBM watsonx.ai Studio as a Service

Except as noted below, the terms of the IBM Cloud Service Description apply.

1. Cloud Service

The functionality of the watsonx.ai platform is comprised of watsonx.ai Studio (formerly known as Watson Studio Paygo) as well as watsonx.ai Runtime (formerly known as Watson Machine Learning). If Client wishes to deploy the watsonx.ai platform, entitlements must be acquired to both watsonx.ai Studio and watsonx.ai Runtime.

1.1 Offerings

The Client may select from the following available offerings.

1.1.1 IBM watsonx.ai Studio as a Service

This Cloud Service is an integrated development environment offering a suite of tools and capabilities that is designed to assist data scientists to enhance their productivity.

This Cloud Service allows the Client to analyze data using RStudio, Jupyter notebooks, SPSS Modeler, Data Refinery and other tools in a configured, collaborative environment.

RStudio is integrated in the offering and provides a development environment for working with R.

This Cloud Service:

- provides Jupyter notebooks which are a web-based environment for interactive computing. Client is able to run small pieces of code that processes data, and then view the results of that computation inside of the notebook;
- includes Projects that allow Client to orient a team of collaborators around a set of notebooks, data sets, articles, models, and analysis work streams;
- provides the ability to create Projects that include data flow execution capabilities along with sampling and profiling; and
- includes Decision Optimization as a capability, for the user, to write models, using docplex API in Jupyter Notebooks, or using a graphical interface, to create or import an existing model. Models supported are LP/MIP models, CPO models, OPL models, or Python models. CPLEX and CP Optimizer engines are embedded in watsonx.ai Studio.

1.1.2 IBM watsonx.ai Synthetic Data Generator

The Cloud Service enables the generation of synthetic tabular data with statistical properties designed to reflect a Client's existing data set or a custom data schema, with the ability to anonymize table columns, preview the synthetic output, and export the output to an existing database or to a supported file type.

2. Data Processing and Protection Data Sheets

The Data Sheet applicable for this service and the terms of this section provides the details and terms, including Client responsibilities, around use of this service. The following Data Sheet(s) apply to this service:

<https://www.ibm.com/terms/?id=95E9BEA0B35711E7A9EB066095601ABB>

To view the compliance certification available by region, please consult the documentation:

<https://datapatform.cloud.ibm.com/docs/content/wsj/getting-started/regional-datactr.html?context=wx&locale=en&audience=wdp>

3. Service Levels and Technical Support

3.1 Service Level Agreement

The service level agreement set forth in the base IBM Cloud Service Description applies to this service.

3.2 Technical Support

The support terms set forth in the base IBM Cloud Service Description apply to this service.

4. Charges

4.1 Charge Metrics

The charge metric(s) for the Cloud Service are specified in the Transaction Document.

The following charge metrics apply to this Cloud Service:

- Authorized User is a unique user authorized to access to the Cloud Services in any manner directly or indirectly (for example, through a multiplexing program, device or application server) through any means.
- Instance is each access to specific configuration of the Cloud Services.
- Capacity Unit-Hour is an hour of use of the selected Cloud Services capacity type that will be multiplied by the specified number of required capacity units for such capacity type.

For the purposes of this Cloud Service, a minimum charge of 0.16 Capacity Unit Hours will apply each time a Capacity Type is initiated.

4.2 Partial Monthly Charges

Each Instance is billed on a monthly basis. Partial month deployment/usage will be pro-rated.

5. Additional Terms

For Cloud Service Agreements (or equivalent base cloud agreements) executed prior to January 1, 2019, the terms available at <https://www.ibm.com/acs> apply.

5.1 Auto-generated Notebook

The auto-generated notebooks are subject to the International License Agreement for Non-Warranted Programs (or equivalent) and License Information document for Watson Studio Auto-generated Notebook ("License Terms"), such agreements located in the link below. Specifically, the Source Components and Sample Materials clause included in the License Information document for Watson Studio Auto-generated Notebook applies to the auto-generated notebooks. By downloading, copying, accessing, or otherwise using the materials, Client agrees to the License Terms found at <http://www.ibm.com/terms/?id=L-AMCU-BYC7LF>.

5.2 Anaconda Packages

The Cloud Service includes certain Anaconda-built data science and machine learning packages from the Anaconda Repositories which are made available from the Cloud Service for the convenience of Client.

These Anaconda Packages are considered Non-IBM Products and are subject to their respective license agreement.

Client is also entitled to access the Anaconda Distribution (formerly known as Anaconda Individual Edition) (<https://www.anaconda.com/download>) and related services, and repositories collectively, the "Repositories".

Client's use of the Anaconda Distribution and the Anaconda Packages are provided on an 'as is' basis as Non-IBM Products and are subject to their respective license agreement under the Anaconda End User License Agreement (<https://legal.anaconda.com/policies/en/>) and Anaconda Terms of Service (<https://legal.anaconda.com/policies/en/?name=terms-of-service#terms-of-service>), as applicable, without any restriction on (i) use for Commercial Purposes, and (ii) redistributing or embedding in a commercial product for redistribution or resale.

Access and use of the Anaconda Distribution and the Anaconda Packages is only for purposes of integration with IBM watsonx.ai Runtime, watsonx.ai Studio, the watsonx.ai Studio component in IBM Cloud Pak for Data, and IBM watsonx.ai offerings.

5.3 Restricted Materials – Libraries and Models

The Cloud Service includes libraries and models ("IBM Materials") that may be used with Jupyter notebooks provided with the Cloud Service. IBM and its licensors are the sole owners of the IBM Materials.

Client is permitted to use and modify the IBM Materials only in support of Client's internal use of the Cloud Service in accordance with the terms of the Service Description. Client may not reverse assemble, reverse compile, translate, or reverse engineer the IBM Materials, except as expressly permitted by law without the possibility of contractual waiver. The IBM Materials may not be used, copied, extracted from, or further distributed separately from the IBM Cloud Service with which they were provided.

5.4 Synthetic Data Generator – Data Quality Terms

The following prevails over anything to the contrary in the base Cloud Service terms between the parties: IBM does not warrant the quality of synthetic data generated by the Cloud Service. Client acknowledges that the synthetic data is not real data and is produced from a model that requires the user to define the characteristics derived from a source data set selected by the Client, and that the monitoring, assessment, and curation of the data used as a source data set is an important aspect of the ethical and responsible use of artificial intelligence solutions.

