

Service Description

IBM BigInsights on Cloud

This Service Description describes the Cloud Service IBM provides to Client. Client means the company and its authorized users and recipients of the Cloud Service. The applicable Quotation and Proof of Entitlement (PoE) are provided as separate Transaction Documents.

1. Cloud Service

IBM BigInsights on Cloud provides pre-installed baremetal servers with IBM Open Platform and optional modules, ready to use. The base offering includes the following infrastructure features:

- Hardware servers,
- Networking infrastructure
- Internet facing firewall
- An operating system
- An LDAP server
- Metadata DB Hadoop Key Management Store

The Cloud Service is architected is based on IBM Open Platform with Apache Hadoop which includes::

- Hadoop including MapReduce, YARN and HDFS (Hadoop Distributed File System)
- Ambari cluster manager
- Avro
- HBase
- Hive
- Knox
- Oozie
- Open JDK
- Pig
- Parquet (Hadoop)
- Parquet (format)
- Spark
- Snappy
- Sqoop
- Slider
- Zookeeper

The Cloud Service provides data node and management node configurations. Both types of nodes are required for every cluster.

- a. Data nodes will hold data to be analyzed and run the programs developed by Client to analyze the data. The data (with the exception of any data IBM may, but is not committed to, provide, as part of the Cloud Service in the future and permitted software, programs and applications (see the Section titled "Client Obligations") will be provided by Client.
- b. Management nodes are comprised of services (like NameNode, Job Tracker, Management console, HBase Master etc), which help in managing and monitoring service on all the data nodes (like datanode, Hbase region, task tracker etc)

The Cloud Service is offered based on the configuration that best meets Client's needs. The available configurations of the nodes are:

Offering Type	Small Configuration	Medium Configuration	Large Configuration
Data Node	<ul style="list-style-type: none"> • 2 x 10 cores 2650 v3 Sandy bridge • 64 GB RAM • 16 TB (OS disks) • 20 TB (internal storage disks) 	<ul style="list-style-type: none"> • 2 x 10 cores 2650 v3 Sandy bridge • 128 GB RAM • 16 TB (OS disks) • 28 TB (internal storage disks) 	<ul style="list-style-type: none"> • 2 x 12 cores 2690 v3 Sandy bridge • 192 GB RAM • 16 TB (OS disks) • 32 TB (internal storage disks)
Management Node	<ul style="list-style-type: none"> • 2 x 10 cores 2650 v3 Sandy bridge • 64 GB RAM • 16 TB (OS disks) 	<ul style="list-style-type: none"> • 2 x 10 cores 2650 v3 Sandy bridge • 128 GB RAM • 16 TB (OS disks) 	<ul style="list-style-type: none"> • 2 x 12 cores 2690 v3 Sandy bridge • 256 GB RAM • 16 TB (OS disks)

1.1 Optional Services

1.1.2 IBM BigInsights on Cloud Data Scientist Module

This optional offering includes the following features:

- Big R
- Machine Learning
- Text Analytics
- Big SQL
- BigSheets

1.1.3 IBM BigInsights on Cloud Analyst Module

This optional offering includes the following features:

- Big SQL
- BigSheets

1.2 IBM Bluemix

IBM Bluemix, IBM's open standards cloud platform for building, running, and managing apps and services, is a technical pre-requisite for the Cloud Service. New users can register for access via the online registration form: <https://console.ng.bluemix.net/registration>.

2. Security Description

This Cloud Service follows IBM's data security and privacy principles for IBM SaaS which are available at <https://www.ibm.com/cloud/resourcecenter/content/80> and any additional terms provided in this section. Any change to IBM's data security and privacy principles will not degrade the security of the Cloud Service.

2.1 Personal Information and Regulated Data

This Cloud Service is not designed to any specific security requirements for regulated content, such as personal information or sensitive personal information. Client is responsible to determine if this Cloud Service meets Clients needs with regard to the type of content Client uses in connection with the Cloud Service. The Cloud Service may not be used for the transmission or storage of any Protected Health Information (PHI) protected under the Health Information Portability and Accountability Act of 1996 (HIPAA) unless (a) IBM and Client have entered into an applicable Business Associate Agreement, and (b) IBM provides Client with express written confirmation that the Cloud Service can be used with PHI.

2.2 Encryption

The primary method of encryption of user data is through the native HDFS encryption available in Hadoop. Client is responsible for identifying the data to be encrypted and executing the required steps to ensure that the data is encrypted.

3. Service Level Objective

IBM provides the following service level objective ("SLO") for the Cloud Service, after IBM makes the Cloud Service available to Client.

IBM will provide an SLO of 99.9% service availability measured by the following:

- a. The ability to deploy an application either through the portal or command line.
- b. The ability to connect to an application service
- c. The availability of an individual service instance

There are no service level credits or reporting offered at this time.

4. Technical Support

Technical support for the Cloud Service is provided via email, online forums, and an online problem reporting system which is available on the Client portal at <https://support.ibmcloud.com>. IBM will make available the IBM Software as a Service Support Handbook which provides technical support contact information and other information and processes. Technical support is offered with the Cloud Service and is not available as a separate offering.

Severity	Severity Definition	Response Time Objectives During Support Hours	Response Time Coverage
1	Critical business impact/service down: Business critical functionality is inoperable or critical interface has failed. This usually applies to a production environment and indicates an inability to access services resulting in a critical impact on operations. This condition requires an immediate solution. Severity 1 issues require that Client is available to help IBM diagnose issues during the 24X7 period otherwise the issue is downgraded to Severity 2.	Within 1 hour	24x7
2	Significant business impact: A service feature or function is severely restricted in its use or Client is in jeopardy of missing business deadlines.	Within 2 business hours	M-F business hours
3	Minor business impact: Indicates the service or functionality is usable and it is not presenting a critical impact on operations.	Within 4 business hours	M-F business hours
4	Minimal business impact: An inquiry or non-technical request.	Within 1 business day	M-F business hours

5. Entitlement and Billing Information

5.1 Charge Metrics

The Cloud Service is available under the charge metric specified in the Transaction Document:

- Instance is a unit of measure by which the Cloud Service can be obtained. An Instance is access to a specific configuration of the Cloud Service. Sufficient entitlements must be obtained for each Instance of the Cloud Service made available to access and use during the measurement period specified in Client's Proof of Entitlement (PoE).

5.2 Partial Month Charges

A partial month charge as specified in the Transaction Document may be assessed on a pro-rated basis.

6. Term and Renewal Options

The term of the Cloud Service begins on the date IBM notifies Client of their access to the Cloud Service, as documented in the PoE. The PoE will specify whether the Cloud Service renews automatically, proceeds on a continuous use basis, or terminates at the end of the term.

For automatic renewal, unless Client provides written notice not to renew at least 90 days prior to the term expiration date, the Cloud Service will automatically renew for the term specified in the PoE.

For continuous use, the Cloud Service will continue to be available on a month to month basis until Client provides 90 days written notice of termination. The Cloud Service will remain available to the end of the calendar month after such 90 day period.

7. Enabling Software

This Cloud Service includes enabling software, which may be used only in connection with Client's use of the Cloud Service and only for the Cloud Service term.

The following IBM software entitlements are included as enabling software under the terms of their applicable IBM program licenses, in addition to the limitations below:

InfoSphere DataStage

Entitlement: 280 PVUs (as defined in the License Information document for InfoSphere DataStage)

Use Restriction: Client may only use InfoSphere DataStage to populate data into the Cloud Service.

InfoSphere DataStage and QualityStage Designer

Entitlement: 2 Concurrent Users (as defined in the License Information document for InfoSphere DataStage and QualityStage Designer)

Use Restriction: Client may only use InfoSphere DataStage and QualityStage Designer to populate data into the Cloud Service.

8. Cloud Service Offering Additional Terms

8.1 Cookies

Client is aware and agrees that IBM may, as part of the normal operation and support of the Cloud Service, collect personal information from Client (your employees and contractors) related to the use of the Cloud Service, through tracking and other technologies. IBM does so to gather usage statistics and information about effectiveness of our Cloud Service for the purpose of improving user experience and/or tailoring interactions with Client. Client confirms that it will obtain or have obtained consent to allow IBM to process the collected personal information for the above purpose within IBM, other IBM companies and their subcontractors, wherever we and our subcontractors do business, in compliance with applicable law. IBM will comply with requests from Client's employees and contractors to access, update, correct or delete their collected personal information

8.2 Twitter Data

8.2.1 Definitions

"Analysis Report" means the results or output created by Client's applications on the Cloud Service from the process of analyzing and deriving information from the Twitter Content. Such results may not contain Twitter Content, but may contain content that is duplicative to that in the Twitter Content, such as search terms and references to the subject matter of the Tweets.

"Tweet ID" means the unique identification number generated by Twitter for each Tweet.

"Tweets" means a public posting with a text body of no more than 140 characters made by any end user of the Twitter service "Twitter Content" means Tweets and Tweet IDs, Twitter end user public profile information, and any derivative works thereof.

8.2.2 Access and Use of Twitter Content

The Cloud Service provides a means for Client to select and access Twitter Content for Client's use solely within the Cloud Service. Twitter Content is neither owned nor controlled by IBM. Twitter Content may include materials that are illegal, inaccurate, misleading, indecent, or otherwise objectionable. IBM or its suppliers have no obligation to review, filter, verify, edit or remove any Twitter Content. However, IBM or its suppliers may, at their sole discretion, do so. Client agrees, upon IBM's request, to promptly delete or modify any Tweet that may be stored on the Cloud Service.

8.2.3 Volume of Twitter Data

If Client has acquired entitlements to a Cloud Service configuration that includes a minimum of fifteen (15) nodes as part of the Cloud Service, then Client may be given access to Twitter Content.

8.2.4 Restrictions on Twitter Content

Client shall not:

- a. Use Twitter Content in violation of any applicable law including but not limited to privacy laws or for any unauthorized or improper purposes.
- b. Download or remove Twitter Content from the Cloud Service. However, Analysis Reports may be downloaded or removed from the Cloud Service.
- c. Make available to any third party any Twitter Content included in the Cloud Service.
- d. Use the Twitter Content for any other purpose except for the limited purpose allowed for in the Cloud Service.
- e. Aggregate, cache, or store location data or geographic information contained in the Twitter Content separately from the Tweet to which it is associated, or use the location data or geographic data included in the Twitter Content for any reason except to identify the location tagged in the Tweet.
- f. Comingle Twitter Content with other data, unless the Twitter Content will always be clearly attributable to Twitter.
- g. Use Twitter Content provided as part of the Cloud Service to perform analysis on a small group of individuals or a single individual for any unlawful or discriminatory purpose.
- h. Display the Twitter Content that Client receives access to as a part of Cloud Service.
- i. Use the Twitter Content or analysis therefrom in the Cloud Service as part of an advertising network.
- j. Use the Twitter Content for the purpose of performing analysis that is a regularly-produced, time-based series of measurements made using the same, or similar, methodologies for the purpose of comparing television program performance over time, or against a defined set or subset of other television programs.
- k. Use any aggregate Twitter user metrics, such as number of users or accounts, obtained while accessing and using the Twitter Content as part of the Cloud Service for any purpose.

8.2.5 Termination by IBM

Client's access and use of the Twitter Content shall cease upon termination of the Cloud Service. Notwithstanding the foregoing, in addition to the rights of suspension and termination in the Agreement, IBM may cease providing access to the Twitter Content at any time without notice and without the obligation to provide Client a refund, credit, or other compensation.

8.2.6 Warranty and Indemnification Disclaimer for Twitter Content

NOTWITHSTANDING THE WARRANTY SET FORTH IN THE AGREEMENT, THE TWITTER CONTENT IS PROVIDED SOLELY "AS IS", "AS AVAILABLE" WITH ALL FAULTS, AND CLIENT'S USE OF THE TWITTER CONTENT IS AT ITS SOLE RISK. IBM DOES NOT MAKE, AND HEREBY DISCLAIMS, ANY AND ALL OTHER EXPRESS AND IMPLIED WARRANTIES, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, PERFORMANCE, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, TITLE, AND ANY WARRANTIES ARISING FROM COURSE OF DEALING, USAGE, OR TRADE PRACTICE, IN CONNECTION WITH THE TWITTER CONTENT. IBM DOES NOT WARRANT THAT THE ACCESS TO THE TWITTER CONTENT WILL BE UNINTERRUPTED, OR ERROR-FREE. THIS DISCLAIMER OF WARRANTY MAY NOT BE VALID IN SOME JURISDICTIONS AND CLIENT MAY HAVE WARRANTY RIGHTS UNDER LAW WHICH MAY NOT BE WAIVED OR DISCLAIMED. ANY SUCH WARRANTY EXTENDS ONLY FOR THIRTY (30) DAYS FROM THE EFFECTIVE DATE OF THIS AGREEMENT (UNLESS SUCH LAW PROVIDES OTHERWISE).

ANY OBLIGATION FOR IBM TO INDEMNIFY CLIENT UNDER THE AGREEMENT DOES NOT APPLY IN ANY WAY TO CLIENT'S ACCESS AND USE OF THE TWITTER CONTENT.

8.2.7 United States Government Use

The Twitter Content is a "commercial item" as that term is defined at 48 C.F.R. 2.101, consisting of "commercial computer software" and "commercial computer software documentation" as such terms are used in 48 C.F.R. 12.212. Any use, modification, derivative, reproduction, release, performance, display, disclosure or distribution of the Twitter Content by any government entity is prohibited, except as expressly permitted by the terms of this ToU. Additionally, any use by U.S. government entities must be in accordance with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4. If Client uses the Twitter Content in Client's official capacity as an employee or representative of a U.S., state or local government entity and Client is legally unable to accept the jurisdiction, venue or other clauses herein,

then those clauses do not apply to such entity, but only to the extent as required by applicable law. Contractor/manufacture is Twitter, Inc. 1355 Market Street, Suite 900, San Francisco, California 94103

8.3 Client Obligations

Client is responsible for:

- following IBM's initial configuring of the IBM Open Platform components and any BigInsights modules selected, monitoring, configuring and managing the IBM Open Platform components which run under Ambari as well as the BigInsights modules. Examples of these components and software include, but are not limited to, HBase, Hive, Zookeeper, Oozie, BigSQL, BigSheets, Big R etc. Client has flexibility to choose which components to run, but is therefore responsible for monitoring and stopping or starting those components.
- limiting access to members of their organization to the Cloud Service by sharing the URL, username and password available after the Cloud Service cluster is provisioned and also granting and revoking access (through entries in LDAP)
- developing, as necessary or required, programs and applications on the service to analyze data and obtain insights from it. Client is also responsible for the quality and performance of such programs or applications developed.
- use and maintenance of any permissible software or data that is added by or on behalf of Client to the cluster. IBM may provide assistance, but will not maintain, move or remove Client added permissible software and data that fills a partition or affects the functioning of the service.
- Using the Hadoop data encryption feature available on the SaaS to encrypt data as per business requirements.
- Checking the following link periodically to be informed about a scheduled non-disruptive or disruptive downtime - <https://developer.ibm.com/bluemix/support/#status>
- backing up all data, metadata, configuration files and environment parameters as per business requirements so as to ensure continuity
- restoring data, metadata, configuration files and environment parameters from any backup, to ensure continuity, in an eventuality of cluster failure of any type including but not limited to data center or pod failure, server failure or hard disk failure or software failures.
- ensuring continuity, compatibility and performance of the Cloud Service platform after installing permissible software (including any open source packages), or upgrading IBM Open Platform components or BigInsights software.

Client may:

- only install and run software on the Cloud Service which is required to perform or support analysis of data through the pre-installed IBM Open Platform or BigInsights software. Client is prohibited from installing or using any other software of any kind on the Cloud Service.
- install additional open source packages for use with the Cloud Service on the nodes. IBM is not obligated to support these packages and is not responsible for impacts to the performance of the Cloud Service due to any such additional packages. If Client installed software is determined by IBM to compromise security or security compliance guidelines of the infrastructure, server or environment, it may be disabled or removed by IBM.

8.4 IBM Operational Responsibilities

As part of the Cloud Service, IBM will:

- provide and manage servers, storage and networking infrastructure for the cluster.
- provide an initial configuration of the IBM Open Platform components and any BigInsights modules selected
- provide and manage an internet facing and internal firewall for protection and isolation.
- monitor and manage the following components on the Cloud Service:
 - (1) Network components
 - (2) Servers and their local storage
 - (3) Operating System

- (4) Hadoop Key Management Store
 - (5) LDAP on the Hadoop cluster
 - (6) Nagios
 - (7) Ganglia
 - (8) Ambari cluster manager
- provide maintenance patches, including appropriate security patches for the operating system, the IBM Open Platform, and any BigInsights modules selected (but not for any software or components separately added by the Client).

Regular maintenance that should not require any system downtime (“non-disruptive” maintenance) and maintenance that may require some system downtime and restarting (“disruptive” maintenance”) will be performed at the scheduled times published at: <https://developer.ibm.com/bluemix/support/#status>. Any changes to the scheduled maintenance times will be posted with at least 24 hours advance notice. Client will be notified of additional scheduled downtimes at least 24 hours in advance.