

## Service Description

### IBM Watson Data Platform Plan

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

Client may deploy any combination of the listed services up to the monetary value of Client's Cloud Service entitlement(s).

IBM Watson Data Platform Plan Standard provides Client access to the data and analytics services listed below:

- a. **IBM Analytics for Apache Spark** - An Apache Spark service with built-in machine learning, an integrated notebook service and connectors for common data sources
  - IBM Analytics for Apache Spark Reserved Enterprise
- b. **IBM BigInsights on Cloud** - An Apache Hadoop development platform as a service
  - IBM Open Platform with Apache Spark and Apache Hadoop on Cloud Small Data Node
  - IBM Open Platform with Apache Spark and Apache Hadoop on Cloud Medium Data Node
  - IBM Open Platform with Apache Spark and Apache Hadoop on Cloud Large Data Node
  - IBM Open Platform with Apache Spark and Apache Hadoop on Cloud Small Management Node
  - IBM Open Platform with Apache Spark and Apache Hadoop on Cloud Medium Management Node
  - IBM Open Platform with Apache Spark and Apache Hadoop on Cloud Large Management Node
  - IBM BigInsights on Cloud
- c. **IBM Bluemix Lift** – Quickly, securely and reliably migrate database from on-premises data center to an IBM Bluemix cloud data property
  - IBM Bluemix Lift Subscription One Data Target Unlimited Inbound Data
- d. **IBM Compose** - Run web and mobile apps on fully managed, hand-picked open source databases with an integrated database-as-a-platform service
  - IBM Compose Enterprise Transactional for Softlayer
  - IBM Compose Enterprise Large Transactional for Softlayer
  - IBM Compose Enterprise Transactional for AWS
  - IBM Compose Enterprise Large Transactional for AWS
- e. **IBM Cloudant** - A NoSQL database platform built for the cloud
  - IBM Cloudant Dedicated SMB Node
  - IBM Cloudant Bare Metal Node
  - IBM Cloudant Virtual Machine Node
  - IBM Cloudant Dedicated I2 Virtualized Node
- f. **IBM dashDB** - A cloud data warehouse, purpose-built for analytics
  - IBM dashDB Enterprise 64.1 Virtual Machine
  - IBM dashDB Enterprise 256.4 Bare Metal
  - IBM dashDB Enterprise MPP 4TB Bare Metal
  - IBM dashDB Enterprise for Transactionas 12.128.1400
  - IBM dashDB Enterprise for Transactionas 2.8.500
  - IBM dashDB Enterprise MPP 32.244.1400 for AWS

- g. **IBM Bluemix Data Connect** - Self-service data preparation and data movement service for developers and IT architects.
  - IBM Bluemix Data Connect Personal
  - IBM Bluemix Data Connect Professional
  - IBM Bluemix Data Connect Enterprise
- h. **IBM Streaming Analytics** - Ingest, analyze, monitor, and correlate data as it arrives from real-time data sources.
  - IBM Streaming Analytics Standard Plan

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.

## 1.1 **IBM DataFirst Watson Data Platform Plan Advisory Workshop:**

IBM DataFirst Bluemix Data and Analytics Plan Advisory Workshop is included with the Cloud Service. The workshop will include service overviews, a review of timing, milestones, dates, Client resource commitments, finalization of use cases, scoping of data migration and ongoing data source requirements, the completion of sizing, scoping requirements, provisioning forms and a playback and confirmation by the Client of the overall provisioning of the Cloud Service. Client will be required to provide IBM all required provisioning forms prior to the completion of the workshop.

The workshop may include a combination of on-site activities at Client's facility and remotely delivered activities, all of which not to exceed more than 40 hours total. The workshop expires 60 calendar days from the date of the purchase of entitlement(s) for the Cloud Service regardless of whether all hours have been used.

### **Client's Responsibilities for on-site activities:**

IBM's performance is dependent upon Client's management and fulfillment of its responsibilities as specified below, at no charge to IBM. Any delay in performance of these responsibilities may impact IBM's ability to provision the Cloud Service. Client will:

- a. provide IBM with the address and room details of the Client's facility where the workshop will take place;
- b. provide a suitable meeting room for the number of attendees, including an adequate number of power sockets, an overhead projector and whiteboard and/or flipchart and connectivity to necessary systems and environments, including Internet connectivity if required.
- c. If applicable, provide safe access, suitable office space, supplies, furniture, high speed connectivity to the Internet, and other facilities for IBM's personnel while working at Client's location;
- d. be responsible for the secure storage of all IBM-owned hardware and software while on Client's premises and ensure all rooms containing IBM equipment are locked at all times.

## 2. **Optional Entitlement**

### 2.1 **IBM Watson Data Platform Plan Standard Add-On**

Client may purchase a subscription to IBM Watson Data Platform Plan Standard Add-On which provides the ability to deploy more of the data and analytics services listed in section 1.

## 3. **Security Description**

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

## 4. **Technical Support**

Technical support for the Cloud Service is provided via online forums and an online problem reporting system.. 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	<p>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.</p> <p>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.</p>	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 in Client's PoE or Transaction Document.

### 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. Cloud Service Offering Additional Terms

### 7.1 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 the Client (Client's 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 the Client. Client confirms that they 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.

## **7.2 IBM Bluemix**

IBM Bluemix is a technical pre-requisite for utilizing the BigInsights on Cloud service. New users can register for access via the online registration form: <https://console.ng.bluemix.net/registration/>.

## **7.3 Terms applicable to Cloud Service for AWS**

If Client deploys a service that is specified as "for AWS", the following terms apply:

The Cloud Service application layer and Client's data and content are hosted on a third party cloud services infrastructure and platform not managed by IBM. The Cloud Service infrastructure, certain aspects of the Cloud Service platform, and related services, including: data center, servers, storage and network; application and data backup; firewall and threat detection; and APIs for application deployment, monitoring and operation (collectively, the "Third Party Cloud Platform Services") are hosted and managed by the third party provider. Accordingly, notwithstanding anything in this Service Description or the base services agreement under which this Cloud Service is provided (e.g. the IBM Cloud Services Agreement) (the "Base Agreement"):

- a. Obligations of IBM relating to data security and data protection in the Base Agreement, in the description of the Cloud Service's security practices in Section 2 above (Security Description), and in the referenced IBM Data Security and Privacy Principles, do not apply to the Third Party Cloud Platform Services or to the Cloud Service to the extent dependent on the Third Party Cloud Platform Services. The Cloud Service may not be used for the transmission, storage or processing of any Protected Health Information or of any personal data of individuals residing in the European Union.
- b. If the provider of the Third Party Cloud Platform Services notifies IBM that it has withdrawn or terminated its services or IBM's or Client's access to such services, IBM may terminate the Cloud Service immediately upon the effective date of such termination by the third party provider by providing notice of termination to Client.
- c. IBM makes no warranties or conditions, express or implied, and IBM will have no liability to Client, regarding the Third Party Cloud Platform Services or to the Cloud Service to the extent dependent on the Third Party Cloud Platform Services.
- d. Client agrees to indemnify, defend and hold IBM harmless from and against any claims, damages, losses, liabilities, costs, and expenses (including reasonable attorneys' fees) arising out of or relating to claims against IBM by the Third Party Cloud Platform Services provider) concerning: (a) Client's use of the Cloud Services; (b) breach of this Service Description, the Base Agreement or violation of applicable law by Client; (c) Client's content or the combination of Client's content with other applications, content or processes, including any claim involving alleged infringement or misappropriation of third-party rights by Client's content or by the use, development, design, production, advertising or marketing of Client's content; or (d) a dispute between IBM and Client.

## **7.4 Provisioning and Configuration Modification:**

If Client has not provided IBM with the required provisioning information within 60 days of the purchase date, IBM will use default technical data and Client profile information to configure the Cloud Service for immediate access. IBM will notify Client with provisioning details and Client's PoE will specify the term start date.

The Cloud Service deployed configuration may be changed once per 30 calendar days upon Client's written request and will be provisioned by IBM within 12 calendar days.

## **8. Offering Specific Terms**

The following sections outline any terms associated with each of the individual offerings available as part of the Cloud Service. In the event of conflict with Sections 1 through 6 of the Service Description, the below terms apply. For the purposes of the subsections below, references to Cloud Service refer to the individual offering, not the IBM Bluemix Data and Analytics Plan.

### **8.1 IBM Analytics for Apache Spark**

IBM Analytics for Apache Spark provides a managed Spark environment accessible on-demand. With this service, Clients get access to Apache Spark's next-generation capabilities. Users can start with a small amount of capacity for experimentation, increase capacity for development and then scale to production,

all within the same environment. The service is immediately ready for analysis, skipping setup hurdles thus providing a quick start to Client's Spark efforts.

The Cloud Service defines capacity in the form of a fixed quantity of Spark Executors. Each Spark Executor is allocated a slice of computing resources by the Spark cluster, with maximum specifications of 12.5GB of memory and 1 physical CPU core (or equivalent) of processor. Capacity can be increased by adding multiple instances of the Cloud Service. The IBM Analytics for Apache Spark Reserved Enterprise offering includes a maximum of 30 Spark Executors.

## 8.2 IBM BigInsights on Cloud

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 based on open source packages from IBM Open Platform with Apache Spark and Apache Hadoop.

- a. IBM BigInsights on Cloud provides data nodes, management node configurations. Management and data nodes are required for every cluster. 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
<b>Data Node</b>	<ul style="list-style-type: none"> <li>● 2 x 10 cores 2650 v3 Sandy bridge</li> <li>● 64 GB RAM</li> <li>● 16 TB (OS disks)</li> <li>● 20 TB (internal storage disks)</li> </ul>	<ul style="list-style-type: none"> <li>● 2 x 10 cores 2650 v3 Sandy bridge</li> <li>● 128 GB RAM</li> <li>● 16 TB (OS disks)</li> <li>● 28 TB (internal storage disks)</li> </ul>	<ul style="list-style-type: none"> <li>● 2 x 12 cores 2690 v3 Sandy bridge</li> <li>● 256 GB RAM</li> <li>● 16 TB (OS disks)</li> <li>● 32 TB (internal storage disks)</li> </ul>
<b>Management Node</b>	<ul style="list-style-type: none"> <li>● 2 x 10 cores 2650 v3 Sandy bridge</li> <li>● 64 GB RAM</li> <li>● 16 TB (OS disks)</li> </ul>	<ul style="list-style-type: none"> <li>● 2 x 10 cores 2650 v3 Sandy bridge</li> <li>● 128 GB RAM</li> <li>● 16 TB (OS disks)</li> </ul>	<ul style="list-style-type: none"> <li>● 2 x 12 cores 2690 v3 Sandy bridge</li> <li>● 256 GB RAM</li> <li>● 16 TB (OS disks)</li> </ul>

### 8.2.1 IBM BigInsights on Cloud

This offering includes the following features:

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

## 8.2.2 Twitter Data

### a. 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.

### b. 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.

### c. 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.

### d. Restrictions on Twitter Content

Client shall not:

- (1) Use Twitter Content in violation of any applicable law including but not limited to privacy laws or for any unauthorized or improper purposes.
- (2) Download or remove Twitter Content from the Cloud Service. However, Analysis Reports may be downloaded or removed from the Cloud Service.
- (3) Make available to any third party any Twitter Content included in the Cloud Service.
- (4) Use the Twitter Content for any other purpose except for the limited purpose allowed for in the Cloud Service.
- (5) 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.
- (6) Comingle Twitter Content with other data, unless the Twitter Content will always be clearly attributable to Twitter.
- (7) 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.
- (8) Display the Twitter Content that Client receives access to as a part of Cloud Service.
- (9) Use the Twitter Content or analysis therefrom in the Cloud Service as part of an advertising network.
- (10) 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.
- (11) 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.

### e. 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.

f. **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.

g. **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/manufacturer is Twitter, Inc. 1355 Market Street, Suite 900, San Francisco, California 94103.

### 8.2.3 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 Cloud Service 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.

#### 8.2.4 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) 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.

#### 8.2.5 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.



## 8.2.6 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.

## 8.2.7 Personal Information and Regulated Data

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

## 8.3 IBM Bluemix Lift

IBM Bluemix Lift is a cloud based database service which allows Client to migrate its database from a private, on-premises data center to a cloud data property.

### 8.3.1 IBM Bluemix Lift Subscription One Data Target Unlimited Inbound Data

The Cloud Service supports an unlimited amount of data transmitted inbound to an IBM cloud data offering target. Each Instance of the Cloud Service may only be used with a single cloud database target instance or single massively parallel processing (MPP) target database cluster. The Client may not use a single Cloud Service Instance with more than one database target instance or target cluster. Additionally, the Cloud Service only operates in conjunction with IBM cloud database offerings.

### 8.3.2 Enabling Software

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

- IBM Secure Gateway for Bluemix
- IBM Bluemix Lift Slingshot

## 8.4 IBM Compose

The Cloud Service provides the Compose Data Platform as an IBM managed service hosted on Softlayer. An Instance consists of a dedicated 3-node cluster on which the Client can provision and operate any combination of the supported data services up to the capacity of the Instance including, but not limited to, the following:

- MongoDB
- Elasticsearch
- Redis
- PostgreSQL
- RethinkDB
- etcd
- RabbitMQ

Clients may also integrate the Cloud Service with 3rd party services, including, but not limited to, the following:

- New Relic – The Cloud Service utilizes MeetMe Inc's agent to integrate with New Relic's plugin API to allow in depth analytics and monitoring. This feature is available for MongoDB, Redis, PostgreSQL and Elastic Search. The Client is responsible for registering with New Relic directly.
- Syslog-NG - This feature forwards logs to an external log provider via Syslog-NG.

The Cloud Service is offered in the following configurations:

### 8.4.1 IBM Compose Enterprise Transactional for Softlayer

A 3 node cluster with each node as follows. Bare metal server with Intel Xeon E5-2690 8 Cores, 2.90 GHz, 16 vCPUs, 64GB RAM, 2x 960GB SSD in RAID1, Bonded 2x1Gbps Public and Private Network Uplinks.

#### **8.4.2 IBM Compose Enterprise Large Transactional for Softlayer**

A 3 node cluster with each node as follows. Bare metal server with Dual Intel Xeon E5-2690 v3, 12 Cores, 2.60 GHz, 256GB RAM, 2x 1TB SATA in RAID1 for OS, 4 TB RAID 10 SSD, Bonded 2x10Gbps Public and Private Network Uplinks.

#### **8.4.3 IBM Compose Enterprise Transactional for AWS**

A 3 node cluster with each node as follows. Virtual r3.2xlarge with Intel Xeon E5-2670 v2 (Ivy Bridge) Processors 8 vCPUs, 61GB RAM, 915GB GP SSD EBS.

#### **8.4.4 IBM Compose Enterprise Large Transactional for AWS**

A 3 node cluster with each node as follows. Virtual private server with r3.8xlarge Intel Xeon E5-2670 v2 (Ivy Bridge) Processors, 32 vCPUs, 244GB RAM, 3.66TB GP SSD.

#### **8.4.5 Cloud Service Capacity**

The capacity of an Instance depends on the type and number of service Deployments, where a Deployment is defined as an individual instance of a supported data service deployed to an Instance of the Cloud Service. Each data service has a set ratio between allocated RAM, Disk and IOPs and the HA replication architecture is service dependent. MongoDB, Elasticsearch, PostgreSQL and RethinkDB have a ratio of .1GB RAM per 1GB of disk; Redis has a ratio of 256MB RAM per 256MB of disk; etcd and RabbitMQ have a ratio of 256MB RAM per 1GB of disk.

New service Deployments can be instantiated until 80% of the Instance RAM has been allocated to existing Deployments. At that point existing Deployments can continue to grow as long as RAM is available.

An interactive sizing tool is available to calculate capacity based on the number and type of databases deployed.

#### **8.4.6 Access to Public Pay As You Go Services**

Client will have access to the Compose.io Pay As You Go (PayGo) services from the Cloud Service. These PayGo services are not included in the Cloud Service and are governed by the Compose, Inc Terms of Service available at <https://help.compose.io/docs/terms-of-service/>. Any usage of PayGo services will be billed directly by Compose, Inc, an IBM Company.

### **8.5 IBM Cloudant Dedicated Cluster**

The following base features are included in each of the node configurations described below:

- Client receives access to IBM Cloudant Data Layer software, a JavaScript Object Notation (JSON) database. The software includes the ability to store and retrieve data via one (1) Cloudant RESTful API, access to client and mobile SDK libraries, built-in full-text indexing, multi-master replication, and real time analytics via MapReduce.
- For each Instance of the Cloud Service, IBM will install and configure IBM Cloudant Data Layer software components on one (1) server. Each server will be located in a supported Infrastructure as a Service (IaaS) provider's data center. Client can select the IaaS provider and the data center locations from a list defined by IBM.
- Client must acquire a minimum of three (3) Instances per cluster. IBM will configure the associated servers into one or more clusters. IBM, at its sole discretion, will configure one or more load balancers to control system workloads.

#### **8.5.1 IBM Cloudant Bare Metal Node**

In addition to the base features described above, the following applies to this Cloud Service:

- For each Instance of the Cloud Service, Client receives access to one (1) server dedicated to their activities which is not shared with any other Client.

#### **8.5.2 IBM Cloudant Dedicated SMB Node**

In addition to the base features described above, the following applies to this Cloud Service:

- For each Instance of the Cloud Service, Client receives access to one (1) server dedicated to their activities which is not shared with any other Client.

#### **8.5.3 IBM Cloudant Virtual Machine Node**

In addition to the base features described above, the following applies to this Cloud Service offering:

- For each Instance of the Cloud Service, Client receives access to one (1) server that is configured as a virtual instance. This virtual instance will share hardware and network resources with other clients but the virtual instance will be exclusive to Client.

#### **8.5.4 IBM Cloudant Dedicated i2 Virtualized Node**

In addition to the base features described above, the following applies to this Cloud Service offering:

- For each Instance of the Cloud Service, Client receives access to one (1) server that is configured as a virtual instance. This virtual instance will share hardware and network resources with other clients but the virtual instance will be exclusive to Client.

#### **8.5.5 Security Description**

The IBM Cloudant Dedicated service is US-EU Safe Harbor certified.

#### **8.5.6 Data Protection**

If the Client makes Personal Data available to the Cloud Service, as between themselves and IBM, the Client is considered the sole controller of the Personal Data, and appoints IBM as a processor to process (as those terms are defined in EU Directive 95/46/EC) such Personal Data. IBM will only process such Personal Data to the extent required to make the Cloud Service available in accordance with this Service Description. IBM shall reasonably cooperate with the Client in its fulfillment of any legal requirement, including providing the Client with access to the Personal Data. The Client agrees that IBM may process content made available to the Cloud Service, including any Personal Data, across country borders to IBM in the United States. If IBM changes the way it processes or secures Personal Data and such change causes the Client to be noncompliant with applicable data protection laws, the Client may terminate the affected Cloud Service within 30 days of IBM's notification of the change.

#### **8.5.7 Personal Information and Regulated Data**

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

#### **8.5.8 Technical Support**

Technical support is offered with the Cloud Service and is not available as a separate offering. A Cloudant Dedicated Client is entitled to 24/7 email support with a response time of 1 hour.

More information about hours of availability, email addresses, online problem reporting systems, problem priorities and severities, and other technical support communication vehicles and processes are described in the IBM Cloudant section of the IBM Software as a Service (SaaS) Support Handbook.

### **8.6 IBM dashDB Enterprise**

The following DashDB services are covered by this Cloud Service Description:

- IBM dashDB Enterprise 64.1 Virtual Machine
- IBM dashDB Enterprise 256.4 Bare Metal
- IBM dashDB Enterprise MPP 4TB Bare Metal
- IBM dashDB Enterprise for Transactions 12.128.1400
- IBM dashDB Enterprise for Transactions 2.8.500
- IBM dashDB Enterprise MPP 32.244.1400 for AWS

IBM dashDB is a managed service which assists in the development of online analytical application processing. The service includes a database that stores user data in a structured format. This database can be accessed and modeled according to the user's requirements through the service's console. The service's console allows users to create tables, load data into the tables and query the data loaded by the user. Tools and environments to help develop, store and share analytical reports are included with the service. The service includes samples and documentation to help create the online analytical applications.

The dashDB Enterprise Transactional plans provide a dashDB database that is optimized for online transaction processing (OLTP). The service includes a database that stores user data in a structured

format. This database can be accessed and modeled according to the user's requirements through the service's console. The service's console allows users to create tables, load data into the tables and query the data loaded by the user.

IBM dashDB Enterprise MPP is a clustered configuration of the IBM dashDB service, comprising a set of independent database nodes managed together as a single system. Database operations are parallel processed across all the database nodes. IBM will configure the associated nodes into one cluster.

### 8.6.1 Security Description

#### a. **Personal Information and Regulated Data**

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

#### b. **User Management**

When the Cloud Service is provisioned, a single administrative user is created for Client. The dashDB console provides administrative users the ability to create additional users. Client is wholly responsible for managing the users defined via the console and the level of access assigned to them.

#### c. **Direct access to the Cloud Services data store**

Both administrative and regular users are able to directly access the Cloud Service data store using IBM DB2 client programs run outside the Cloud Service. Client is wholly responsible for ensuring that such access is secured according to Client's requirements. For example, Client may configure the clients such that SSL is used to protect network traffic.

#### d. **Table level access control**

The Cloud Service allows Client to manage the access rights associated with some database objects, such as tables. Client is wholly responsible for assigning, managing and reviewing these access rights.

#### e. **Audit Reports**

Monitoring reports are made available to Client through the dashDB console. Client is wholly responsible for accessing and interpreting the reports to determine if the activity they reflect is authorized.

#### f. **Encryption**

The Cloud Service databases are automatically encrypted. The encryption uses Advanced Encryption Standard (AES) in Cipher-Block Chaining (CBC) mode with a 256 bits key. Also, database backup images are automatically compressed and encrypted. Backup images are encrypted using AES in CBC mode with 256 bit keys.

### 8.6.2 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.

### **8.6.3 Beta Functions**

Some functions, features or components of the Cloud Service are early release or preview technology, and may be identified within the Cloud Service as "Beta" ("Beta Functions"). These Beta Functions may be used as part of Client's permitted use of the Cloud Service, subject to the limitations and conditions of this section. Use of Beta Functions is at Client's own risk and provided without obligation of support of any kind. Beta Functions are provided "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF TITLE, NON-INFRINGEMENT OR NON-INTERFERENCE AND ANY IMPLIED WARRANTIES AND CONDITIONS OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Beta Functions may not be made generally available by IBM as or in any product or offering. IBM may withdraw or terminate access to Beta Functions at any time, without notice. Client should take precautions to avoid any loss of data that might result when the Beta Functions can no longer be used. Any feedback or suggestions regarding the Beta Functions provided by Client to IBM may be freely used, copied, modified, and incorporated in the development, distribution, deployment and sale of IBM's products and services.

## **8.7 IBM Bluemix Data Connect**

IBM DataWorks is a cloud based, fully managed data preparation and movement service, that enables analysts, developers, data scientists and data engineers to put data to work. DataWorks allows technical and non-technical users to discover, cleanse, standardize, transform, and move data in support of application development and analytic use cases.

### **8.7.1 IBM Bluemix Data Connect Personal**

Per Instance, 100 activity runs or up to 1 TB of outbound data bandwidth per month. Outbound bandwidth is defined as the amount of data copied to a target.

### **8.7.2 IBM Bluemix Data Connect Professional**

Per Instance, 500 activity runs or up to 5 TB of outbound data bandwidth per month. Outbound bandwidth is defined as the amount of data copied to a target.

### **8.7.3 IBM DataWorks Enterprise**

Per Instance, 5000 activity runs or up to 50 TB of outbound data bandwidth per month. Outbound bandwidth is defined as the amount of data copied to a target.

## **8.8 IBM Streaming Analytics**

The IBM Streaming Analytics Standard Plan Cloud Service allows Clients to ingest, analyze, monitor, and correlate data as it arrives from real-time data sources. A single streams node instance is a virtual server with 1 CPU, 4 Cores, 12 GB RAM and 25 GB Disk.