

Release Notes

Community 1.5.1 release

The Release Notes provides the correlation between the releases and corresponding new features, enhancement and bug fixes.

Contents

Change List for 1.5.1 Release (22 June 2016)	3
Bug Fixes	3
Tasks.....	3
Change List for 1.5.0 Release (5 May 2015).....	3
Tasks.....	3
Bug Fixes	3
Enhancements	4
Features Available in 1.5.0 Release (5 May 2015)	4
Change List for 1.4.1 Release (16 February 2015)	4
Bug Fixes	4
Enhancements	5
Change List for 1.4.0 Release (11 December 2014).....	5
Tasks.....	5
Bug Fixes	5
Change List for 1.3.0 Release (5 November 2014).....	6
Enhancements	6
Bug Fixes	6
Change List for 1.2.1 Release (25 August 2014)	6
Enhancements	6
Bug Fixes	6
Change List for 1.2.0 Release (21 July 2014).....	6
Enhancements	6
Bug Fixes	7
Change List for 1.1.1 Release (10 July 2014).....	7
Bug Fixes	7
Change List for 1.1.0 Release (30 June 2014)	7
Enhancements	7
Bug Fixes	7
Features Available in 1.0.0 Release (5 June 2014).....	8

Change List for 1.5.1 Release (22 June 2016)

Bug Fixes

Bug ID	Details
JUM-10	No error is shown during validation if a wrong username is mentioned in the YAML.
JUM-100	Ability to build Jumbune docker image from the latest SNAPSHOT codebase.
JUM-102	Not able to deploy Jumbune when login with username hadoop.
JUM-117	No Password Attempts on Jumbune deployment (Distribution Creation).
JUM-118	Threw Stacktrace on console incase of authentication failure.
JUM-119	Password Attempts on Jumbune Agent Deployment.
JUM-120	Console logger for the namenode username by default.

Tasks

Task ID	Details
JUM-121	Upgrade Apache Commons Collections to v3.2.2
JUM-122	Modified For minor code refactor & exception handling.

Change List for 1.5.0 Release (5 May 2015)

Tasks

Bug ID	Details
JUM-58	Create and publish docker image of Jumbune
JUM-91	Create a yarn communication layer which can be reused by various components
JUM-93	Add a feature to Data Suite for enabling Data Profiling
JUM-94	Add a feature to Data Suite for providing Data Quality Timelines

Bug Fixes

Feature ID	Details
JUM-86	Simplify & Refactor YamlConfig to JobConfig, remove YamlLoader
JUM-87	Empty responses in two way calls resets barrier, empty responses has to avoided and to be handled on caller side

JUM-88	Job validation is unnecessarily done for historical jobs too
JUM-89	IdleResponseHandler is no longer required
JUM-90	For Yarn, Hist files given doesn't have reliable values for cpu and memory stats, revert from hist file to top file
JUM-95	Remove APIInvokeHints, provide a dynamic way by providing method names to calls
JUM-97	Data Validation needs to show all tuples and violated tuples count, null violation

Enhancements

Bug ID	Details
JUM-92	Provide an option to have declarative deployment through both command line options or configuration file
JUM-96	Add rumen processing logic for processing on MapR hadoop distribution
JUM-98	Deployment on Jumbune should support verbose mode to enable debug level logs

Features Available in 1.5.0 Release (5 May 2015)

The following features are available in the 1.5.0 Release:

- **Data Profiling:** Data Profiling Module helps user to classify data on the basis of different criteria. It enables user to put numeric constraints on different fields and then produces comparison statistics accordingly.
- **Data Quality Timeline:** Data Quality Timeline provides a user friendly way to analyze incremental data. It proves itself highly useful in the environments where high volumes data is ingested periodically. It assesses quality of data and powers user with ineffable insights.

Change List for 1.4.1 Release (16 February 2015)

Bug Fixes

Feature ID	Details
JUM-9	No validation done for the presence of the Job jar on the Jumbune machine.
JUM-77	Performance degradation of debugger jobs after instrumentation
JUM-79	Jumbune requests for password during deployment, even if password-less ssh is present

JUM-80	Need to refactor some classes related to validations to make them inheritable
JUM-81	Cluster Monitoring slots alignment issue on UI
JUM-82	Need to add support for Mapred user for cdh distributions of hadoop
JUM-83	Deployment doesn't support relative path to jar

Enhancements

Bug ID	Details
JUM-72	Data Validation does not work for recursive data path.

Change List for 1.4.0 Release (11 December 2014)

Tasks

Bug ID	Details
JUM-40	Using User policies
JUM-64	Support for Mapr 3.x
JUM-68	Data validation report should be persisted on the /tmp folder of the HDFS
JUM-69	Add serialization of distribution type so that its accessible to Jumbune container

Bug Fixes

Feature ID	Details
JUM-4	Jumbune execution does not pass the Validation phase for CDH
JUM-8	Data load partition and replica management are not working for Jumbune on CDH4
JUM-34	Unable to get resource statistics in job profiler after rumen processing
JUM-54	Exhaustive Deploy attempts in case of authentication errors during deployment
JUM-55	Multiple stats with name starting with "tag.*" appear and values are not shown
JUM-56	Flow debugger module fails on 0.23.x version of hadoop
JUM-62	Unused JIT jquery library requires to be removed.
JUM-63	Remove help framework.
JUM-65	Remove additional hadoop core jar from CLI
JUM-66	Export to excel feature is not working on YARN version
JUM-67	Cluster monitor looks for Resource Manager stats on non-RM nodes

JUM-70	Supported versions are not only limited to Apache, hence needs to be updated
JUM-71	Jumbune now doesn't require static Hadoop libraries

Change List for 1.3.0 Release (5 November 2014)

Enhancements

Bug ID	Details
JUM-24	Support for YARN based Hadoop distributions.
JUM-53	Jumbune job configuration in JSON format for web and shell based execution

Bug Fixes

Feature ID	Details
JUM-47	Refactor config and loader classes, make them reusable
JUM-41	JMX LRU connector cache is inadequately utilized during cluster monitoring
JUM-31	Jumbune CLI not functioning properly

Change List for 1.2.1 Release (25 August 2014)

Enhancements

Bug ID	Details
JUM-29	Various Jumbune dependency up-gradation.

Bug Fixes

Feature ID	Details
JUM-26	Cluster monitor fails after 3-4 hours of continuous execution and throws an 'out of memory' error
JUM-27	Data validation report is shown only for one type of violation
JUM-30	Password authentication not performed accurately during deployment
JUM-32	Jumbune is not started as hotspot server JVM

Change List for 1.2.0 Release (21 July 2014)

Enhancements

Feature ID	Details
------------	---------

JUM-12	Ability to Profile MapReduce jobs that are initialized outside Jumbune
------------------------	--

Bug Fixes

Bug ID	Details
JUM-25	Kill command not shutting down top command on the cluster nodes.

Change List for 1.1.1 Release (10 July 2014)

Bug Fixes

Bug ID	Details
JUM-14	The contents of the "jobjars" folder does not get deleted if cluster monitoring module is enabled along with the others.
JUM-15	Top file in the tmp directory get dumped with too much information and this could affect performance on long running jobs
JUM-23	Multiple initialisations of nio event group for jumbune remoter.

Change List for 1.1.0 Release (30 June 2014)

Enhancements

Feature Id	Details
JUM-11	Upgradation of Netty from 3.x to 4.x
JUM-13	Export feature for Debugger and data validation

Bug Fixes

Bug ID	Details
JUM-2	Cluster monitoring module fails if Master node has only NN and JT daemons running on it, without DN or TT.
JUM-16	Optimization:Executor service remains held even after completion of submitted Jumbune job.
JUM-17	Attributes from HTTP session are not cleaned up while home button is clicked.
JUM-18	Found instances of unreleased buffer readers
JUM-19	Optimization: Recurring buffer writer writer & flush can be optimized
JUM-20	Top dumps are not getting cleared
JUM-21	Cleanup of the job folder on the agent side is getting called for each slave.

Features Available in 1.0.0 Release (5 June 2014)

The following features are available in the 1.0.0 Release:

- **Flow Debugging:** This component verifies the flow of input records in user's map reduce implementation and provides In-depth analysis of Map Reduce flow thereby helping to identifying the bottle necks in the map reduce phase.
- **Cluster Monitoring:** This component provides on demand Hadoop JMX and Node resource statistics. This component provides network latency across Hadoop nodes, per file per node wise replica placement, rack aware nodes view as well as data load partition across the nodes.
- **JVM Profiling:** This component provides Map Reduce Phase wise stats which includes per job performance of JVM, data flow rate and resource usage. This component also provides per job heap sites and CPU cycles for Mapper and Reducer.
- **HDFS Data Validation:** This component validates inconsistencies in the HDFS data per record basis in the form of null, data type & regex checks.
- **NIO Based Transport Framework:** This framework is responsible for handling the interacting between the above components and Hadoop.