## Alibaba Cloud

E-MapReduce Product Introduction

Document Version: 20220627

(-) Alibaba Cloud

### Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

- You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
- 2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
- 3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
- 4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).
- 5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
- ${\bf 6. \ \ Please \ directly \ contact \ Alibaba \ Cloud \ for \ any \ errors \ of \ this \ document.}$

## **Document conventions**

Style	Description	Example
<u>Nanger</u>	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	Danger:  Resetting will result in the loss of user configuration data.
Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice:  If the weight is set to 0, the server no longer receives new requests.
? Note	A note indicates supplemental instructions, best practices, tips, and other content.	Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings> Network> Set network type.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click OK.
Courier font	Courier font is used for commands	Run the cd /d C:/window command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	bae log listinstanceid  Instance_ID
[] or [a b]	This format is used for an optional value, where only one item can be selected.	ipconfig [-all -t]
{} or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}

## **Table of Contents**

1.What is E-MapReduce?	06
2.Benefits	09
3.Architecture	10
4.Scenarios	12
5.Limits	15
6.Technical support	18
7.Release version	21
7.1. Overview	21
7.2. Release notes (EMR V5.X series)	34
7.2.1. Release notes for EMR V5.6.X	34
7.2.2. Release notes for EMR V5.5.X	37
7.2.3. Release notes for EMR V5.4.X	42
7.2.4. Release notes for EMR V5.3.X	47
7.2.5. Release notes for EMR V5.2.X	50
7.3. Release notes (EMR V4.X series)	54
7.3.1. Release notes for EMR V4.10.X	54
7.3.2. Release notes for EMR V4.9.X	59
7.3.3. Release notes for EMR V4.8.X	61
7.3.4. Release notes for EMR V4.6.X	66
7.3.5. Release notes for EMR V4.5.X	69
7.3.6. Release notes for EMR V4.4.X	72
7.3.7. Release notes for EMR V4.3.X	75
7.4. Release notes (EMR V3.X series)	78
7.4.1. Release notes for EMR V3.40.X	78
7.4.2. Release notes for EMR V3.39.X	81
7.4.3. Release notes for EMR V3.38.X	87

7.4.4. Release notes for EMR V3.37.X	95
7.4.5. Release notes for EMR V3.36.X	99
7.4.6. Release notes for EMR V3.35.X	02
7.4.7. Release notes for EMR V3.34.X	03
7.4.8. Release notes for EMR V3.33.X	06
7.4.9. Release notes of EMR V3.32.X1	08
7.4.10. Release notes of EMR V3.30.X	09
7.4.11. Release notes of EMR V3.29.X1	11
7.4.12. Release notes of EMR V3.28.X1	12
7.4.13. Release notes of EMR V3.27.X1	14
7.4.14. Release notes of EMR V3.26.X1	16
7.4.15. Release notes of EMR V3.25.X1	17
7.4.16. Release notes of EMR V3.24.X 1	18
7.4.17. Release notes of EMR V3.23.X1	20
7.4.18. Release notes of EMR V3.22.X1	22
7.4.19. Release notes of versions earlier than E-MapReduce V 1	26

## 1.What is E-MapReduce?

Alibaba Cloud E-MapReduce (EMR) is a big data processing solution provided by Alibaba Cloud.

### Introduction

EMR is built on top of Alibaba Cloud Elastic Compute Service (ECS) and developed based on open source Apache Hadoop and Apache Spark. It allows you to conveniently use peripheral systems in the Hadoop and Spark ecosystems to analyze and process data. EMR can also read data from or write data to other Alibaba Cloud storage systems and database systems, such as Object Storage Service (OSS) and ApsaraDB RDS.

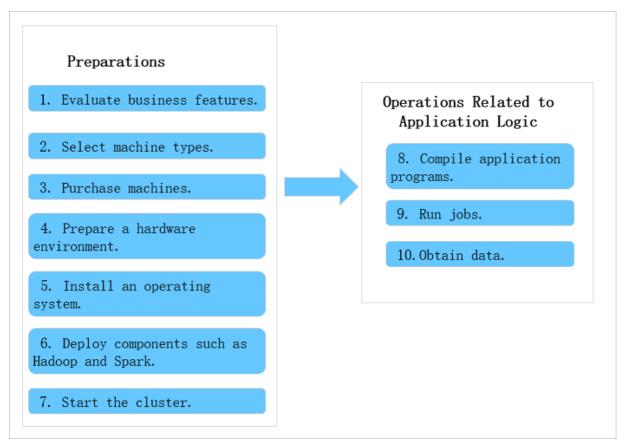
Smart Data is a storage service for the EMR Jindo engine. Smart Data provides centralized storage, optimized caching, and accelerated cache computing for EMR computing engines. It extends the storage features.



- For more information about Apache Hadoop, visit the Apache Hadoop official website.
- For more information about Apache Spark, visit the Apache Spark official website.
- For more information about Apache Hive, visit the Apache Hive official website.
- For more information about Apache HBase, visit the Apache HBase official website.
- For more information about Smart Data, see Smart Data.

### Advantages over open source big data ecosystems

If you use an open source distributed processing system, such as Hadoop or Spark, to process data without using EMR, you must perform all the steps in the following figure.



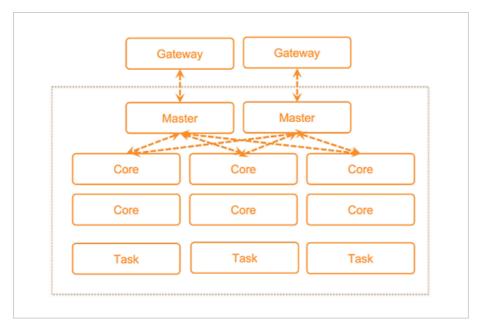
In this procedure, only the last three steps are related to your application logic. The first seven steps are all preparations, which are complex and time-consuming. EMR integrates all the required cluster management tools to provide the following features: host selection, environment deployment, cluster building, cluster configuration, cluster running, job configuration, job running, cluster management, and performance monitoring. This frees you from all the tedious procurement, preparation, and O&M work required to build clusters. You need only to focus on the processing logic of your applications.

EMR also offers different combinations of cluster services to meet your business requirements. For example, to perform daily data measurement and batch computing, you need only to run the Hadoop service for EMR. If you also want to perform stream computing and real-time computing, you can add the Spark service.

### Composition of EMR clusters

Clusters are the core user-oriented component of EMR. An EMR cluster is a Hadoop, Flink, Druid, or ZooKeeper cluster that is deployed on one or more ECS instances. For example, a Hadoop cluster consists of some daemon processes, such as NameNode, DataNode, ResourceManager, and NodeManager. These daemon processes run on the ECS instances of the cluster. Each ECS instance corresponds to a node.

The following figure shows a Hadoop cluster that consists of master nodes, core nodes, and task nodes and is associated with a gateway cluster.



- Master nodes: run the following services and processes: NameNode of HDFS, JournalNode of HDFS, ZooKeeper, ResourceManager of YARN, and HMaster of HBase. You can determine whether to create a high-availability (HA) cluster based on the use scenario of the cluster. We recommend that you create a non-HA cluster in a test environment and create an HA cluster in a production environment. In an HA cluster, you can deploy two or three master nodes. If you deploy two master nodes, JournalNode of HDFS and ZooKeeper are deployed on the emr-worker-1 node. We recommend that you deploy three master nodes when you create an HA cluster in a production environment.
- Core nodes: run the DataNode process of HDFS and the NodeManager process of YARN to support data storage and computing. Core nodes do not support auto scaling.
- Task nodes: run the NodeManager process of YARN to support computing. You can flexibly add or remove task nodes based on your business requirements.
- Gateway cluster: is configured with Hadoop client files. You can use a gateway cluster to submit jobs to the Hadoop cluster associated with this gateway cluster. This way, you do not need to log on to the Hadoop cluster, which avoids security issues and environment isolation issues. You must create a Hadoop cluster first, and then create a gateway cluster and associate it with the Hadoop cluster.

## 2.Benefits

Alibaba Cloud E-MapReduce (EMR) provides you with convenient and controllable open source big data services at the enterprise level. It allows you to easily deploy open source big data services, such as Hadoop, Spark, Flink, Kafka, and HBase.

### Stable and reliable open source components

- Open source components are used. Each version of EMR provides the latest versions of open source components. For more information about the mappings between EMR versions and open source component versions, see Overview.
- EMR is fully adapted to open source components and has eliminated the version compatibility issues among various open source components.
- EMR provides an enhanced deployment environment in Alibaba Cloud for open source components. This ensures a much higher performance than that in the open source community.

### Cost effectiveness

- Compared with traditional HDFS-based clusters that use fixed configurations, EMR clusters use auto scaling and tiered storage mechanisms, which help reduce costs by more than 50%.
- You can create preemptible instances. Compared with the pay-as-you-go billing method, preemptible instances help reduce fees by 50% to 80%. For more information, see Overview.

### Ease of use

- You can create or scale out a cluster within minutes. You do not need to manually deploy or start services.
- EMR provides a comprehensive monitoring and alerting system, which covers the hardware and Hadoop service of clusters. You can configure alert templates.

### Scalability

- Compute-storage separation: Computing and storage are decoupled to support the elastic use of resources.
- Custom cluster environment: You can use bootstrap actions and cluster scripts to flexibly configure a cluster environment and deploy third-party optimization or cluster management tools in EMR. For more information, see Manage bootstrap actions and Cluster scripts.
- Self-managed maintenance: You can log on to the master node of a cluster, view the logs and deployment environment of the cluster, and optimize the configurations. For more information, see Common file paths.
- Auto scaling: EMR can automatically scale in or out clusters based on your business requirements.

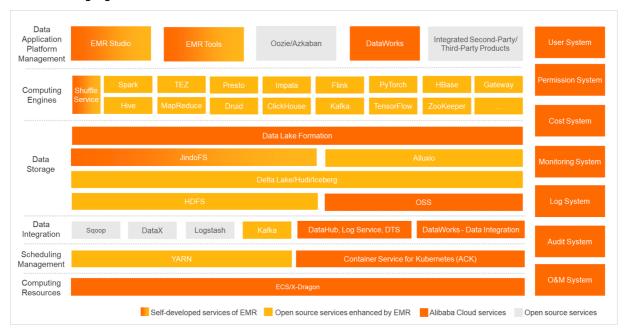
### Deep integration

- You can deploy EMR clusters based on Alibaba Cloud Elastic Compute Service (ECS) and Container Service for Kubernetes (ACK). Various ECS instance types are supported. You can determine the instance types to use based on your business requirements. For more information, see ECS instances.
- EMR is integrated into DataWorks. You can use EMR as a job computing and data storage engine in DataWorks.
- Data Lake Formation (DLF) is integrated into EMR. In data lake scenarios, EMR allows you to manage metadata for multiple engines in a centralized manner.

## 3.Architecture

This topic describes the architecture of E-MapReduce (EMR).

The following figure shows the architecture of EMR.



EMR consists of four types of services:

• Open source services

Apache big data services, such as Hadoop, Hive, and HBase, are integrated into EMR. The versions of the open source services are updated with EMR versions. For more information, see release notes in Overview.

Notice You are not allowed to update the version of a service in an existing EMR cluster.

• Open source services enhanced by EMR

EMR enhances the performance and features of some open source services. Examples:

- Spark Streaming SQL is added to Spark. This significantly improves the performance of Spark. For more information, see Common keywords.
- The Z-ordering and Data Skipping features are added to Delta Lake. For more information, see Overview.
- Self-developed services of EMR

EMR provides the following self-developed services, which ensure that open source components and services can better run on the Alibaba Cloud infrastructure:

- Shuffle Service is an extended component of EMR. It is used to optimize the shuffle operations of computing engines. For more information, see ESS overview.
- Smart Data optimizes storage, caching, and computing for various EMR computing engines in a centralized manner and extends storage features. For more information, see Smart Data.
- Alibaba Cloud services

EMR connects to both open source big data ecosystems and the Alibaba Cloud ecosystem. You can deploy EMR clusters on Alibaba Cloud Elastic Compute Service (ECS) instances or Container Service for Kubernetes (ACK) clusters and store data in Alibaba Cloud Object Storage Service (OSS). You can learn and use Machine Learning Platform for AI (PAI) in an EMR Data Science cluster. EMR is integrated into DataWorks, and you can use EMR as a job computing engine or data storage engine in DataWorks.

## 4.Scenarios

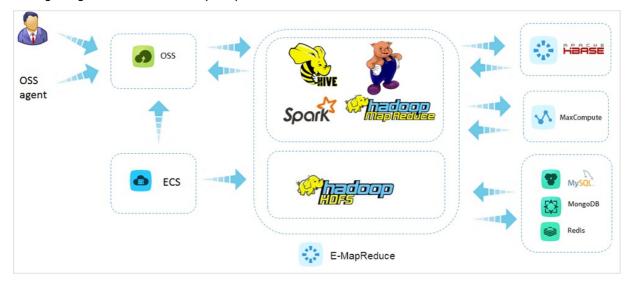
E-MapReduce (EMR) clusters are suitable for all the scenarios supported by the Hadoop ecosystem and Spark.

EMR is a cluster service based on Hadoop and Spark. You can use the Alibaba Cloud Elastic Compute Service (ECS) instances on which EMR clusters are deployed as your dedicated physical machines. The following examples show the classic scenarios of EMR:

- Batch processing of data
- Ad hoc queries for data analysis
- Online analysis of large amounts of data
- Processing of streaming data

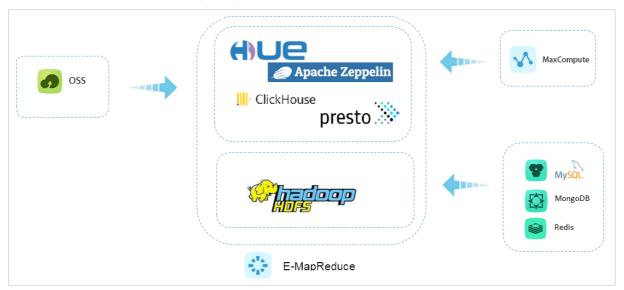
### Batch processing of data

You can synchronize a large number of logs to the data nodes of EMR. Then, you can use tools such as Hue and a mainstream computing framework such as Hive, Spark, or Presto to get a quick insight into the data. You can also use tools such as Sqoop to load data from ApsaraDB RDS or other storage engines to EMR. Then, you can analyze the data and synchronize the results to ApsaraDB RDS or other storage engines. This feature helps implement data visualization.



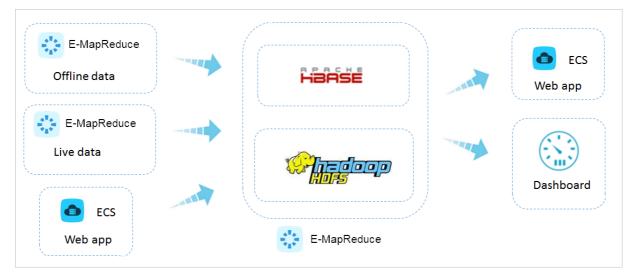
Ad hoc queries for data analysis

EMR allows you to directly import massive data to or use an external table to import massive data to online analytical processing (OLAP) analysis engines, such as ClickHouse, Presto, and Impala, for efficient, real-time, and flexible data analysis. This feature is suitable for business scenarios such as user persona analysis, audience selection, BI reporting, and business analysis.



### Online analysis of large amounts of data

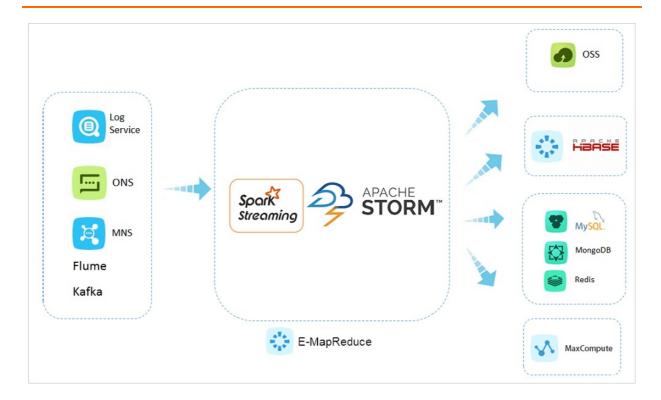
EMR analyzes petabytes of structured, semi-structured, or unstructured data generated by web apps and mobile apps. This allows web apps or data visualization services to visualize data in real time based on the analysis results obtained from EMR.



### Processing of streaming data

EMR allows you to use and process real-time streaming data from services such as Log Service, Message Queue, Message Service (MNS), and Apache Kafka based on Spark Streaming and Storm.

EMR analyzes streaming data in fault-tolerant mode and writes analysis results to Object Storage Service (OSS) or HDFS.



## 5.Limits

When you use E-MapReduce (EMR), cluster instability may occur or clusters may become unavailable due to unexpected operations. Take note of the information in this topic to avoid these issues. This topic describes the limits of EMR.

Notice When you use EMR, you must perform all operations in the EMR console. We recommend that you do not perform operations in the Elastic Compute Service (ECS) console. This may cause cluster instability or abnormalities. Take note of the information in this topic. In the case that you perform the operations that are prohibited, you will bear the consequences and responsibilities.

### High-risk operations (prohibited)

Operation	Possible result	Suggestion
Delete or modify the hosts file that is stored in the etc/directory.	You cannot find the services that run on the nodes of your cluster, which causes service exceptions.	Add information to the hosts file.
Modify parameters in component configuration files in the ECS console.	After specific services are restarted, the settings of the parameters that are modified are overwritten.	Modify the parameters in the EMR console.
Redeploy ECS instances in the ECS console.	The EMR service is affected.	submit a ticket.
Attach disks to the nodes of your EMR cluster in the ECS console.	The disks are unavailable because EMR cannot recognize and initialize the disks.	submit a ticket to replace disks if
Detach disks from the nodes of your EMR cluster in the ECS console.	This may cause data loss because EMR is unaware of the disk detaching operation.	necessary.
Remove core nodes in the ECS console.	This causes data loss, and execution failures of jobs on the removed nodes.	None

Operation	Possible result	Suggestion
Remove master nodes in the ECS console.	<ul> <li>For a high availability (HA) cluster, if you remove master nodes, the switchover of HDFS NameNode HA, YARN ResourceManager, or HBase HMaster fails. In this case, you must purchase a new EMR cluster to migrate data or tasks.</li> <li>For a non-HA cluster, if you remove the master node, the cluster becomes unavailable, and you cannot migrate data or tasks.</li> </ul>	None
Remove task nodes in the ECS console.	The jobs that you run on the removed nodes fail.	None
Stop the MySQL service of the master node. (Type is set to Built-in MySQL when you create an EMR cluster.)	The MySQL service deployed on the emr-header-1 node is associated with Hive MetaStore, Oozie, and Ranger. If you stop the MySQL service, the associated components cannot access the specific database.	None
Change the password of the root user that is used to access the MySQL service deployed on the emr-header-1 node. (Type is set to Built-in MySQL when you create an EMR cluster.)	The associated component such as Hue or Ranger fails.	None
Modify the security group to which ECS instances belong when an EMR cluster is running.	<ul> <li>The network connection between nodes is abnormal.</li> <li>Components become unavailable.</li> </ul>	None

Operation	Possible result	Suggestion	
		Change the billing method in the EMR console. For more information, see Switch from pay-as-you-go to subscription.	
Change the billing method in the ECS console.	After you change the billing method, you can no longer change the new billing method back to the original one.	Notice You cannot change the billing method from subscription to pay-as-you-go in the EMR console. If you want to change the billing method from subscription to pay-as-you-go, .	

### FAQ

Problem description	Solution
Insufficient disk capacity	Increase the capacity of a single disk or add core nodes in the EMR console. EMR clusters do not support the addition of disks.
Excess disk capacity	Purchase a new cluster and release the original one. For more information, see Create a cluster. EMR clusters do not support scaledown of disk capacity.
Insufficient computing capabilities	Add task nodes in the EMR console. For more information, see Scale out a cluster.
Excess computing capabilities	<ul> <li>Resolve this issue based on the billing method of your cluster.</li> <li>For a pay-as-you-go cluster, remove one or more task nodes from the cluster in the EMR console.</li> <li>For a subscription cluster, stop the NodeManager of YARN on a specific task node, change the billing method of the ECS instance that serves as the task node to pay-as-you-go in the ECS console, and then release the instance.</li> </ul>
Outdated component versions	Purchase a cluster of a later version. You are not allowed to update a specific component of an existing cluster. For more information, see Create a cluster.
Conversion from a non-HA cluster to an HA cluster	Non-HA clusters cannot be converted to HA clusters. We recommend that you purchase an HA cluster.
Deployment of third-party software or services on EMR	We recommend that you use bootstrap actions to install third-party software or third-party services when you create a cluster.  If you manually install third-party software or third-party services after you create a cluster, you must reinstall the software or services when you add nodes.

## 6.Technical support

Alibaba Cloud E-MapReduce (EMR) is a solution provided by Alibaba Cloud. EMR is built on top of Alibaba Cloud Elastic Compute Service (ECS) or Container Service for Kubernetes (ACK) in semi-hosted mode. You are granted all management and operation permissions on EMR clusters that you created. This topic describes the EMR technical support for you from the following aspects:

- Support scope
- Seek help from Alibaba Cloud technical support
- SLA

### Support scope

- Services supported
  - The EMR technical support team provides support for the use of the EMR console, provides related consultation services, and handles stability issues and product defects related to the console. The EMR console excludes components that are deployed in EMR clusters. You can perform the following operations in the EMR console:
    - Purchase an EMR cluster.
    - Create an EMR cluster.
    - Scale an EMR cluster.
    - Add services to an EMR cluster.
    - Configure auto scaling rules.
    - Use the monitoring and alerting feature.
    - Develop data.
  - The EMR technical support team provides support for the use of self-developed components in EMR, provides related consultation services, and handles stability issues and product defects related to the self-developed components. EMR provides the following self-developed components:
    - Smart Dat a
    - Bigboot
    - ESS
    - DLF-Auth
  - The EMR technical support team assists in troubleshooting issues that occur on open source components in EMR and provides solutions in an iterative schedule based on the product plan. EMR supports the following open source components:
    - Hadoop components, including Hadoop Distributed File System (HDFS), YARN, and MapReduce
    - Components for data integration, such as Flume and Sqoop
    - Components for data storage, such as HDFS, HBase, Kudu, Iceberg, Delta Lake, Hudi, and Kafka
    - Compute engine components, such as Spark, Hive, Tez, Druid, Flink, and Phoenix
    - Online analytical processing (OLAP) components, such as Presto, ClickHouse, and Impala
    - Other components, such as ZooKeeper, Knox, Ganglia, OpenLDAP, Hue, Zeppelin, Livy, Superset, Ranger, Alluxio, and Oozie

- The EMR technical support team assists in troubleshooting issues that occur on other Alibaba
  Cloud services on which EMR depends. The service teams of these Alibaba Cloud services are
  responsible for the handling of service defects and provide solutions to meet service requirements.
  EMR depends on the following Alibaba Cloud services:
  - ECS or ACK
  - Object Storage Service (OSS)
  - CloudMonitor
- Services not supported
  - The EMR technical support team provides reliable O&M features such as monitoring and alerting
    and parameter configuration. However, the EMR technical support team does not perform O&M
    operations for cluster or component management. The O&M operations include cluster capacity
    planning, component performance optimization, and specific O&M operations on a cluster, such as
    service restart and parameter modification.
  - The EMR technical support team does not provide technical consultation on the development of business applications on top of EMR components or troubleshoot related issues. The development covers the following processes: business implementation, code development, job performance tuning, and business migration.
  - The EMR technical support team does not troubleshoot issues in a single job if no obvious exceptions in cluster components occur or no service defects exist.
  - The EMR technical support team does not provide support for non-standard capabilities of EMR, such as upgrade of the cluster version or component configurations, scale-in of core nodes, and reduction of disk capacity.
  - The EMR technical support team does not troubleshoot the issues that occur on third-party components.
  - The EMR technical support team does not troubleshoot the issue of cluster instability or cluster unavailability due to unexpected operations. For more information about unexpected operations, see Limits.

### Seek help from Alibaba Cloud technical support

If you have questions, we recommend that you submit a ticket on the Alibaba Cloud official website to seek help from Alibaba Cloud technical support

- 1. Log on to the ticket system.
- 2. Click the vicon below More Product And Service Category.
- 3. Click E-MapReduce in the More section.
- 4. On the **New Ticket** page, configure the parameters.
- 5. Click Submit.

You can scan the following QR code to join the DingTalk group of EMR. In this group, you can discuss issues about EMR with other users.



**Note** This DingTalk group is used only for communication between users. If you have any questions, submit a ticket.

### SLA

For more information about SLA, see Service Level Comparison.

# 7.Release version 7.1. Overview

This topic describes the following information about the released versions of E-MapReduce (EMR): version format, version upgrade records, and release notes.

### Version format

EMR versions follow the Va.b.c format, where:

- a indicates that major changes are made in a new version.
- b indicates that changes are made on some components in a new version.
- c indicates that several bugs are fixed in a new version and the version is forward compatible.

The following example describes the version number changes in a version upgrade:

- Upgrade from version 1.0.0 to 2.0.0: The value of a changes, which indicates that major changes are made in the new version. You must conduct a test to verify that all existing jobs can run properly after the upgrade.
- Upgrade from version 1.0.0 to 1.1.0: The value of b changes, which indicates that minor changes are made on some components in the new version. Most features remain unchanged after the upgrade. You must conduct a test to verify that all existing jobs can run properly after the upgrade.
- Upgrade from version 1.0.0 to 1.0.1: The value of c changes, which indicates that the new version is backward compatible with the previous version.

The following content describes the bundled software, cluster creation, and cluster upgrade of each EMR version:

- Bundled software: The software and software version are fixed in each EMR version. We recommend that you do not change the software version.
- Cluster creation: After you select an EMR version and create a cluster of this version, the cluster version is not automatically upgraded. Subsequent upgrades to the image used by this version do not affect the existing clusters. The new version of the image is used to create new clusters.
- Cluster upgrade: When you upgrade the version of your cluster, for example, from version 1.0.X to 1.1.X, you must conduct a test to verify that the existing jobs in the cluster can run properly in the new software environment.

### Version upgrade records (EMR V5.X series)

- Release notes for EMR V5.6.X
- Release notes for EMR V5.5.X
- Release notes for EMR V5.4.X
- Release notes for EMR V5.3.X
- Release notes for EMR V5.2.X

### Version upgrade records (EMR V4.X series)

- Release notes for EMR V4.10.X
- Release notes for EMR V4.9.X
- Release notes for EMR V4.8.X

- Release notes for EMR V4.6.X
- Release notes for EMR V4.5.X
- Release notes for EMR V4.4.X
- Release notes for EMR V4.3.X

### Version upgrade records (EMR V3.X series)

- Release notes for EMR V3.40.X
- Release notes for EMR V3.39.X
- Release notes for EMR V3.38.X
- Release notes for EMR V3.37.X
- Release notes for EMR V3.36.X
- Release notes for EMR V3.35.X
- Release notes for EMR V3.34.X
- Release notes for EMR V3.33.X
- Release notes of EMR V3.32.X
- Release notes of EMR V3.30.X
- Release notes of EMR V3.29.X
- Release notes of EMR V3.28.X
- Release notes of EMR V3.27.X
- Release notes of EMR V3.26.X
- Release notes of EMR V3.25.X
- Release notes of EMR V3.24.X
- Release notes of EMR V3.23.XRelease notes of EMR V3.22.X
- Release notes of versions earlier than E-MapReduce V3.22.X

### EMR V4.3.0 to EMR V4.6.0

Version	EMR V4.6.0	EMR V4.5.0	EMR V4.4.1	EMR V4.3.0
Release date	2021.1	2020.12	2020.8	2020.5
Hadoop	3.2.1	3.1.3	3.1.3	3.1.3
HDFS	3.2.1	3.1.3	3.1.3	3.1.3
YARN	3.2.1	3.1.3	3.1.3	3.1.3
Hive	3.1.2	3.1.2	3.1.2	3.1.1
Spark	2.4.7	2.4.5	2.4.5	2.4.5
Knox	1.1.0	1.1.0	1.1.0	1.1.0
Tez	0.9.2	0.9.2	0.9.2	0.9.2

Version	EMR V4.6.0	EMR V4.5.0	EMR V4.4.1	EMR V4.3.0
Ganglia	3.7.2	3.7.2	3.7.2	3.7.2
Sqoop	1.4.7	1.4.7	1.4.7	1.4.7
SmartData	3.2.0	3.1.0	2.7.301	2.6.7
Bigboot	3.2.0	3.1.0	2.7.301	2.6.7
ESS	1.0.0	1.0.0	N/A	N/A
Hudi	0.6.0	0.6.0	N/A	N/A
OpenLDAP	2.4.44	2.4.44	2.4.44	2.4.44
Hue	4.4.0	4.4.0	4.4.0	4.4.0
HBase	2.1.9	2.1.9	2.1.9	2.1.9
ZooKeeper	3.5.6	3.5.6	3.5.6	3.5.6
Presto	338	338	331	331
Impala	3.4.0	3.4.0	3.4.0	3.2.0
Zeppelin	0.9.0	0.8.2	0.8.1	0.8.1
Pig	N/A	0.14.0	0.14.0	0.14.0
Flume	1.9.0	1.9.0	1.9.0	1.9.0
Livy	0.6.0	0.6.0	0.6.0	0.6.0
Superset	0.36.0	0.36.0	0.36.0	0.28.1
Ranger	2.1.0	2.1.0	2.0.0	2.0.0
Flink	1.10-vvr-1.0.2	1.10-vvr-1.0.2	1.10-vvr-1.0.2	1.10-vvr-1.0.2
Storm	1.2.2	1.2.2	1.2.2	1.2.2
Kudu	1.11.1	1.11.1	1.11.1	1.11.1
Oozie	5.1.0	5.1.0	5.1.0	5.1.0
Kafka	2.4.1	2.4.1	2.4.1	2.4.1
Kafka-Manager	1.3.3.16	1.3.3.16	1.3.3.16	1.3.3.16

### EMR V3.28.2 to EMR V3.33.0

Version	EMR V3.33.0	EMR V3.32.0	EMR V3.30.0	EMR V3.29.0	EMR V3.28.2
Release date	2021.1	2020.11	2020.6	2020.6	2020.5
Hadoop	2.8.5	2.8.5	2.8.5	2.8.5	2.8.5
Knox	1.1.0	1.1.0	1.1.0	1.1.0	1.1.0
Spark	2.4.7	2.4.5	2.4.5	2.4.5	2.4.5
Hive	2.3.7	2.3.5	2.3.5	2.3.5	2.3.5
Tez	0.9.2	0.9.2	0.9.2	0.9.2	0.9.2
Pig	0.14.0	0.14.0	0.14.0	0.14.0	0.14.0
Sqoop	1.4.7	1.4.7	1.4.7	1.4.7	1.4.7
YARN	2.8.5	2.8.5	2.8.5	2.8.5	2.8.5
HDFS	2.8.5	2.8.5	2.8.5	2.8.5	2.8.5
Flink	1.11-vvr-2.0.1	1.11-vvr-2.0.1	1.11-vvr-2.0.1	1.10-vvr-1.0.4	1.10-vvr-1.0.4
Druid	0.20.0	0.18.1	0.18.0	0.18.0	0.17.1
HBase	1.4.9	1.4.9	1.4.9	1.4.9	1.4.9
Phoenix	4.14.1	4.14.1	4.14.1	4.14.1	4.14.1
ZooKeeper	3.5.6	3.5.6	3.5.6	3.5.6	3.5.6
Livy	0.6.0	0.6.0	0.6.0	0.6.0	0.6.0
Presto	338	338	338	338	331
Storm	1.2.2	1.2.2	1.2.2	1.2.2	1.2.2
Impala	2.12.2	2.12.2	2.12.2	2.12.2	2.12.2
Flume	1.9.0	1.9.0	1.9.0	1.9.0	1.9.0
Hue	4.4.0	4.4.0	4.4.0	4.4.0	4.4.0
Oozie	5.1.0	5.1.0	5.1.0	5.1.0	5.1.0
Zeppelin	0.9.0	0.8.1	0.8.1	0.8.1	0.8.1
Ranger	1.2.0	1.2.0	1.2.0	1.2.0	1.2.0
Ganglia	3.7.2	3.7.2	3.7.2	3.7.2	3.7.2
Smart Dat a	3.2.0	3.1.0	3.0.0	2.7.301	2.7.2

Version	EMR V3.33.0	EMR V3.32.0	EMR V3.30.0	EMR V3.29.0	EMR V3.28.2
Kafka	1.1.1	1.1.1	1.1.1	1.1.1	1.1.1
Kafka- Manager	1.3.3.16	1.3.3.16	1.3.3.16	1.3.3.16	1.3.3.16
Superset	0.36.0	0.36.0	0.36.0	0.35.2	0.35.2
Bigboot	3.2.0	3.1.0	3.0.0	2.7.301	2.7.2
OpenLDAP	2.4.44	2.4.44	2.4.44	2.4.44	2.4.44
Kudu	1.10.0	1.10.0	1.10.0	1.10.0	1.10.0
Flink-Vvp	1.11-2.2.2	1.11-2.2.2	1.11-2.2.2	1.11-2.2.2	1.11-2.2.2
PAI-Alink	1.1.0	1.1.0	1.1.0	1.1.0	1.1.0
Alluxio	2.4.0	2.4.0	N/A	N/A	N/A
Hudi	0.6.0	0.6.0	N/A	N/A	N/A

### EMR V3.25.1 to EMR V3.28.1

Version	EMR V3.28.1	EMR V3.28.0	EMR V3.27.2	EMR V3.26.3	EMR V3.25.1
Release date	2020.6	2020.6	2020.5	2020.4	2020.1
Hadoop	2.8.5	2.8.5	2.8.5	2.8.5	2.8.5
Knox	1.1.0	1.1.0	1.1.0	1.1.0	1.1.0
Spark	2.4.5	2.4.5	2.4.3	2.4.3	2.4.3
Hive	2.3.5	2.3.5	2.3.5	2.3.5	2.3.5
Tez	0.9.2	0.9.2	0.9.2	0.9.2	0.9.2
Pig	0.14.0	0.14.0	0.14.0	0.14.0	0.14.0
Sqoop	1.4.7	1.4.7	1.4.7	1.4.7	1.4.7
YARN	2.8.5	2.8.5	2.8.5	2.8.5	2.8.5
HDFS	2.8.5	2.8.5	2.8.5	2.8.5	2.8.5
Flink	1.10-vvr-1.0.2	1.10-vvr-1.0.2	1.9.1	1.9.1	1.9.1
Druid	0.18.0	0.18.0	0.17.1	0.16.0	0.16.0
HBase	1.4.9	1.4.9	1.4.9	1.4.9	1.4.9

Version	EMR V3.28.1	EMR V3.28.0	EMR V3.27.2	EMR V3.26.3	EMR V3.25.1
Phoenix	4.14.1	4.14.1	4.14.1	4.14.1	4.14.1
ZooKeeper	3.5.6	3.5.6	3.5.6	3.5.6	3.5.6
Livy	0.6.0	0.6.0	0.6.0	0.6.0	0.6.0
Presto	331	331	331	310	310
Storm	1.2.2	1.2.2	1.2.2	1.2.2	1.2.2
Impala	2.12.2	2.12.2	2.12.2	2.12.2	2.12.2
Flume	1.9.0	1.9.0	1.9.0	1.9.0	1.9.0
Hue	4.4.0	4.4.0	4.4.0	4.4.0	4.4.0
Oozie	5.1.0	5.1.0	5.1.0	5.1.0	5.1.0
Zeppelin	0.8.1	0.8.1	0.8.1	0.8.1	0.8.1
Ranger	1.2.0	1.2.0	1.2.0	1.2.0	1.2.0
Ganglia	3.7.2	3.7.2	3.7.2	3.7.2	3.7.2
SmartData	2.7.1	2.7.0	2.6.5	2.6.3	2.2.3
Kafka	1.1.1	1.1.1	1.1.1	1.1.1	1.1.1
Kafka- Manager	1.3.3.16	1.3.3.16	1.3.3.16	1.3.3.16	1.3.3.16
Superset	0.35.2	0.35.2	0.35.2	0.28.1	0.28.1
Bigboot	2.7.1	2.7.0	2.6.5	2.6.3	2.2.3
OpenLDAP	2.4.44	2.4.44	2.4.44	2.4.44	2.4.44
Kudu	1.10.0	1.10.0	1.10.0	1.10.0	1.10.0
Analytics Zoo	N/A	N/A	N/A	0.5.0	0.5.0
Tensorflow on Spark	N/A	N/A	N/A	1.0.1	1.0.1

### EMR V3.22.1 to EMR V3.25.0

Version	EMR V3.25.0	EMR V3.24.3	EMR V3.24.0	EMR V3.23.0	EMR V3.22.1
Release date	2020.1	2019.11	2019.11	2019.9	2019.7
Hadoop	2.8.5	2.8.5	2.8.5	2.8.5	2.8.5

Version	EMR V3.25.0	EMR V3.24.3	EMR V3.24.0	EMR V3.23.0	EMR V3.22.1
Knox	1.1.0	1.1.0	1.1.0	1.1.0	1.1.0
Spark	2.4.3	2.4.3	2.3.5	2.4.3	2.4.3
Hive	2.3.5	2.3.5	2.3.5	2.3.5	2.3.5
Tez	0.9.2	0.9.1	0.9.1	0.9.1	0.9.1
Pig	0.14.0	0.14.0	0.14.0	0.14.0	0.14.0
Sqoop	1.4.7	1.4.7	1.4.7	1.4.7	1.4.7
YARN	2.8.5	2.8.5	2.8.5	2.8.5	2.8.5
HDFS	2.8.5	2.8.5	2.8.5	2.8.5	2.8.5
Flink	1.9.1	1.9.1	1.9.1	1.8.2	1.7.2
Druid	0.16.0	0.16.0	0.16.0	0.15.1	0.14.2
HBase	1.4.9	1.4.9	1.4.9	1.4.9	1.4.9
Phoenix	4.14.1	4.14.1	4.14.1	4.14.1	4.14.1
ZooKeeper	3.5.6	3.5.5	3.5.5	3.5.5	3.5.5
Livy	0.6.0	0.6.0	0.6.0	0.6.0	0.6.0
Presto	310	0.228	0.228	0.221	0.221
Storm	1.2.2	1.2.2	1.2.2	1.2.2	1.2.2
Impala	2.12.2	2.12.2	2.12.2	2.12.2	2.12.2
Flume	1.9.0	1.9.0	1.9.0	1.9.0	1.9.0
Hue	4.4.0	4.4.0	4.4.0	4.4.0	4.4.0
Oozie	5.1.1	5.1.0	5.1.0	5.1.0	5.1.0
Zeppelin	0.8.1	0.8.1	0.8.1	0.8.1	0.8.1
Ranger	1.2.0	1.2.0	1.2.0	1.2.0	1.2.0
Ganglia	3.7.2	3.7.2	3.7.2	3.7.2	3.7.2
Tensorflow on Spark	1.0.1	1.0.1	1.0.1	1.0.0	1.0.0
Analytics Zoo	0.5.0	0.5.0	N/A	0.5.0	0.5.0
Kafka	1.1.1	1.1.1	1.1.1	1.1.1	1.1.1

Version	EMR V3.25.0	EMR V3.24.3	EMR V3.24.0	EMR V3.23.0	EMR V3.22.1
Kafka- Manager	1.3.3.16	1.3.3.16	1.3.3.16	1.3.3.16	1.3.3.16
Superset	0.28.1	0.28.1	0.28.1	0.28.1	0.28.1
Bigboot	2.2.3	2.2.2	2.2.1	2.0.2	2.0.1
Smart Dat a	2.2.3	2.2.2	2.2.1	2.0.2	2.0.1
OpenLDAP	2.4.44	2.4.44	2.4.44	2.4.44	2.4.44
Kudu	1.10.0	1.10.0	1.10.0	1.10.0	1.10.0

## Version upgrade records (earlier than EMR V3.22.X)

### • EMR V3.19.0 to EMR V3.21.0

Version	EMR V3.19.0	EMR V3.19.1	EMR V3.20.0	EMR V3.21.0
Release date	2019.3	2019.4	2019.5	2019.6
Hadoop	2.8.5	2.8.5	2.8.5	2.8.5
Knox	1.1.0	1.1.0	1.1.0	1.1.0
ApacheDS	2.0.0	2.0.0	2.0.0	2.0.0
Spark	2.4.1	2.4.1	2.4.2	2.4.3
Hive	3.1.1	3.1.1	3.1.1	3.1.1
Tez	0.9.1	0.9.1	0.9.1	0.9.1
Pig	0.14.0	0.14.0	0.14.0	0.14.0
Sqoop	1.4.7	1.4.7	1.4.7	1.4.7
YARN	2.8.5	2.8.5	2.8.5	2.8.5
HDFS	2.8.5	2.8.5	2.8.5	2.8.5
Flink	1.7.2	1.7.2	1.7.2	1.7.2
Druid	0.13.0	0.13.0	0.13.0	0.14.2
HBase	1.4.9	1.4.9	1.4.9	1.4.9
Phoenix	4.14.1	4.14.1	4.14.1	4.14.1
ZooKeeper	3.4.13	3.4.13	3.4.13	3.4.13
Livy	0.6.0	0.6.0	0.6.0	0.6.0

Version         EMR V3.19.0         EMR V3.19.1         EMR V3.20.0         EMR V3.21.0           Presto         0.213         0.213         0.213         0.213           Storm         1.2.2         1.2.2         1.2.2         1.2.2           Impala         2.12.2         2.12.2         2.12.2         2.12.2           Flume         1.8.0         1.8.0         1.8.0         1.8.0           Hue         4.1.0         4.1.0         4.4.0         4.4.0           Oozie         4.2.0         4.2.0         5.1.0         5.1.0           Zeppelin         0.8.0         0.8.1         0.8.1         0.8.1           Ranger         1.2.0         1.2.0         1.2.0         1.2.0           Ganglia         3.7.2         3.7.2         3.7.2         3.7.2           OS         CentOS 7.4         CentOS 7.4         CentOS 7.4         CentOS 7.4         CentOS 7.4           Tensorflow         N/A         N/A         1.8.0         1.8.0         1.8.0           Kafka         2.11-1.1.1         1.1.1         2.11         1.1.1         2.11         1.1.1         1.1.1         1.1.1         1.1.1         1.1.1         2.11         1.2.0         0.2.0 <th></th> <th></th> <th></th> <th></th> <th></th>					
Storm       1.2.2       1.2.2       1.2.2       1.2.2       1.2.2         Impala       2.12.2       2.12.2       2.12.2       2.12.2         Flume       1.8.0       1.8.0       1.8.0         Hue       4.1.0       4.1.0       4.1.0       4.4.0         Oozie       4.2.0       4.2.0       5.1.0       5.1.0         Zeppelin       0.8.0       0.8.1       0.8.1         Ranger       1.2.0       1.2.0       1.2.0       1.2.0         Ganglia       3.7.2       3.7.2       3.7.2       3.7.2         OS       CentOS 7.4       CentOS 7.4       CentOS 7.4       CentOS 7.4       CentOS 7.4         Tensorflow       N/A       N/A       1.8.0       1.8.0         Kafka       2.11-1.1.1       1.1.1       2.11       1.1.1         Superset       0.28.1       0.28.1       0.28.1       0.28.1         Jupyter       N/A       N/A       4.4.0       4.4.0         Analytics Zoo       N/A       N/A       1.0.0       1.0.0         OpenLDAP       N/A       N/A       N/A       N/A       N/A	Version	EMR V3.19.0	EMR V3.19.1	EMR V3.20.0	EMR V3.21.0
Impala         2.12.2         2.12.2         2.12.2         2.12.2         2.12.2           Flume         1.8.0         1.8.0         1.8.0         1.8.0         1.8.0           Hue         4.1.0         4.1.0         4.1.0         4.4.0         4.4.0           Oozie         4.2.0         4.2.0         5.1.0         5.1.0         5.1.0           Zeppelin         0.8.0         0.8.1         0.8.1         0.8.1         0.8.1           Ranger         1.2.0         3.7.2         3.7.2         3.7.2         3.7.2         3.7.2         3.7.2         3.7.2         2.11         1.1.1         1.1.1         2.11         1.1.1         1.1.1         1.1.1         2.11         1.1.1         1.1.1         2.11         1.1.1         1.1.1         3.7.2         3.7.2         3.7.2         3.7.2         3.7.2	Presto	0.213	0.213	0.213	0.213
Flume       1.8.0       1.8.0       1.8.0       1.8.0         Hue       4.1.0       4.1.0       4.1.0       4.4.0         Oozie       4.2.0       4.2.0       5.1.0       5.1.0         Zeppelin       0.8.0       0.8.0       0.8.1       0.8.1         Ranger       1.2.0       1.2.0       1.2.0       1.2.0         Ganglia       3.7.2       3.7.2       3.7.2       3.7.2         OS       CentOS 7.4       CentOS 7.4       CentOS 7.4       CentOS 7.4         Tensorflow       N/A       N/A       1.8.0       1.8.0         Kafka       2.11-1.1.1       1.1.1       2.11       1.1.1         Superset       0.28.1       0.28.1       0.28.1       0.28.1         Jupyter       N/A       N/A       4.4.0       4.4.0         Analytics Zoo       N/A       N/A       1.0.0       0.5.0         Bigboot       N/A       N/A       N/A       N/A       N/A	Storm	1.2.2	1.2.2	1.2.2	1.2.2
Hue       4.1.0       4.1.0       4.1.0       4.4.0         Oozie       4.2.0       4.2.0       5.1.0       5.1.0         Zeppelin       0.8.0       0.8.1       0.8.1         Ranger       1.2.0       1.2.0       1.2.0       1.2.0         Ganglia       3.7.2       3.7.2       3.7.2       3.7.2         OS       CentOS 7.4       CentOS 7.4       CentOS 7.4       CentOS 7.4         Tensorflow       N/A       N/A       1.8.0       1.8.0         Kafka       2.11-1.1.1       1.1.1       2.11       1.1.1         Superset       0.28.1       0.28.1       0.28.1       0.28.1         Jupyter       N/A       N/A       4.4.0       4.4.0         Analytics Zoo       N/A       N/A       1.0.0       1.0.0         Bigboot       N/A       N/A       N/A       N/A       N/A	Impala	2.12.2	2.12.2	2.12.2	2.12.2
Oozie       4.2.0       4.2.0       5.1.0       5.1.0         Zeppelin       0.8.0       0.8.0       0.8.1       0.8.1         Ranger       1.2.0       1.2.0       1.2.0       1.2.0         Ganglia       3.7.2       3.7.2       3.7.2       3.7.2         OS       CentOS 7.4       CentOS 7.4       CentOS 7.4       CentOS 7.4         Tensorflow       N/A       N/A       1.8.0       1.8.0         Kafka       2.11-1.1.1       1.1.1       2.11       1.1.1         Superset       0.28.1       0.28.1       0.28.1       0.28.1         Jupyter       N/A       N/A       4.4.0       4.4.0         Analytics Zoo       N/A       N/A       1.0.0       1.0.0         Bigboot       N/A       N/A       N/A       N/A       N/A	Flume	1.8.0	1.8.0	1.8.0	1.8.0
Zeppelin       0.8.0       0.8.0       0.8.1       0.8.1         Ranger       1.2.0       1.2.0       1.2.0       1.2.0         Ganglia       3.7.2       3.7.2       3.7.2       3.7.2         OS       CentOS 7.4       CentOS 7.4       CentOS 7.4       CentOS 7.4         Tensorflow       N/A       N/A       1.8.0       1.8.0         Kafka       2.11-1.1.1       1.1.1       2.11       1.1.1         Superset       0.28.1       0.28.1       0.28.1       0.28.1         Jupyter       N/A       N/A       4.4.0       4.4.0         Analytics Zoo       N/A       N/A       1.0.0       0.5.0         Bigboot       N/A       N/A       N/A       N/A       N/A         OpenLDAP       N/A       N/A       N/A       N/A       N/A	Hue	4.1.0	4.1.0	4.1.0	4.4.0
Ranger       1.2.0       1.2.0       1.2.0       1.2.0         Ganglia       3.7.2       3.7.2       3.7.2       3.7.2         OS       CentOS 7.4       CentOS 7.4       CentOS 7.4       CentOS 7.4         Tensorflow       N/A       N/A       1.8.0       1.8.0         Kafka       2.11-1.1.1       1.1.1       2.11       1.1.1         Superset       0.28.1       0.28.1       0.28.1       0.28.1         Jupyter       N/A       N/A       4.4.0       4.4.0         Analytics Zoo       N/A       N/A       0.2.0       0.5.0         Bigboot       N/A       N/A       N/A       N/A         OpenLDAP       N/A       N/A       N/A       N/A	Oozie	4.2.0	4.2.0	5.1.0	5.1.0
Ganglia       3.7.2       3.7.2       3.7.2       3.7.2         OS       CentOS 7.4       CentOS 7.4       CentOS 7.4       CentOS 7.4         Tensorflow       N/A       N/A       1.8.0       1.8.0         Kafka       2.11-1.1.1       1.1.1       2.11       1.1.1         Superset       0.28.1       0.28.1       0.28.1       0.28.1         Jupyter       N/A       N/A       4.4.0       4.4.0         Analytics Zoo       N/A       N/A       0.2.0       0.5.0         Bigboot       N/A       N/A       1.0.0       1.0.0         OpenLDAP       N/A       N/A       N/A       N/A	Zeppelin	0.8.0	0.8.0	0.8.1	0.8.1
OS         CentOS 7.4         CentOS 7.4         CentOS 7.4         CentOS 7.4           Tensorflow         N/A         N/A         1.8.0         1.8.0           Kafka         2.11-1.1.1         1.1.1         2.11         1.1.1           Superset         0.28.1         0.28.1         0.28.1         0.28.1           Jupyter         N/A         N/A         4.4.0         4.4.0           Analytics Zoo         N/A         N/A         0.2.0         0.5.0           Bigboot         N/A         N/A         1.0.0         1.0.0           OpenLDAP         N/A         N/A         N/A         N/A	Ranger	1.2.0	1.2.0	1.2.0	1.2.0
Tensorflow         N/A         N/A         1.8.0         1.8.0           Kafka         2.11-1.1.1         1.1.1         2.11         1.1.1           Superset         0.28.1         0.28.1         0.28.1         0.28.1           Jupyter         N/A         N/A         4.4.0         4.4.0           Analytics Zoo         N/A         N/A         0.2.0         0.5.0           Bigboot         N/A         N/A         1.0.0         1.0.0           OpenLDAP         N/A         N/A         N/A         N/A	Ganglia	3.7.2	3.7.2	3.7.2	3.7.2
Kafka       2.11-1.1.1       1.1.1       2.11       1.1.1         Superset       0.28.1       0.28.1       0.28.1       0.28.1         Jupyter       N/A       N/A       4.4.0       4.4.0         Analytics Zoo       N/A       N/A       0.2.0       0.5.0         Bigboot       N/A       N/A       1.0.0       1.0.0         OpenLDAP       N/A       N/A       N/A       N/A	OS	CentOS 7.4	CentOS 7.4	CentOS 7.4	CentOS 7.4
Superset         0.28.1         0.28.1         0.28.1           Jupyter         N/A         N/A         4.4.0         4.4.0           Analytics Zoo         N/A         N/A         0.2.0         0.5.0           Bigboot         N/A         N/A         1.0.0         1.0.0           OpenLDAP         N/A         N/A         N/A         N/A	Tensorflow	N/A	N/A	1.8.0	1.8.0
Jupyter         N/A         N/A         4.4.0         4.4.0           Analytics Zoo         N/A         N/A         0.2.0         0.5.0           Bigboot         N/A         N/A         1.0.0         1.0.0           OpenLDAP         N/A         N/A         N/A         N/A	Kafka	2.11-1.1.1	1.1.1	2.11	1.1.1
Analytics Zoo         N/A         N/A         0.2.0         0.5.0           Bigboot         N/A         N/A         1.0.0         1.0.0           OpenLDAP         N/A         N/A         N/A         N/A	Superset	0.28.1	0.28.1	0.28.1	0.28.1
Bigboot         N/A         N/A         1.0.0         1.0.0           OpenLDAP         N/A         N/A         N/A         N/A	Jupyter	N/A	N/A	4.4.0	4.4.0
OpenLDAP N/A N/A N/A N/A	Analytics Zoo	N/A	N/A	0.2.0	0.5.0
	Bigboot	N/A	N/A	1.0.0	1.0.0
Kudu N/A N/A N/A N/A	OpenLDAP	N/A	N/A	N/A	N/A
	Kudu	N/A	N/A	N/A	N/A

### • EMR V3.15.0 to EMR V3.18.1

Version	EMR V3.15.0	EMR V3.16.0	EMR V3.17.0	EMR V3.18.1
Release date	2018.11	2018.12	2019.1	2019.2
Hadoop	2.7.2	2.7.2-1.3.2	2.7.2	2.8.5
Knox	0.13.0	0.13.0	1.1.0	1.1.0
ApacheDS	2.0.0	2.0.0	2.0.0	2.0.0
Spark	2.3.2	2.3.2-1.0.1	2.3.2	2.3.2
Hive	2.3.3	2.3.3-1.0.3	2.3.3	3.1.1

Version	EMR V3.15.0	EMR V3.16.0	EMR V3.17.0	EMR V3.18.1
Tez	0.9.1	0.9.1-1.0.2	0.9.1	0.9.1
Pig	0.14.0	0.14.0	0.14.0	0.14.0
Sqoop	1.4.7	1.4.7-1.0.0	1.4.7	1.4.7
YARN	2.7.2	2.7.2	2.7.2	2.8.5
HDFS	2.7.2	2.7.2	2.7.2	2.8.5
Flink	1.4.0	1.6.2-1.0.0	1.6.2	1.6.2
Druid	0.12.3	0.12.3-1.0.1	0.12.3	0.13.0
HBase	1.1.1	1.1.1-1.0.2	1.1.1	1.4.9
Phoenix	4.10.0	4.10.0-1.0.0	4.10.0	4.14.1
ZooKeeper	3.4.13	3.4.13	3.4.13	3.4.13
Livy	0.5.0	0.5.0	0.5.0	0.5.0
Presto	0.208	0.208	0.213	0.213
Storm	1.1.2	1.2.2	1.2.2	1.2.2
Impala	2.10.0	2.10.0-1.0.0	2.12.2	2.12.2
Flume	N/A	1.8.0	1.8.0	1.8.0
Hue	4.1.0	4.1.0	4.1.0	4.1.0
Oozie	4.2.0	4.2.0	4.2.0	4.2.0
Zeppelin	0.8.0	0.8.0	0.8.0	0.8.0
Ranger	1.0.0	1.0.0	1.0.0	1.0.0
Ganglia	3.7.2	3.7.2	3.7.2	3.7.2
OS	CentOS 7.4	Cent OS 7.4	CentOS 7.4	Cent OS 7.4
Tensorflow	1.8.0	1.8.0	1.8.0	N/A
Kafka	2.11-1.0.1	2.11-1.1.0	2.11-1.1.1	2.11-1.1.1
Superset	0.27.0	0.28.1	0.28.1	0.28.1
Jupyter	4.4.0	4.4.0	4.4.0	N/A

### • EMR V3.11.0 to EMR V3.14.0

Version	EMR V3.11.0	EMR V3.12.0	EMR V3.13.0	EMR V3.14.0
Release date	2018.6	2018.7	2018.8	2018.10
Hadoop	2.7.2-emr-1.2.14	2.7.2-emr-1.2.14	2.7.2	2.7.2
Knox	0.13.0	0.13.0	0.13.0	0.13.0
ApacheDS	2.0.0	2.0.0	2.0.0	2.0.0
Spark	2.2.1	2.3.1	2.3.1	2.3.1
Hive	2.3.3	2.3.3	2.3.3	2.3.3
Tez	0.9.1	0.9.1	0.9.1	0.9.1
Pig	0.14.0	0.14.0	0.14.0	0.14.0
Sqoop	1.4.6	1.4.7	1.4.7	1.4.7
YARN	2.7.2	2.7.2	2.7.2	2.7.2
HDFS	2.7.2	2.7.2	2.7.2	2.7.2
Flink	1.4.0	1.4.0	1.4.0	1.4.0
Druid	0.11.0	0.12.0	0.12.2	0.12.3
HBase	1.1.1	1.1.1	1.1.1	1.1.1
Phoenix	4.10.0	4.10.0	4.10.0	4.10.0
ZooKeeper	3.4.11	3.4.12	3.4.12	3.4.13
Presto	0.188	0.188	0.208	0.208
Storm	1.1.2	1.1.2	1.1.2	1.1.2
Impala	2.10.0	2.10.0	2.10.0	2.10.0
Hue	4.1.0	4.1.0	4.1.0	4.1.0
Oozie	4.2.0	4.2.0	4.2.0	4.2.0
Zeppelin	0.7.3	0.7.3	0.8.0	0.8.0
Ranger	0.7.3	1.0.0	1.0.0	1.0.0
Ganglia	3.7.2	3.7.2	3.7.2	3.7.2
OS	CentOS 7.4	CentOS 7.4	CentOS 7.4	CentOS 7.4

Version	EMR V3.11.0	EMR V3.12.0	EMR V3.13.0	EMR V3.14.0
Tensorflow	N/A	N/A	1.8.0	1.8.0
Kafka	N/A	N/A	2.11-1.0.1	2.11-1.0.1
Superset	N/A	N/A	0.25.6	0.25.6
Jupyter	N/A	N/A	N/A	N/A
Analytics Zoo	N/A	N/A	N/A	N/A

### • EMR V3.7.1 to EMR V3.10.1

Version	EMR V3.7.1	EMR V3.8.1	EMR V3.9.1	EMR V3.10.1
Release date	2018.1	2018.1	2018.2	2018.4
Hadoop	2.7.2-emr-1.2.10	2.7.2-emr-1.2.12	2.7.2-emr-1.2.13	2.7.2-emr-1.2.14
Knox	0.13.0	0.13.0	0.13.0	0.13.0
ApacheDS	2.0.0	2.0.0	2.0.0	2.0.0
Spark	2.2.1	2.2.1	2.2.1	2.2.1
Hive	2.3.2	2.3.2	2.3.2	2.3.2
Tez	0.8.4	0.8.4	0.8.4	0.9.1
Pig	0.14.0	0.14.0	0.14.0	0.14.0
Sqoop	1.4.6	1.4.6	1.4.6	1.4.6
YARN	2.7.2	2.7.2	2.7.2	2.7.2
HDFS	2.7.2	2.7.2	2.7.2	2.7.2
Flink	N/A	1.4.0	1.4.0	1.4.0
Druid	N/A	N/A	0.11.0	0.11.0
HBase	1.1.1	1.1.1	1.1.1	1.1.1
Phoenix	4.10.0	4.10.0	4.10.0	4.10.0
ZooKeeper	3.4.11	3.4.11	3.4.11	3.4.11
Presto	0.188	0.188	0.188	0.188
Storm	1.0.1	1.0.1	1.0.1	1.1.2
Impala	2.10.0	2.10.0	2.10.0	2.10.0

Version	EMR V3.7.1	EMR V3.8.1	EMR V3.9.1	EMR V3.10.1
Hue	3.12.0	3.12.0	3.12.0	4.1.0
Oozie	4.2.0	4.2.0	4.2.0	4.2.0
Zeppelin	0.7.1	0.7.1	0.7.1	0.7.1
Ganglia	3.7.2	3.7.2	3.7.2	3.7.2
OS	CentOS 7.4	CentOS 7.4	CentOS 7.4	CentOS 7.4

### • EMR V2.9.2 to EMR V2.11.0

Version	EMR V2.9.2	EMR V2.10.0	EMR V2.11.0
Release date	2018.2	2018.4	2018.7
Hadoop	2.7.2-emr-1.2.12	2.7.2-emr-1.2.12	2.7.2-emr-1.2.12
Spark	1.6.3	1.6.3	1.6.3
Hive	2.3.2	2.3.2	2.3.3
Tez	0.8.4	0.9.1	0.9.1
Pig	0.14.0	0.14.0	0.14.0
Sqoop	1.4.6	1.4.6	1.4.6
Hue	3.12.0	4.1.0	4.1.0
Zeppelin	0.7.1	0.7.1	0.7.3
HBase	1.1.1	1.1.1	1.1.1
Phoenix	4.10.0	4.10.0	4.10.0
Storm	1.0.1	1.1.2	1.1.2
Presto	0.188	0.188	0.188.0
Impala	2.10.0	2.10.0	2.10.0
ZooKeeper	3.4.6	3.4.11	3.4.11
Oozie	4.2.0	4.2.0	4.2.0
Ranger	0.7.1	0.7.1	0.7.3
Ganglia	3.7.2	3.7.2	3.7.2
OS	CentOS 7.4	CentOS 7.4	CentOS 7.4

#### • EMR V1.0.0 to EMR V1.3.0

Version	EMR V1.0.0	EMR V1.1.0	EMR V1.2.0	EMR V1.3.0
Release date	2015.11	2016.3	2016.4	2016.5
Hadoop	2.6.0	2.6.0	2.6.0	2.6.0-emr-1.1.1
Spark	1.4.1	1.6.0	1.6.1	1.6.1
Hive	1.0.1	1.0.1	2.0.0	2.0.0
Pig	0.14.0	0.14.0	0.14.0	0.14.0
Sqoop	N/A	N/A	N/A	1.4.6
Hue	N/A	N/A	N/A	3.9.0
Zeppelin	N/A	N/A	N/A	0.5.6
HBase	N/A	N/A	1.1.1	1.1.1
Phoenix	N/A	N/A	N/A	N/A
ZooKeeper	N/A	N/A	3.4.6	3.4.6
Ganglia	3.7.2	3.7.2	3.7.2	3.7.2

### ? Note About Hadoop versions:

To ensure compatibility with Alibaba Cloud Object Storage Service (OSS), the emr-core component is deployed based on an open source Hadoop version. Existing APIs remain unchanged. The version number of the emr-core component is suffixed to the end of the Hadoop version.

## 7.2. Release notes (EMR V5.X series)

### 7.2.1. Release notes for EMR V5.6.X

This topic describes the release notes for E-MapReduce (EMR) V5.6.X, including the release date, updates, and release version information.

### Release date

April 21, 2022 for EMR V5.6.0

### **Updates**

Service	Description
JindoData	The JindoData service is added. The service version is 4.3.0.
JindoSDK	JindoSDK is updated to 4.3.0.

Service	Description	
---------	-------------	--

Spark	Spark is updated to 3.2.1.
Hive	The following issue is fixed: After speculative execution is enabled for Hive on Tez, both the original task and the speculative task are committed.
Presto	The following issue is fixed: After a Hadoop cluster is initialized and the Presto service is added to the cluster, the Presto service fails to start.
DeltaLake	Subqueries can be nested in DML statements.
Hudi	Hudi is updated to 0.10.1.
Iceberg	Iceberg is updated to 0.13.1.
YARN	A feature that allows you to run an ApplicationMaster only on a core node is added.
HBase	The issue of the HBase 2.3.4 kernel is fixed.
ZooKeeper	The configuration of Java Virtual Machine (JVM) parameters is optimized.
Impala	Impala is adapted to JindoSDK 4.3.0.
Sqoop	The PostgreSQL version is updated.
Zeppelin	The issue that the Java Database Connectivity (JDBC) interpreter fails to start is fixed.
Ranger	Delta tables and Hudi tables are supported by the Spark plugin of Ranger 1.2.0.
Flume	Flume is adapted to JindoSDK 4.3.0.
Oozie	Log4j is updated to 2.17.2.
DLF-Auth	DLF-Auth is updated to 2.0.0.

### Release version information

### ○ Hadoop cluster ○ Shuffle Service cluster

Service	Version
HDFS	3.2.1
YARN	3.2.1

Service	Version
Hive	3.1.2
Spark	3.2.1
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
JindoData	4.3.0
DLF-Auth	2.0.0
Iceberg	0.13.1
Hudi	0.10.1
DeltaLake	1.1.0
OpenLDAP	2.4.44
Hue	4.9.0
JindoSDK	4.3.0
HBase	2.3.4
ZooKeeper	3.6.3
Presto	358
Impala	3.4.0
Zeppelin	0.10.1
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	2.1.0
RSS	1.0.0
Alluxio	2.5.0
Kudu	1.14.0

Service	Version
Oozie	5.2.1

Service	Version
RSS	1.0.0

# 7.2.2. Release notes for EMR V5.5.X

This topic describes the release notes for E-MapReduce (EMR) V5.5.X, including the release dates, updates, and release version information.

### Release dates

Version	Date
EMR V5.5.1	March 25, 2022
EMR V5.5.0	February 15, 2022

### **Updates**

#### **EMR V5.5.1**

Note Only OLAP clusters support this version. You can create an OLAP cluster in the new EMR console.

Service	Description
ClickHouse	The default values of some parameters are changed.
StarRocks	StarRocks is updated to 2.1.1.

### **EMR V5.5.0**

Service	Description	
SmartData	The conject is no longer used	
Bigboot	The service is no longer used.	
RSS	<ul> <li>EMR Remote Shuffle Service (ESS) is upgraded to Remote Shuffle Service (RSS). For more information, see RSS.</li> <li>The service functionality and stability are enhanced.</li> </ul>	

Service	Description
JindoSDK	<ul> <li>The architecture of SmartData is upgraded to JindoData.</li> <li>EMR is integrated with JindoSDK for JindoData 4.0.0 for the first time. JindoData connects to Alibaba Cloud Object Storage Service (OSS) and the Alibaba Cloud OSS-HDFS service.</li> </ul>
	<ul> <li>An IF expression can be used in the COUNT DISTINCT function, and the CASE WHEN syntax for the COUNT DISTINCT function is optimized.</li> <li>To use this feature, set</li> </ul>
	spark.sql.optimizer.rewriteConditionalDistinctAggregates to true.
	<ul> <li>Fallback from Shuffle Hash Join to Sort Merge Join is supported.</li> </ul>
	To use this feature, set spark.sql.join.preferSortMergeJoin to false and set spark.sql.join.enableShuffledHashJoinFallback to true.
Spark	<ul> <li>Automatic merging of small files in non-dynamic partitions is supported. The data of merged files is written to tables.</li> </ul>
	To use this feature, set spark.sql.adaptive.merge.output.small.files.enabled to true.
	<ul> <li>The concurrency in scenarios in which the GROUPING SETS clause or the DISTINCT function is used can be automatically adjusted.</li> </ul>
	To use this feature, set spark.sql.execution.optimizeExpand to true.
	Hive on Spark is optimized.
	The syntax of the time travel feature is supported.
	Spark is adapted to JindoSDK.
Tez	Tez is adapted to JindoSDK.
	The issue about batch deletion that occurs on Hive Jindo is fixed.
Hive	• The out of memory (OOM) issue that occurs on HiveServer2 is fixed.
	Hive on Spark is optimized.
	Hive is adapted to JindoSDK.

Service	Description
Presto	<ul> <li>Presto is updated to 358.</li> <li>By default, the MySQL, Iceberg, Hudi, Phoenix, Kudu, and Delta connector types are added. The default configurations of these connector types are updated.</li> <li>Data lake analysis is supported.</li> <li>User-defined functions (UDFs) can be dynamically loaded.</li> <li>Presto is adapted to JindoSDK.</li> </ul>
Delta Lake	<ul> <li>Version update</li> <li>Delta Lake is updated to 1.1.0 and compatible with Spark 3.2.0.</li> <li>All features that are released for commercial use are migrated to Delta Lake 1.1.0.</li> <li>Metadata management</li> <li>The feature of synchronizing metadata changes to metastores is optimized.</li> <li>The statistics on table data are automatically reported to metastores.</li> <li>SQL</li> <li>The SPARTITION SQL syntax is supported.</li> <li>SQL statements can be executed to overwrite data in dynamic partitions.</li> <li>The ADD COLUMN statement can be used to add columns to specified locations (FIRST and AFTER).</li> <li>Enhanced table management capabilities</li> <li>The file size can be dynamically adjusted based on the table size. By default, this feature is enabled.</li> <li>The auto-vacuum feature is supported and enabled by default. Concurrent vacuum operations are supported.</li> <li>The logic of automatic compaction is optimized. By default, the automatic compaction feature is disabled.</li> <li>The Z-ordering syntax is added. Z-ordering-based data processing is accelerated.</li> </ul>
Hudi	<ul> <li>Hudi is updated to 0.10.0.</li> <li>Spark 3.2.0 is supported.</li> <li>JindoFS in block storage mode is supported.</li> </ul>
HDFS	HDFS is adapted to JindoSDK.

Service	Description
YARN	<ul><li>YARN is adapted to RSS memory configurations.</li><li>YARN is adapted to JindoSDK.</li></ul>
Flume	Flume is adapted to JindoSDK.
Impala	Impala is adapted to JindoSDK.
Ranger	<ul><li>Spark 3.2.0 is supported.</li><li>Presto 358 is supported.</li></ul>
HBase	<ul> <li>The issue about default parameter settings is fixed.</li> <li>The issue about the date format of garbage collection (GC) logs is fixed.</li> </ul>
ClickHouse	<ul> <li>The transaction feature is added. For more information, see Use a transaction.</li> <li>The Hadoop Distributed File System (HDFS) and OSS disk types are added to support the separation of hot and cold data. For more information, see Separate hot and cold data by using HDFS and Separate hot and cold data by using OSS.</li> <li>When you use an engine of the Replicated*MergeTree type, the following disk types are supported: zero copy provided by Amazon S3, OSS, and HDFS.</li> <li>The logic of processing data when the ClickHouse component stops working is optimized.</li> </ul>
Iceberg	<ul><li>Iceberg is updated to 0.13.0.</li><li>Presto 358 is supported.</li></ul>
DLF-Auth	<ul><li>Spark 3.2.0 is supported.</li><li>Presto 358 is supported.</li></ul>

Note To view the information about an OLAP cluster, you must log on to the new EMR console.

### ○ Hadoop cluster ○ ClickHouse cluster ○ Shuffle Service cluster ○ OLAP cluster

Service	Version
HDFS	3.2.1
YARN	3.2.1

Service	Version
Hive	3.1.2
Spark	3.2.0
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
DLF-Auth	1.0.4
Iceberg	0.13.0
Hudi	0.10.0
Delta Lake	1.1.0
OpenLDAP	2.4.44
Hue	4.9.0
JindoSDK	4.0.0
HBase	2.3.4
ZooKeeper	3.6.3
Presto	358
Impala	3.4.0
Zeppelin	0.10.2
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	2.1.0
RSS	1.0.0
Alluxio	2.5.0
Kudu	1.14.0
Oozie	5.2.1

Service	Version
ZooKeeper	3.6.3
Ganglia	3.7.2
ClickHouse	21.3.13.9

Service	Version
RSS	1.0.0

Service	Version
ClickHouse	21.3.13.9.2.9
StarRocks	2.1.1
ZooKeeper	3.6.3

# 7.2.3. Release notes for EMR V5.4.X

This topic describes the release notes for E-MapReduce (EMR) V5.4.X, including the release dates, updates, and release version information.

### Release dates

Version	Date
EMR V5.4.3	December 2021
EMR V5.4.2	December 2021
EMR V5.4.1	November 2021
EMR V5.4.0	October 2021

### **Updates**

#### • EMR V5.4.3

The Log4j security vulnerability of all related components is fixed. For more information, see Vulnerability announcement | RCE vulnerability in Apache Log4j 2.

Service	Description
Presto	The Log4j security vulnerability of the Elasticsearch connector is fixed.

Service	Description
Data Lake Formation (DLF) Metastore	<ul> <li>By default, the logging feature for metastores is disabled.         In earlier versions, the feature is enabled by default.     </li> <li>The error caused by an excessively long getTableStats URI for a metastore is fixed.</li> </ul>
Delta Lake	The issue that schema changes fail to be synchronized to a metastore is fixed.
Sqoop	The issue that precision loss for the DECIMAL data type occurs when you use Sqoop to import data to HCatalog tables is fixed.

#### • EMR V5.4.2

Service	Description
SmartData	<ul> <li>SmartData is updated to 3.8.0. For more information, see SmartData 3.8.X overview.</li> <li>Authentication and authorization based on Kerberos and Ranger can be used to manage permissions on data in OSS.</li> </ul>

### • EMR V5.4.1

Service	Description
Smart Dat a	SmartData is updated to 3.7.3. For more information, see SmartData 3.7.X overview.
Oozie	The issue that Jetty Server of Oozie fails to start due to JAR package conflicts in high availability (HA) scenarios is fixed.
Impala	The issue that the no such method error message appears when you query data in DLF metadata tables is fixed.
DLF-Auth	DLF-Auth is updated to 1.0.1.

### • EMR V5.4.0

Service	Description
SmartData	SmartData is updated to 3.7.2. For more information, see SmartData 3.7.X overview.

Service	Description
Spark	<ul> <li>Spark is updated to 3.1.2.</li> <li>In Spark 3.x, the Distinct computing performance is optimized for Spark SQL. The optimization feature is triggered if an aggregation operator contains multiple count(distinct case when) methods.</li> <li>The array-index out of bounds error that is returned when some required statistics for Adaptive Query Execution (AQE) are missing is fixed.</li> <li>Errors related to AQE and data caching in specific scenarios are fixed.</li> </ul>
Hive	In JindoFS in block storage mode, the metadata of multiple Hive tables can be optimized at the same time. By default, this feature is disabled.
Presto	Storage handlers can be used to query data in Delta tables.
Delta Lake	<ul> <li>Delta Lake is updated to 1.0.0.</li> <li>The same Delta Lake connectors are used in Hive 2 and Hive 3.</li> <li>The error that is returned when you use Delta Lake connectors to query data from multi-level partitioned tables is fixed.</li> <li>The SQL syntax that supports various features, such as Data Skipping, Optimize, and Z-ordering, can be used.</li> <li>Metadata can be synchronized to a metastore.</li> </ul>
Hudi	<ul> <li>Hudi is updated to 0.9.0.</li> <li>The issue about the compatibility of sql.extension between Delta Lake and Hudi is fixed.</li> <li>Note Spark 3.1.2 is supported.</li> </ul>
HDFS	By default, the reserved space of NameNode adaptively increases. This way, NameNode enters the Safe mode in a timely manner when the disk space is insufficient.
Storm	The service is no longer used.
Zeppelin	Zeppelin is updated to 0.10.0.
Hue	<ul> <li>The issue that YARN Job Browser sometimes cannot present or terminate jobs is fixed.</li> <li>YARN Job Browser is accessible by default.</li> <li>The Presto protocol is supported by default.</li> </ul>

Service	Description
Druid	The following issue is fixed: After a server is unexpectedly shut down, the related node fails to restart because a PID file is not deleted.
ClickHouse	<ul> <li>Some default configurations are updated.</li> <li>Clusters can be scaled out.</li> <li>The MetaChecker feature is supported.</li> <li>Object Storage Service (OSS) table engines and OSS table functions can be used to read data.</li> </ul>
Iceberg	<ul> <li>Iceberg is updated to 0.12.0-1.0.1.</li> <li>Errors related to the Hive runtime dependency are fixed.</li> </ul>
Knox	The issue that the first access to the Spark UI fails is fixed.
DLF-Auth	The service is added.  The permissions of using Hive or Spark to access DLF can be configured. The supported version is 1.0.0.

○ Hadoop cluster ○ Kafka cluster ○ ClickHouse cluster

Service	Version
HDFS	3.2.1
YARN	3.2.1
Hive	3.1.2
Spark	3.1.2
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	EMR V5.4.0: 3.7.2 EMR V5.4.1: 3.7.3
Bigboot	EMR V5.4.2: 3.8.0
Iceberg	0.12.0
DLF-Auth	1.0.0

Service	Version
Hudi	0.9.0
Delta Lake	1.0.0
OpenLDAP	2.4.44
Hue	4.9.0
HBase	2.3.4
ZooKeeper	3.6.3
Presto	338
Impala	3.4.0
Zeppelin	0.10.0
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	2.1.0
ESS	1.2.0
Alluxio	2.5.0
Kudu	1.14.0
Oozie	5.2.1

Service	Version
ZooKeeper	3.6.3
Ganglia	3.7.2
Kafka	2.4.1
Kafka-Manager	1.3.3.16
OpenLDAP	2.4.44
Knox	1.1.0
Ranger	2.1.0

Service	Version
ZooKeeper	3.6.3
Ganglia	3.7.2
ClickHouse	21.3.13.9

# 7.2.4. Release notes for EMR V5.3.X

This topic describes the release notes for E-MapReduce (EMR) V5.3.X, including the release dates, updates, and release version information.

### Release dates

Version	Date
EMR V5.3.1	September 2021
EMR V5.3.0	August 2021

# **Updates**

#### • EMR V5.3.1

Service	Description
SmartData	SmartData is updated to 3.7.1.
Hue	The issue that Impala cannot be used in a high-security cluster is fixed.
Kudu	Kerberos is supported.
HBase	<ul> <li>The time required to restart HBase in a high-security cluster is reduced.</li> <li>The issue that Spark 3.1.1 cannot be integrated with HBase is fixed.</li> <li>The Graceful Stop process is optimized.</li> </ul>

#### • EMR V5.3.0

Service	Description
SmartData	SmartData is updated to 3.7.0.
Spark	The incompatibility issue between Spark and Delta Lake is fixed.

Service	Description
Hive	In JindoFS in block storage mode, the metadata of multiple Hive tables can be optimized at the same time. By default, this feature is disabled.
Delta Lake	<ul> <li>The partitioning feature of Delta Lake is supported.</li> <li>The incompatibility issue between the command and Spark 3.1.1 is fixed.</li> </ul>
YARN	<ul> <li>Information about app IDs, CPU utilization, and memory usage is added to the REST APIs of containers for nodes.</li> <li>The issue that the Application Master (AM) logs of an automatically released node cannot be viewed is fixed.</li> <li>The issue that a cluster cannot be accessed due to historical state store data is fixed.</li> <li>The data of a node that is automatically released based on the decommissioning logic of auto scaling can be deleted.</li> <li>The Graceful Decommission logic of auto scaling is optimized. The node on which NodeManager runs is marked deprecated only after the NodeManager process is complete.</li> </ul>
ZooKeeper	ZooKeeper is updated to 3.6.3.
Flink	<ul> <li>The SmartData service is added.</li> <li>Object Storage Service (OSS) can be accessed in password-free mode when you log on to a DataFlow-Flink cluster in SSH mode and submit a job in the cluster.</li> </ul>
Impala	The issue that the LIST operation is repeatedly performed on directly deleted OSS partition directories is fixed.
Hue	The UI display exception that occurs when you use Hue together with Oozie is fixed.
Kudu	Kudu is updated to 1.14.0.
ClickHouse	ClickHouse is updated to 21.3.13.9.
Iceberg	The lceberg service is added. lceberg 0.12.0 is used.

○ Hadoop cluster ○ ClickHouse cluster ○ Shuffle Service cluster ○ Kafka cluster

Service	Version
HDFS	3.2.1
YARN	3.2.1
Hive	3.1.2
Spark	3.1.1
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	EMR V5.3.0: 3.7.0
Bigboot	EMR V5.3.1: 3.7.1
Iceberg	0.12.0
Hudi	0.8.0
Delta Lake	0.8.0
OpenLDAP	2.4.44
Hue	4.9.0
HBase	2.3.4
ZooKeeper	3.6.3
Presto	338
Impala	3.4.0
Zeppelin	0.9.0
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	2.1.0
Storm	1.2.2
ESS	1.2.0

Service	Version
Alluxio	2.5.0
Kudu	1.14.0
Oozie	5.2.1

Service	Version
ZooKeeper	3.6.3
Ganglia	3.7.2
ClickHouse	21.3.13.9

Service	Version
ESS	1.2.0

Service	Version
ZooKeeper	3.6.3
Ganglia	3.7.2
Kafka	2.4.1
Kafka-Manager	1.3.3.16
OpenLDAP	2.4.44
Knox	1.1.0
Ranger	2.1.0

# 7.2.5. Release notes for EMR V5.2.X

E-MapReduce (EMR) V5.2.X is the first stable version of the EMR V5.X series. This topic describes the release notes for EMR V5.2.X, including the release date, updates, and release version information.

#### Release date

July 16, 2021 for EMR V5.2.1

### **Updates**

Service	Description
---------	-------------

Service	Description	
SmartData	SmartData is updated to 3.6.1.	
	For more information, see SmartData 3.6.X overview.	
	The issue that the output of the show create table command based on Data Lake Formation (DLF) metadata is inaccurate is fixed.	
	<ul> <li>The default parameters of Hive are optimized to improve the performance of Hive jobs.</li> </ul>	
Hive	• In the EMR console, the parameter names on the <b>hive-env</b> tab of the <b>Configure</b> tab for the Hive service are changed to uppercase. This facilitates the use of the parameters.	
	<ul> <li>The issue that user-defined functions (UDFs) cause HiveServer2 memory leak is fixed.</li> </ul>	
	<ul> <li>The error message that is reported because of the incompatibility between the file system and Hive metastore when you write data to a Hive table is optimized.</li> </ul>	
HDFS	The data compression algorithm Zstandard is supported.	
	Delta Lake is updated to 0.8.0.	
Delta Lake	Spark 3 is supported.	
Flink	Flink is updated to 1.12-vvr-3.0.2.	
	Hudi is updated to 0.8.0.	
Hudi	Hudi can be integrated with Spark SQL.	
	Notice In EMR V5.2.1, Spark (3.1.1) and Kudu (1.11.1) are incompatible with each other.	
	<ul> <li>Delta Lake and Hudi are supported.</li> <li>Remote Shuffle Service is supported.</li> <li>Livy is supported.</li> </ul>	
Spark	<ul> <li>In the EMR console, the parameter names on the spark- defaults tab of the Configure tab for the Spark service are optimized.</li> </ul>	
	<ul> <li>The cost-based optimization (CBO), dynamic partition pruning, and Z-order features are optimized. The performance of these features is 50% higher than in Spark 3.</li> <li>Log Service, DataHub, and Message Queue for Apache RocketMQ can be used as data sources.</li> </ul>	
Tez	The default parameters of Tez are optimized to improve the performance of Tez jobs.	

Service	Description
Ranger	<ul> <li>The warning error contained in logs about starting Spark in Ranger is fixed.</li> <li>The issue that user information fails to be automatically synchronized after Ranger is connected to a Lightweight Directory Access Protocol (LDAP) server is fixed.</li> </ul>
Knox	<ul><li>Knox is adapted to Kudu.</li><li>Knox is adapted to HBase.</li></ul>
Kafka	<ul> <li>The Cruise Control component can be used to provide the balance feature for Kafka clusters.</li> <li>Disks for Kafka clusters are hot-swappable. You can replace a damaged disk without the need to stop the Kafka broker of your cluster.</li> <li>The default values of some parameters are changed.</li> </ul>
Phoenix	The issue that no JDBC driver is found when Hive or Spark SQL is used to access Phoenix tables is fixed.
EMR Remote Shuffle Service (ESS)	Spark 3 is supported.

○ Hadoop cluster ○ Shuffle Service cluster ○ Kafka cluster

Service	Version
HDFS	3.2.1
YARN	3.2.1
Hive	3.1.2
Spark	3.1.1
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	3.6.1
Bigboot	3.6.1
Hudi	0.8.0

Service	Version
OpenLDAP	2.4.44
Hue	4.9.0
HBase	2.3.4
ZooKeeper	3.6.2
Presto	338
Impala	3.4.0
Zeppelin	0.9.0
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	2.1.0
Storm	1.2.2
ESS	1.0.0
Alluxio	2.5.0
Kudu	1.11.1
Oozie	5.1.0

Service	Version
ESS	1.0.0

Service	Version
ZooKeeper	3.6.2
Ganglia	3.7.2
Kafka	2.4.1
Kafka-Manager	1.3.3.16
OpenLDAP	2.4.44
knox	1.1.0

Service	Version
Ranger	2.1.0

# 7.3. Release notes (EMR V4.X series)

# 7.3.1. Release notes for EMR V4.10.X

This topic describes the release notes for E-MapReduce (EMR) V4.10.X, including the release date and updates.

#### Release date

March 23, 2022 for EMR V4.10.0

### **Updates**

Service	Description
SmartData	The service is no longer used.
Bigboot	
JindoSDK	<ul> <li>The architecture of SmartData is upgraded to JindoData.</li> <li>EMR is integrated with JindoSDK for JindoData 4.0.0 for the first time. JindoData connects to Alibaba Cloud Object Storage Service (OSS) and the Alibaba Cloud OSS-HDFS service.</li> </ul>

Service	Description
Spark	<ul> <li>Spark is updated to 2.4.8.</li> <li>The issue that adaptive execution does not take effect in some scenarios is fixed.</li> <li>The issue that statistical aggregate functions are used in different manners in Spark and Hive is fixed.</li> <li>The issue that Spark cannot read valid data of the type from a Hive ORC table is fixed.</li> <li>The default configurations of Thrift Server are optimized.</li> <li>In the EMR console, the parameter names on the sparkdefaults tab of the Configure tab for the Spark service are optimized.</li> <li>Hive on Spark is optimized.</li> <li>The array-index out of bounds error that is returned when some required statistics for Adaptive Query Execution (AQE) are missing is fixed.</li> <li>Errors related to AQE and data caching in specific scenarios are fixed.</li> <li>Log4j Metrics Appender is removed because the configuration is invalid.</li> <li>The null pointer exception that occurs when SparkContext is started is fixed.</li> <li>The data compression algorithm Zstandard is supported.</li> </ul>
Hive	<ul> <li>The issue about HiveServer2 memory leaks caused by userdefined functions (UDFs) is fixed.</li> <li>The issue that the output of the Show Create Table command based on Data Lake Formation (DLF) metadata is inaccurate is fixed.</li> <li>The default parameters of Hive are optimized to improve the performance of Hive jobs.</li> <li>In the EMR console, the parameter names on the hive-env tab of the Configure tab for the Hive service are changed to uppercase. This facilitates the use of the parameters.</li> <li>The error message that is reported because of the incompatibility between the file system and Hive metastore when you write data to a Hive table is optimized.</li> <li>In JindoFS in block storage mode, the metadata of multiple Hive tables can be optimized at the same time. By default, this feature is disabled.</li> </ul>
Ranger	<ul> <li>The warning error contained in logs about starting Spark in Ranger is fixed.</li> <li>The issue that user information fails to be automatically synchronized after Ranger is connected to a Lightweight Directory Access Protocol (LDAP) server is fixed.</li> </ul>

Service	Description
HDFS	<ul> <li>The data compression algorithm Zstandard is supported.</li> <li>By default, the reserved space of NameNode adaptively increases. This way, NameNode enters the Safe mode at the earliest opportunity when the disk space is insufficient.</li> </ul>
YARN	<ul> <li>Information about app IDs, CPU utilization, and memory usage is added to the RESTful APIs of containers for nodes.</li> <li>The issue that the Application Master (AM) logs of an automatically released node cannot be viewed is fixed.</li> <li>The issue that a cluster cannot be accessed due to historical state store data is fixed.</li> <li>The data of a node that is automatically released based on the decommissioning logic of auto scaling can be deleted.</li> <li>The Graceful Decommission logic of auto scaling is optimized. The node on which NodeManager runs is marked deprecated only after the NodeManager process is complete.</li> </ul>
Knox	<ul> <li>Knox is adapted to Kudu.</li> <li>Knox is adapted to HBase.</li> <li>The issue that the first access to the Spark UI fails is fixed.</li> </ul>
Tez	The default parameters of Tez are optimized to improve the performance of Tez jobs.
Sqoop	The issue that precision loss for the DECIMAL data type occurs when you use Sqoop to import data to HCatalog tables is fixed.

Service	Description
	Metadata management
	<ul> <li>The built-in catalog of Spark, instead of an API operation that is called by using the Hive CLI, is used to synchronize metadata and partition information.</li> </ul>
	<ul> <li>The statistics on table data are automatically reported to metastores.</li> </ul>
	• SQL
	<ul> <li>The syntax of the time travel feature is supported.</li> </ul>
	<ul> <li>The DROP PARTITION SQL syntax is supported.</li> </ul>
Delta Lake	<ul> <li>The ADD COLUMN statement can be used to add columns to specified locations (FIRST and AFTER).</li> </ul>
	Enhanced table management capabilities
	<ul> <li>The file size can be dynamically adjusted based on the table size. By default, this feature is enabled.</li> </ul>
	<ul> <li>The auto-vacuum feature is supported and enabled by default. Concurrent vacuum operations are supported.</li> </ul>
	<ul> <li>The logic of automatic compaction is optimized. By default, the automatic compaction feature is disabled.</li> </ul>
	<ul> <li>The Z-ordering syntax is added. Z-ordering-based data processing is accelerated.</li> </ul>
	Hudi is updated to 0.10.0.
Hudi	• The issue about the compatibility of sql.extension between Delta Lake and Hudi is fixed.
	The service is added.
lceberg	The supported version is 0.13.0.
Hue	The issue that garbled characters are displayed when Hue is used to query historical records is fixed.
	<ul> <li>The UI display exception that occurs when you use Hue together with Oozie is fixed.</li> </ul>
	• The issue that YARN Job Browser sometimes cannot present or terminate jobs is fixed.
	YARN Job Browser is accessible by default.
	The Presto protocol is supported by default.
D.E.A. II	The service is added.
DLF-Auth	The supported version is 1.0.4.

Service	Description
HBase	<ul> <li>The time required to restart HBase in a high-security cluster is reduced.</li> <li>The issue that Spark 3.1.1 cannot be integrated with HBase is fixed.</li> <li>The Graceful Stop process is optimized.</li> </ul>
ZooKeeper	ZooKeeper is updated to 3.6.3.
Presto	<ul> <li>Presto is updated to 358.</li> <li>User-defined functions (UDFs) can be dynamically loaded. For more information, see Dynamically load UDFs.</li> <li>Data lake analysis is supported.</li> </ul>
Impala	<ul> <li>The issue that the LIST operation is repeatedly performed on directly deleted OSS partition directories is fixed.</li> <li>The following issue is fixed: The no such method error message appears when you query data in DLF metadata tables.</li> </ul>
Zeppelin	Zeppelin is updated to 0.10.0.
Oozie	The issue that Jetty Server of Oozie fails to start due to JAR package conflicts in high availability (HA) scenarios is fixed.

### C Hadoop cluster

Service	Version
HDFS	3.2.1
YARN	3.2.1
Hive	3.1.2
Spark	2.4.8
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
DLF-Auth	1.0.4
Iceberg	0.13.0

Service	Version
Hudi	0.10.0
Delta Lake	0.6.1
OpenLDAP	2.4.44
Hue	4.9.0
JindoSDK	4.0.0
HBase	2.3.4
ZooKeeper	3.6.3
Presto	358
Impala	3.4.0
Zeppelin	0.10.2
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	2.1.0
RSS	1.0.0
Alluxio	2.5.0
Kudu	1.14.0
Oozie	5.2.1

# 7.3.2. Release notes for EMR V4.9.X

This topic describes the release notes for E-MapReduce (EMR) V4.9.X, including the release date and updates.

### Release date

April 21, 2021 for EMR V4.9.0

# **Updates**

59

Service	Description
SmartData	SmartData is updated to 3.5.0.  For more information, see SmartData 3.5.X.
Spark	<ul> <li>The issue that adaptive execution does not take effect in some scenarios is fixed.</li> <li>The issue that statistical aggregate functions are used in different manners in Spark and Hive is fixed.</li> <li>The issue that Spark cannot read valid data of the CHAR type from a Hive ORC table is fixed.</li> </ul>
HDFS	The SM4 encryption algorithm is supported.
Hue	Hue is updated to 4.9.0.
Alluxio	Alluxio is updated to 2.5.0.
Livy	Livy is updated to 0.7.1.

○ Hadoop cluster ○ Shuffle Service cluster ○ Kafka cluster

Service	Version
HDFS	3.2.1
YARN	3.2.1
Hive	3.1.2
Spark	2.4.7
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	3.5.0
Bigboot	3.5.0
Hudi	0.6.0
OpenLDAP	2.4.44
Hue	4.9.0

61

Service	Version
HBase	2.3.4
ZooKeeper	3.6.2
Presto	338
Impala	3.4.0
Zeppelin	0.9.0
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	2.1.0
Flink	1.10-vvr-1.0.2
Storm	1.2.2
Alluxio	2.5.0
ESS	1.0.0
Kudu	1.11.1
Oozie	5.1.0

Service	Version
ESS	1.0.0

ZooKeeper	3.6.2
Ganglia	3.7.2
Kafka	2.4.1
Kafka-Manager	1.3.3.16
OpenLDAP	2.4.44
Knox	1.1.0
Ranger	2.1.0

# 7.3.3. Release notes for EMR V4.8.X

This topic describes the release notes for E-MapReduce (EMR) V4.8.X, including the release date and updates.

### Release date

March 15, 2021 for EMR V4.8.0

# **Updates**

Service	Description
Smart Dat a	SmartData is updated to 3.4.0.  For more information, see SmartData 3.4.X.
Spark	<ul> <li>Some default configurations are optimized.</li> <li>Performance is optimized. Window-based top-k queries can be pushed down.</li> <li>The capability of reading data from and writing data to Hive tables in the CSV or JSON format is enhanced.</li> <li>All the column names of a table can be omitted in the ANALYZE statement.</li> <li>LDAP authentication can be enabled or disabled with a click.</li> <li>For more information, see Manage LDAP authentication.</li> <li>Spark Beeline is easier to use.</li> </ul>
Hive	<ul> <li>Some default configurations are optimized.</li> <li>Performance is optimized. The cost-based optimization (CBO) feature is enhanced.</li> <li>LDAP authentication can be enabled or disabled with a click.</li> <li>For more information, see Manage LDAP authentication.</li> </ul>
YARN	The risk caused by unauthorized access from a Hadoop cluster to the YARN web UI is fixed. If you access the YARN web UI by using SSH Tunnel, you no longer need to explicitly specify user.name in the URL.
Tez	Some default configurations are optimized.

63

Service	Description
Ranger	<ul> <li>The issue caused by filter pushdown in Spark is fixed.</li> <li>The issue that prevents Presto from being enabled after you disable Presto in Ranger is fixed.</li> <li>LDAP authentication can be enabled or disabled with a click.</li> <li>For more information, see Manage LDAP authentication.</li> </ul>
Hue	LDAP authentication can be enabled or disabled with a click.  For more information, see Manage LDAP authentication.
Impala	<ul> <li>Impala is updated to 3.4.0.</li> <li>Shiro is updated to 1.7.0.</li> <li>Metadata stored in Alibaba Cloud Data Lake Formation (DLF) is supported.</li> <li>Data in the Delta format can be queried.</li> <li>LDAP authentication can be enabled or disabled with a click.</li> <li>For more information, see Manage LDAP authentication.</li> <li>The exception that occurs when you use the INSERT OVERWRITE statement to overwrite data stored in OSS is fixed.</li> </ul>
Hudi	<ul> <li>SQL statements can be executed to query data in Hudi tables.</li> <li>The issue that causes the query results on some data to be inaccurate is fixed.</li> <li>Partition pruning is supported if you query data in a Copy on Write table of Hudi by using Spark.</li> <li>The bucket-based index mechanism is supported to improve write performance.</li> </ul>

Service	Description
Delta Lake	<ul> <li>The issue that prevents metadata from being synchronized to a Hive metastore based on existing Delta tables is fixed.</li> <li>The issue that prevents the MERGE statement from parsing asterisks (* ) in data is fixed.</li> <li>The issue that causes an error to be reported when you convert data in the Parquet format into a Delta table and create table metadata is fixed.</li> <li>The issue that causes the OPTIMIZE command to fail if no files need to be compacted is fixed.</li> <li>A subquery can be used as the source in the MERGE statement.</li> <li>Data can be cached if you use Presto to query data in a Delta table. This improves query efficiency.</li> <li>Impala can be used to query data in Delta tables.</li> </ul>
ESS	<ul> <li>Exceptions in the shuffle read stage, such as ClosedChannelException, IndexOutOfBoundsException, and excessive offheap memory usage, are fixed.</li> <li>The issue that causes NullPointerException (NPE) to be reported after metric monitoring is enabled is fixed.</li> </ul>
HAS	The issue that prevents the admin.keytab file from being initiated again after a HAS installation error is reported is fixed.
Presto	LDAP authentication can be enabled or disabled with a click.  For more information, see Manage LDAP authentication.
HBase	<ul> <li>HBase is updated to 2.2.6.</li> <li>Access control based on Ranger is no longer supported.</li> </ul>
Sqoop	Files in the Parquet format can be imported to OSS.
Superset	<ul> <li>The issue that prevents the admin user from logging on to the web UI is fixed.</li> <li>Datasets are compatible with Druid clusters.</li> <li>Spark SQL datasets are no longer supported.</li> </ul>

Service	Description
Knox	<ul> <li>Access to Presto by using Knox is supported.</li> <li>The issue that causes the Druid web UI to be inaccessible is fixed.</li> <li>The limit that you can access the Ranger web UI based on HTTP only by using Knox in high security mode is removed.</li> </ul>

○ Hadoop cluster ○ Shuffle Service cluster ○ Kafka cluster

Service	Version
HDFS	3.2.1
YARN	3.2.1
Hive	3.1.2
Spark	2.4.7
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	3.4.0
Bigboot	3.4.0
Hudi	0.6.0
OpenLDAP	2.4.44
Hue	4.4.0
HBase	2.3.4
ZooKeeper	3.5.6
Presto	338
Impala	3.4.0
Zeppelin	0.9.0
Flume	1.9.0

Service	Version
Livy	0.6.0
Superset	0.36.0
Ranger	2.1.0
Flink	1.10-vvr-1.0.2
Storm	1.2.2
Alluxio	2.4.1
ESS	1.0.0
Kudu	1.11.1
Oozie	5.1.0

Service	Version
ESS	1.0.0
ZooKeeper	3.5.6
Ganglia	3.7.2
Kafka	2.4.1
Kafka-Manager	1.3.3.16
OpenLDAP	2.4.44
Knox	1.1.0
Ranger	2.1.0

# 7.3.4. Release notes for EMR V4.6.X

This topic describes the release notes for E-MapReduce (EMR) V4.6.X, including the release date and updates.

### Release date

January 15, 2021 for EMR V4.6.0

### **Updates**

Service	Description
Smart Dat a	SmartData is updated to 3.2.0.
	For more information, see SmartData 3.2.X.
	• Spark is updated to 2.4.7.
	• jQuery is updated to 3.5.1.
Spark	<ul> <li>Spark is compatible with Hive to automatically update table and partition sizes.</li> </ul>
	<ul> <li>Spark metadata and job running information can be sent to DataWorks.</li> </ul>
Hive	<ul> <li>Metadata from Alibaba Cloud Data Lake Formation (DLF) in an HCatalog table is supported.</li> </ul>
Tilve	Hive metadata and job running information can be sent to DataWorks.
	The statistics feature for Hive is added.
Metastore	<ul> <li>Metadata from Alibaba Cloud DLF in an HCatalog table is supported.</li> </ul>
	Methods to obtain an STS token are optimized.
UD FG	• jQuery is updated to 3.5.1.
HDFS	HDFS is updated to 3.2.1.
	• YARN is updated to 3.2.1.
YARN	• jQuery is updated to 3.5.1.
	<ul><li>Configurations of Fair Scheduler are adjusted.</li><li>Timeline Server is optimized.</li></ul>
Zeppelin	Zeppelin is updated to 0.9.0.
OpenLDAP	The audit feature is added.
	By default, the SSL port 10636 is enabled.
	Presto can be enabled with one click.
Hue	Presto is supported.
EMRHook	<ul> <li>The EMRHook service is added.</li> <li>hive-hook is used to send Hive metadata and job running information to DataWorks.</li> <li>spark-hook is used to send Spark metadata and job running information to DataWorks.</li> </ul>

O Hadoop cluster O Shuffle Service cluster O Kafka cluster

Service	Version
HDFS	3.2.1
YARN	3.2.1
Hive	3.1.2
Spark	2.4.7
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	3.2.0
Bigboot	3.2.0
Hudi	0.6.0
OpenLDAP	2.4.44
Hue	4.4.0
HBase	2.1.9
ZooKeeper	3.5.6
Presto	338
Impala	3.4.0
Zeppelin	0.9.0
Flume	1.9.0
Livy	0.6.0
Superset	0.36.0
Ranger	2.1.0
Flink	1.10-vvr-1.0.2
Storm	1.2.2

n

Service	Version
ESS	1.0.0
Kudu	1.11.1
Oozie	5.1.0

Service	Version
ESS	1.0.0

ZooKeeper	3.5.6
Ganglia	3.7.2
Kafka	2.4.1
Kafka-Manager	1.3.3.16
OpenLDAP	2.4.44
Knox	1.1.0
Ranger	2.1.0

# 7.3.5. Release notes for EMR V4.5.X

This topic describes the release notes for E-MapReduce (EMR) V4.5.X, including the release date and updates.

#### **EMR V4.5.1**

• Release date

December 13, 2020

- Updates
  - o The issue that occurs when you query partitioned tables by using Hive or Presto is fixed.
  - o EMR V4.5.1 is available only in the China (Hangzhou), China (Shanghai), and China (Beijing) regions.

#### **EMR V4.5.0**

• Release date

December 7, 2020

• New features

Service	Description
---------	-------------

Service	Description
EMR Remote Shuffle Service (ESS)	ESS 1.0.0 is supported.  For more information, see ESS.
Hudi	Hudi 0.6.0 is supported.
Delta Lake	Delta Lake 0.6.1 is supported.

### • Updates

Service	Description
Ranger	<ul><li>Ranger is updated to 2.1.0.</li><li>Ownership-related permissions are supported.</li></ul>
Presto	<ul> <li>Presto is updated to 338.</li> <li>Metadata stored in Alibaba Cloud Data Lake Formation (DLF) is supported.</li> </ul>
Zeppelin	Zeppelin is updated to 0.8.2.
SmartData	SmartData is updated to 3.1.0.  For more information, see SmartData 3.1.x.
Bigboot	Bigboot is updated to 3.1.0.
Hive	<ul> <li>Metadata stored in Alibaba Cloud DLF is supported.</li> <li>Ownership-related permissions of Ranger are supported.</li> </ul>
Spark	Metadata stored in Alibaba Cloud DLF is supported.
DLF Metastore	<ul> <li>The issue that Presto cannot be started in a high-security cluster is fixed.</li> <li>Hive 3 and metadata cache are supported.</li> <li>The issue that occurs when you query data by using Hive or Presto is fixed.</li> </ul>
Impala	Custom configuration of parameters on the <i>catalo gd.flgs, impalad.flgs,</i> and <i>statestored.flgs</i> tabs is supported in the EMR console.
Tez	Vulnerabilities related to autoDeploy on the web UI of Tez are fixed.

Service	Description
OpenLDAP	A rule that determines whether port 10389 is in the waiting state is added.
Hue	Security vulnerabilities of the MySQL backend are fixed.
Kerberos	<ul> <li>Apache Kerby is updated to 2.0.1.</li> <li>The issue that the kadmin principal of an external Kerberos cluster cannot be customized is fixed.</li> </ul>
Sqoop	<ul> <li>File formats such as Parquet, Avro, and ORC are supported.</li> <li>Metadata stored in Alibaba Cloud DLF is supported.</li> </ul>

O Hadoop cluster O Shuffle Service cluster O Kafka cluster

Service	Version
HDFS	3.1.3
YARN	3.1.3
Hive	3.1.2
Spark	2.4.5
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	3.1.0
Bigboot	3.1.0
Hudi	0.6.0
OpenLDAP	2.4.44
Hue	4.4.0
HBase	2.1.9

Service	Version
ZooKeeper	3.5.6
Presto	338
Impala	3.4.0
Zeppelin	0.8.2
Pig	0.14.0
Flume	1.9.0
Livy	0.6.0
Superset	0.36.0
Ranger	2.1.0
Flink	1.10-vvr-1.0.2
Storm	1.2.2
ESS	1.0.0
Kudu	1.11.1
Oozie	5.1.0
Service	Version

Service	Version	
ESS	1.0.0	
ZooKeeper	3.5.6	
Ganglia	3.7.2	
Kafka	2.4.1	
Kafka-Manager	1.3.3.16	
OpenLDAP	2.4.44	
Knox	1.1.0	
Ranger	2.1.0	

# 7.3.6. Release notes for EMR V4.4.X

This topic describes the release notes for E-MapReduce (EMR) V4.4.X, including the release date and updates.

#### Release date

September 15, 2020 for EMR V4.4.1

Service	Description
YARN	<ul> <li>The hadoop/tools/lib directory is deleted from the value of the yarn.application.classpath parameter.</li> <li>Default parameter settings for MapReduce jobs are optimized.</li> </ul>
Hive	Default parameter settings are optimized.
Tez	Delicate parameter sectings and optimized.
Ranger	<ul><li>Impala-based access control is supported.</li><li>The jackson-databind version is updated.</li></ul>
Impala	<ul><li>Integration with Ranger is supported.</li><li>Shiro is updated to 1.6.0.</li></ul>
SmartData	SmartData and Bigboot are updated to 2.7.301.
Bigboot	Smartbaca and bigboot are updated to 2.7.301.
Knox	<ul> <li>The web UI of Tez can be independently viewed. Knox is compatible with Tez on the web UI of YARN.</li> <li>Shiro is updated to 1.6.0.</li> </ul>
EMRDOCT OR	The following issue is fixed: Job information is not collected if a time-based configuration file is empty.
Ganglia	The detection feature of the service RPC port for Hadoop Distributed File System (HDFS) is enabled.
Oozie	<ul> <li>The issue that the web UI cannot be accessed is fixed.</li> <li>The jackson-databind version is updated.</li> </ul>
ZooKeeper	Internal IP addresses can be bound to Elastic Compute Service (ECS) instances to start service ports.
Superset	The startup script is repaired.

Service	Description
Livy	The versions of jackson-databind and Fastjson are updated.
Zepplin	The versions of jackson-databind and Shiro are updated.
HAS	The versions of jackson-databind and Fastjson are updated.
Flume	The Fastjson version is updated.

#### ○ Hadoop cluster ○ Kafka cluster

Service	Version
HDFS	3.1.3
YARN	3.1.3
Hive	3.1.2
Spark	2.4.5
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	2.7.301
Bigboot	2.7.301
OpenLDAP	2.4.44
Hue	4.4.0
HBase	2.1.9
ZooKeeper	3.5.6
Presto	331
Impala	3.4.0
Zeppelin	0.8.1
Pig	0.14.0

Service	Version
Flume	1.9.0
Livy	0.6.0
Superset	0.36.0
Ranger	2.0.0
Flink	1.10-vvr-1.0.2
Storm	1.2.2
Kudu	1.11.1
Oozie	5.1.0
ZooKeeper	3.5.6
Ganglia	3.7.2
Kafka	2.4.1
Kafka-Manager	1.3.3.16
OpenLDAP	2.4.44
Knox	1.1.0
Ranger	2.0.0

# 7.3.7. Release notes for EMR V4.3.X

This topic describes the release notes for E-MapReduce (EMR) V4.3.X, including the release date and updates.

#### Release date

May 20, 2020 for EMR V4.3.0

Service	Description	

Service	Description
Ranger	<ul> <li>Custom deployment of Hadoop Distributed File System (HDFS), Hive, and Spark plug-ins is supported. You can enable plug-ins on required service nodes.</li> <li>The RangerUserSync and RangerAdmin components can be configured in the EMR console.</li> </ul>
Presto	The Kudu client is updated.
Spark	<ul> <li>Spark is updated to 2.4.5.</li> <li>The associated Delta Lake is updated to 0.6.0.</li> <li>The issue that PySpark cannot properly run after Ranger Hive is enabled is fixed.</li> </ul>
HDFS	<ul> <li>The issue that the HDFS_NAMENODE_OPTS parameter does not take effect is fixed.</li> <li>Custom deployment is supported.</li> </ul>
YARN	Custom deployment is supported.
Hive	Custom deployment is supported.
Knox	Information on the web UI of HDFS NameNode in Hadoop 3.X can be viewed.
Zeppelin	The issue that a <i>zepping.keytab</i> file fails to be generated is fixed.
Kafka	Kafka is updated to 2.4.1.
Kudu	Kudu is updated to 1.11.1.
Impala	Issues related to HAProxy are fixed.
Livy	Issues related to xmllint are fixed.
Hue	<ul> <li>Hue can be deployed on gateway clusters.</li> <li>Multiple Hue instances can be enabled on a single node.</li> </ul>

#### ○ Hadoop cluster ○ Kafka cluster

Service	Version
HDFS	3.1.3

Service	Version
YARN	3.1.3
Hive	3.1.1
Spark	2.4.5
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
Smart Dat a	2.6.7
Bigboot	2.6.7
OpenLDAP	2.4.44
Hue	4.4.0
HBase	2.1.9
ZooKeeper	3.5.6
Presto	331
Impala	3.2.0
Zeppelin	0.8.1
Pig	0.14.0
Flume	1.9.0
Livy	0.6.0
Superset	0.28.1
Ranger	2.0.0
Flink	1.10-vvr-1.0.2
Storm	1.2.2
Kudu	1.11.1
Oozie	5.1.0

ZooKeeper	3.5.6

	Ganglia	3.7.2
	Kafka	2.4.1
Kafka	Kafka-Manager	1.3.3.16
	OpenLDAP	2.4.44
	Knox	1.1.0
	Ranger	2.0.0

# 7.4. Release notes (EMR V3.X series) 7.4.1. Release notes for EMR V3.40.X

This topic describes the release notes for E-MapReduce (EMR) V3.40.X, including the release date, updates, and release version information.

#### Release date

April 21, 2022 for EMR V3.40.0

Service	Description
JindoData	The JindoData service is added. The service version is 4.3.0.
JindoSDK	JindoSDK is updated to 4.3.0.
Spark	Spark is updated to 3.2.1.
Hive	<ul> <li>The following issue is fixed: After speculative execution is enabled for Hive on Tez, both the original task and the speculative task are committed.</li> <li>The following issue is fixed: User-defined functions (UDFs) can be called only after you reload the functions.</li> </ul>
Presto	The following issue is fixed: After a Hadoop cluster is initialized and the Presto service is added to the cluster, the Presto service fails to start.
DeltaLake	The incompatibility issue between Delta Lake and Streaming SQL is fixed.
Hudi	Hudi is updated to 0.10.1.
Iceberg	Iceberg is updated to 0.13.1.

Service	Description	
YARN	<ul> <li>A feature that allows you to run an ApplicationMaster only on a core node is added.</li> <li>The issue that the value of the mareduce.map.java.opts parameter does not include taihaodoctor is fixed.</li> </ul>	
ZooKeeper	The configuration of Java Virtual Machine (JVM) parameters is optimized.	
Flink		
Impala	Flink Impala Flume and Druid are adapted to lindoFDV 4.2.0	
Flume	Flink, Impala, Flume, and Druid are adapted to JindoSDK 4.3.0.	
Druid		
Sqoop	The PostgreSQL version is updated.	
Zeppelin	The issue that the Java Database Connectivity (JDBC) interpreter fails to start is fixed.	
Ranger	Hudi tables are supported by the Spark plugin of Ranger 1.2.0.	
Oozie	Log4j is updated to 2.17.2.	
HBase	The issue that the RegionServer of HBase 1.4.9 fails to start is fixed.	
DLF-Auth	DLF-Auth is updated to 2.0.0.	

79

○ Hadoop cluster ○ Druid cluster ○ Presto cluster

Service	Version
HDFS	2.8.5
YARN	2.8.5
Hive	2.3.9
Spark	2.4.8
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7

Service	Version
JindoData	4.3.0
Iceberg	0.13.1
DLF-Auth	2.0.0
Hudi	0.10.1
DeltaLake	0.6.1
OpenLDAP	2.4.44
Hue	4.9.0
JindoSDK	4.3.0
Spark	3.2.1
HBase	1.4.9
ZooKeeper	3.6.3
Presto	358
Impala	3.4.0
Zeppelin	0.10.1
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	1.2.0
Phoenix	4.14.1
RSS	1.0.0
Alluxio	2.5.0
Kudu	1.14.0
Oozie	5.2.1

Service	Version
HDFS	2.8.5
Druid	0.20.1

Service	Version
ZooKeeper	3.6.3
Knox	1.1.0
Ganglia	3.7.2
JindoData	4.3.0
OpenLDAP	2.4.44
JindoSDK	4.3.0
YARN	2.8.5
Superset	0.36.0

Component	Version
Knox	1.1.0
Presto	358
Ganglia	3.7.2
JindoData	4.3.0
lceberg	0.13.1
Hudi	0.10.1
DeltaLake	0.6.1
OpenLDAP	2.4.44
Hue	4.9.0
JindoSDK	4.3.0
Alluxio	2.5.0

# 7.4.2. Release notes for EMR V3.39.X

This topic describes the release notes for E-MapReduce (EMR) V3.39.X, including the release dates, updates, and release version information.

#### Release dates

81

Version	Date
EMR V3.39.2	March 25, 2022

Version	Date
EMR V3.39.1	February 15, 2022

#### Updates

#### EMR V3.39.2

**? Note** Only OLAP and Dataflow clusters support this version. You can create an OLAP or Dataflow cluster in the **new EMR console**.

Service	Description
Flink	<ul> <li>The application performance management (APM) dashboard is optimized. Metrics such as sourceldleTime are added.</li> <li>Alerting by CloudMonitor is supported.</li> </ul>
Kafka	<ul> <li>SSL and Simple Authentication and Security Layer (SASL) configurations are supported.</li> <li>The default values of some parameters are changed.</li> </ul>
ClickHouse	The default values of some parameters are changed.

#### EMR V3.39.1

Service	Description
SmartData	The service is no longer used.
Bigboot	The service is no tonger used.
RSS	<ul> <li>EMR Remote Shuffle Service (ESS) is upgraded to Remote Shuffle Service (RSS). For more information, see RSS.</li> <li>The service functionality and stability are enhanced.</li> </ul>
JindoSDK	<ul> <li>The architecture of SmartData is upgraded to JindoData.</li> <li>EMR is integrated with JindoSDK for JindoData 4.0.0 for the first time. JindoData connects to Alibaba Cloud Object Storage Service (OSS) and the Alibaba Cloud OSS-HDFS service.</li> </ul>
Spark	<ul><li>Hive on Spark is optimized.</li><li>Spark is adapted to JindoSDK.</li></ul>
Tez	Tez is adapted to JindoSDK.
Hive	Hive is adapted to JindoSDK.

Service	Description
Presto	<ul> <li>User-defined functions (UDFs) can be dynamically loaded.</li> <li>The for as of syntax can be used in the time travel feature to query data in a Delta Lake table.</li> <li>An independent Delta Lake catalog is added. Presto provides default configurations for a Delta connector and supports the Optimize, Z-ordering, and Data Skipping features based on the independent Delta Lake catalog.</li> <li>The issue that data in Merge on Read tables of Hudi cannot be queried by using a Hudi connector is fixed. You cannot use a Hive connector to query Merge on Read tables of Hudi.</li> <li>Presto is adapted to JindoSDK.</li> </ul>
Delta Lake	<ul> <li>Metadata management</li> <li>The built-in catalog of Spark, instead of an API operation that is called by using the Hive CLI, is used to synchronize metadata and partition information.</li> <li>The statistics on table data are automatically reported to metastores.</li> <li>SQL</li> <li>The syntax of the time travel feature is supported.</li> <li>The DROP PARTITION SQL syntax is supported.</li> <li>The ADD COLUMN statement can be used to add columns to specified locations (FIRST and AFTER).</li> <li>Enhanced table management capabilities</li> <li>The file size can be dynamically adjusted based on the table size. By default, this feature is enabled.</li> <li>The auto-vacuum feature is supported and enabled by default. Concurrent vacuum operations are supported.</li> <li>The logic of automatic compaction is optimized. By default, the automatic compaction feature is disabled.</li> <li>The Z-ordering syntax is added. Z-ordering-based data processing is accelerated.</li> </ul>
Hudi	Hudi is updated to 0.10.0.
HDFS	HDFS is adapted to JindoSDK.
YARN	YARN is adapted to JindoSDK.
Flume	Flume is adapted to JindoSDK.

Service	Description
Flink	<ul> <li>By default, the <i>lib</i> directory of Flink is uploaded to your HDFS cluster. This way, you can configure the yarn.provided.lib.dirs parameter to use the directory.</li> <li>Flink is adapted to JindoSDK.</li> </ul>
Impala	Impala is adapted to JindoSDK.
Ranger	<ul><li> The issue that Spark History Server fails to start is fixed.</li><li> Ranger is adapted to JindoSDK.</li></ul>
HBase	<ul> <li>The issue about default parameter settings is fixed.</li> <li>The issue about the date format of garbage collection (GC) logs is fixed.</li> <li>The restart issue that occurs if an IP address is configured for RegionServer is fixed.</li> </ul>
Druid	Druid is adapted to JindoSDK.
ClickHouse	<ul> <li>The transaction feature is added. For more information, see         Use a transaction.</li> <li>The logic of processing data when the ClickHouse         component stops working is optimized.</li> </ul>
Iceberg	<ul> <li>Iceberg is updated to 0.13.0.</li> <li>Default configuration items are hidden to improve user experience.</li> </ul>
DLF-Auth	The issue that Spark History Server fails to start is fixed.
StarRocks	The service is added in the new EMR console.  The 2.0.1 version of this service is released.

Note To view the information about an OLAP cluster, you must log on to the new EMR console.

# ○ Hadoop cluster ○ Druid cluster ○ Dataflow cluster ○ ClickHouse cluster ○ Presto cluster ○ OLAP cluster

Service	Version
HDFS	2.8.5
YARN	2.8.5

Service	Version
Hive	2.3.9
Spark	2.4.8
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
Iceberg	0.13.0
DLF-Auth	1.0.4
Hudi	0.10.0
Delta Lake	0.6.1
OpenLDAP	2.4.44
Hue	4.9.0
JindoSDK	4.0.0
Spark	3.2.0
HBase	1.4.9
ZooKeeper	3.6.3
Presto	358
Impala	3.4.0
Zeppelin	0.10.2
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	1.2.0
Phoenix	4.14.1
RSS	1.0.0
Alluxio	2.5.0

Service	Version
Kudu	1.14.0
Oozie	5.2.1

Service	Version
HDFS	2.8.5
Druid	0.20.1
ZooKeeper	3.6.3
Knox	1.1.0
Ganglia	3.7.2
OpenLDAP	2.4.44
JindoSDK	4.0.0
YARN	2.8.5
Superset	0.36.0

Component	Version
HDFS	2.8.5
YARN	2.8.5
ZooKeeper	3.6.3
Knox	1.1.0
Flink	1.13-vvr-4.0.11
OpenLDAP	2.4.44
JindoSDK	4.0.0
Ganglia	3.7.2
Kafka	1.1.1
Kafka-Manager	1.3.3.16

Service	Version
ZooKeeper	3.6.3

Service	Version
Ganglia	3.7.2
ClickHouse	20.8.12.2

Component	Version
Knox	1.1.0
Presto	358
Ganglia	3.7.2
Iceberg	0.13.0
Hudi	0.10.0
Delta Lake	0.6.1
OpenLDAP	2.4.44
Hue	4.9.0
JindoSDK	4.0.0
Alluxio	2.5.0

Service	Version
ClickHouse	20.8.12.2.2.17
StarRocks	2.0.1
ZooKeeper	3.6.3

# 7.4.3. Release notes for EMR V3.38.X

This topic describes the release notes for E-MapReduce (EMR) V3.38.X, including the release dates, updates, and release version information.

#### Release dates

Version	Date
EMR V3.38.3	December 2021
EMR V3.38.2	December 2021
EMR V3.38.1	November 2021

Version	Date
EMR V3.38.0	October 2021

#### Updates

#### • EMR V3.38.3

The Log4j security vulnerability of all related components is fixed. For more information, see Vulnerability announcement | RCE vulnerability in Apache Log4j 2.

Service	Description
Presto	<ul> <li>The error that is reported when Presto in a high availability cluster is used to query data in a Hudi table is fixed.</li> <li>The Log4j security vulnerability of the Elasticsearch connector is fixed.</li> </ul>
DLF Metastore	<ul> <li>By default, the logging feature for metastores is disabled.         In earlier versions, the feature is enabled by default.     </li> <li>The error caused by an excessively long getTableStats URI for a metastore is fixed.</li> </ul>
Delta Lake	The issue that schema changes fail to be synchronized to a metastore is fixed.
Flink	<ul> <li>VVR is updated to 4.0.11. This version supports the following features:</li> <li>The commercial features of Flink Change Data Capture (CDC) are released:</li> <li>Schema evolution is supported.</li> <li>Flink SQL semantics can be used to synchronize an entire database at a time.</li> <li>GeminiStateBackend can be used to store state data in Object Storage Service (OSS).</li> <li>A Hudi connector of Enterprise Edition is provided, and a built-in Data Lake Formation (DLF) catalog is used in the connector to manage metadata.</li> </ul>
Sqoop	The following issue is fixed: Precision loss for the DECIMAL data type occurs when you use Sqoop to import data to HCatalog tables.

#### • EMR V3.38.2

Service	Description
---------	-------------

Service	Description
SmartData	<ul> <li>SmartData is updated to 3.8.0. For more information, see SmartData 3.8.X overview.</li> <li>Authentication and authorization based on Kerberos and Ranger can be used to manage permissions on data in OSS.</li> </ul>

#### • EMR V3.38.1

Service	Description
Smart Dat a	SmartData is updated to 3.7.3. For more information, see SmartData 3.7.X overview.
Spark	<ul> <li>Log4j Metrics Appender is removed because the configuration is invalid.</li> <li>The null pointer exception that occurs when SparkContext is started is fixed.</li> </ul>
Presto	<ul> <li>The following issue is fixed: You must configure host parameters for the Presto service before you can use Presto in a high-availability Hadoop cluster to query data from a Hive table.</li> <li>The following issue is fixed: Presto cannot be started by default when the memory is small.</li> <li>The issue that modifications to worker-jvm do not take effect is fixed.</li> <li>Ranger is supported.</li> </ul>
Impala	The following issue is fixed: The no such method error message appears when you query data in DLF metadata tables.
Ranger	<ul> <li>Presto is supported.</li> <li>The exceptions that occur when you use Ranger to configure the permissions of using Spark to insert data into ORC or Parquet tables are fixed.</li> <li>The following issue is fixed: Hive role permissions configured in Ranger do not take effect after Kerberos is enabled.</li> </ul>
DLF-Auth	<ul> <li>DLF-Auth is updated to 1.0.1.</li> <li>The permissions of using Presto to access DLF can be configured.</li> <li>The issue that data cannot be cached for RAM users is fixed.</li> </ul>

#### • EMR V3.38.0

> Document Version: 20220627

Service	Description	
Smart Dat a	SmartData is updated to 3.7.2. For more information, see SmartData 3.7.X overview.	
Spark	<ul> <li>Spark is updated to 2.4.8.</li> <li>Both of Spark 2.4.8 and Spark 3.1.2 are supported.</li> <li>Note Delta and Remote Shuffle Service are not supported in Spark 3.</li> </ul>	
	<ul> <li>In Spark 3.x, the Distinct computing performance is optimized for Spark SQL. The optimization feature is triggered if an aggregation operator contains multiple count(distinct case when) methods.</li> <li>The array-index out of bounds error that is returned when some required statistics for Adaptive Query Execution (AQE)</li> </ul>	
	<ul> <li>are missing is fixed.</li> <li>Errors related to AQE and data caching in specific scenarios are fixed.</li> </ul>	
Hive	Hive is updated to 2.3.9.	
Presto	<ul> <li>An independent Presto cluster can be created.</li> <li>Presto is updated to 358.</li> <li>Notice This version does not support Ranger.</li> <li>Connectors such as Hudi and MySQL connectors are supported by default, and the default configurations are updated.</li> <li>Auto scaling is supported by Presto clusters.</li> <li>Note To use the auto scaling feature, .</li> <li>Data lake analysis is supported.</li> </ul>	
Delta Lake	<ul> <li>The same Delta Lake connectors are used in Hive 2 and Hive 3.</li> <li>The error that is returned when you use Delta Lake connectors to query data from multi-level partitioned tables is fixed.</li> </ul>	
Hudi	<ul> <li>Hudi is updated to 0.9.0.</li> <li>The issue about the compatibility of sql.extension between Delta Lake and Hudi is fixed.</li> </ul>	

Service	Description
HDFS	By default, the reserved space of NameNode adaptively increases. This way, NameNode enters the Safe mode in a timely manner when the disk space is insufficient.
	<ul> <li>Flink is updated to 1.13-vvr-4.0.10, which corresponds to Apache Flink 1.13.1.</li> <li>Commercial connectors, such as a Hologres connector, are added to Flink.</li> </ul>
Flink	<ul> <li>Metric reporters are added to report metrics on the APM dashboard.</li> </ul>
	<ul> <li>A SchemaRegistry-based Kafka catalog is added to the Kafka connector. This way, you can read data from or write data to existing Kafka topics without the need to execute DDL statements.</li> </ul>
Storm	The service is no longer used.
Zeppelin	Zeppelin is updated to 0.10.0.
Ranger	If Presto 358 is used, Ranger cannot be used to configure permissions related to Presto.
lh-a	<ul> <li>The issue that YARN Job Browser sometimes cannot present or terminate jobs is fixed.</li> </ul>
Hue	<ul><li>YARN Job Browser is accessible by default.</li><li>The Presto protocol is supported by default.</li></ul>
Druid	The following issue is fixed: After a server is unexpectedly shut down, the related node fails to restart because a PID file is not deleted.
	Some default configurations are updated.
	Clusters can be scaled out.
ClickHouse	<ul> <li>The MetaChecker feature is supported.</li> <li>OSS table engines and OSS table functions can be used to read data.</li> </ul>
	<ul> <li>Table-level custom ZooKeeper addresses are supported.</li> </ul>
	The service is added.
Iceberg	The supported versions range from 0.12.0 to 1.0.1.
Knox	The issue that the first access to the Spark UI fails is fixed.
	The service is added.
DLF-Auth	The permissions of using Hive or Spark to access DLF can be configured. The supported version is 1.0.0.

Service	Description
ESS	ESS is updated to 1.2.0.

O Hadoop cluster O Druid cluster O Dataflow cluster O ClickHouse cluster O Presto cluster

Service	Version
HDFS	2.8.5
YARN	2.8.5
Hive	2.3.9
Spark	2.4.8
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	EMR V3.38.0: 3.7.2 EMR V3.38.1: 3.7.3
Bigboot	EMR V3.38.2: 3.8.0
Iceberg	0.12.0
DLF-Auth	1.0.0
Hudi	0.9.0
Delta Lake	0.6.1
OpenLDAP	2.4.44
Hue	4.9.0
Spark	3.1.2
HBase	1.4.9
ZooKeeper	3.6.3
Presto	358
Impala	3.4.0
Zeppelin	0.10.0

Service	Version
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	1.2.0
Phoenix	4.14.1
ESS	1.2.0
Alluxio	2.5.0
Kudu	1.14.0
Oozie	5.2.1

Service	Version
HDFS	2.8.5
Druid	0.20.1
ZooKeeper	3.6.3
Knox	1.1.0
Ganglia	3.7.2
SmartData	3.7.2
Bigboot	5.7.2
OpenLDAP	2.4.44
YARN	2.8.5
Superset	0.36.0

Cluster mode	Component	Version
	HDFS	2.8.5
	YARN	2.8.5
	ZooKeeper	3.6.3
	Knox	1.1.0
Flink		

	Flink	EMR V3.38.0: 1.13-vvr-4.0.10 EMR V3.38.3: 1.13-vvr-4.0.11
	SmartData	EMR V3.38.0: 3.7.2
	Bigboot	EMR V3.38.3: 3.8.0
	OpenLDAP	2.4.44
Kafka	ZooKeeper	3.6.3
	Ganglia	3.7.2
	Kafka	1.1.1
	Kafka-Manager	1.3.3.16
	OpenLDAP	2.4.44
	Knox	1.1.0
	Ranger	1.2.0

Service	Version
ZooKeeper	3.6.3
Ganglia	3.7.2
ClickHouse	20.8.12.2

Component	Version
Knox	1.1.0
Presto	358
Ganglia	3.7.2
SmartData	EMR V3.38.0: 3.7.2 EMR V3.38.1: 3.7.3
Bigboot	EMR V3.38.2: 3.8.0
Hudi	0.9.0
Delta Lake	0.6.1
OpenLDAP	2.4.44
Hue	4.9.0
Alluxio	2.5.0

### 7.4.4. Release notes for EMR V3.37.X

This topic describes the release notes for E-MapReduce (EMR) V3.37.X, including the release dates, updates, and release version information.

#### Release dates

Version	Date
EMR V3.37.1	September 2021
EMR V3.37.0	August 2021

#### **Updates**

#### • EMR V3.37.1

Service	Description
SmartData	SmartData is updated to 3.7.1.
Hue	The issue that Impala cannot be used in a high- security cluster is fixed.
Kudu	Kerberos is supported.

#### • EMR V3.37.0

95

Service	Description
SmartData	SmartData is updated to 3.7.0.
Spark	The incompatibility issue between Spark and Delta Lake is fixed.
	<ul> <li>The connectors for Delta Lake are updated. This way, you can use a storage handler to create tables and query data.</li> </ul>
Delta Lake	<ul> <li>The exception that occurs when you perform an INSERT OVERWRITE operation on a partitioned table is fixed.</li> </ul>
	<ul> <li>The exception that occurs when you perform an OPT IMIZE operation to write virtual fields to a file in the G-SCD scenario is fixed.</li> </ul>

Service	Description
YARN	<ul> <li>Information about app IDs, CPU utilization, and memory usage is added to the REST APIs of containers for nodes.</li> <li>The issue that the Application Master (AM) logs of an automatically released node cannot be viewed is fixed.</li> <li>The data of a node that is automatically released based on the decommissioning logic of auto scaling can be deleted.</li> <li>The Graceful Decommission logic of auto scaling is optimized. The node on which NodeManager runs is marked deprecated only after the NodeManager process is complete.</li> </ul>
ZooKeeper	ZooKeeper is updated to 3.6.3.
Flink	<ul> <li>The SmartData service is added.</li> <li>Object Storage Service (OSS) can be accessed in password-free mode when you log on to a DataFlow-Flink cluster in SSH mode and submit a job in the cluster.</li> </ul>
Impala	The issue that the LIST operation is repeatedly performed on directly deleted OSS directories is fixed.
Hue	The UI display exception that occurs when you use Hue together with Oozie is fixed.
Kudu	Kudu is updated to 1.14.0.
ClickHouse	Some default configurations are updated.

O Hadoop cluster O Druid cluster O Flink cluster O Kafka cluster O ClickHouse cluster

Service	Version
HDFS	2.8.5
YARN	2.8.5
Hive	2.3.8
Spark	2.4.7
Knox	1.1.0
Tez	0.9.2

Service	Version
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	EMR V3.37.0: 3.7.0
Bigboot	EMR V3.37.1: 3.7.1
Hudi	0.8.0
Delta Lake	0.6.1
OpenLDAP	2.4.44
Hue	4.9.0
HBase	1.4.9
ZooKeeper	3.6.3
Presto	338
Impala	3.4.0
Zeppelin	0.9.0
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	1.2.0
Storm	1.2.2
Phoenix	4.14.1
ESS	1.0.0
Alluxio	2.5.0
Kudu	1.14.0
Oozie	5.2.1

Service	Version
HDFS	2.8.5
Druid	0.20.1

Service	Version
ZooKeeper	3.6.3
Knox	1.1.0
Ganglia	3.7.2
SmartData	EMR V3.37.0: 3.7.0
Bigboot	EMR V3.37.1: 3.7.1
OpenLDAP	2.4.44
YARN	2.8.5
Superset	0.36.0

Service	Version
HDFS	2.8.5
YARN	2.8.5
ZooKeeper	3.6.3
Knox	1.1.0
Flink	1.12-vvr-3.0.2
SmartData	EMR V3.37.0: 3.7.0
Bigboot	EMR V3.37.1: 3.7.1
OpenLDAP	2.4.44

Service	Version
ZooKeeper	3.6.3
Ganglia	3.7.2
Kafka	1.1.1
Kafka-Manager	1.3.3.16
OpenLDAP	2.4.44
Knox	1.1.0
Ranger	1.2.0

Service	Version
ZooKeeper	3.6.3
Ganglia	3.7.2
ClickHouse	20.8.12.2

# 7.4.5. Release notes for EMR V3.36.X

This topic describes the release notes for E-MapReduce (EMR) V3.36.X, including the release date, updates, and release version information.

#### Release date

July 16, 2021 for EMR V3.36.1

Service	Description
SmartData	SmartData is updated to 3.6.1.  For more information, see SmartData 3.6.X overview.
Hive	<ul> <li>Hive is updated to 2.3.8.</li> <li>The issue that the output of the command based on Data Lake Formation (DLF) metadata is inaccurate is fixed.</li> <li>The default parameters of Hive are optimized to improve the performance of Hive jobs.</li> <li>In the EMR console, the parameter names on the hive-env tab of the Configure tab for the Hive service are changed to uppercase. This facilitates the use of the parameters.</li> <li>The error message that is reported because of the incompatibility between the file system and Hive metastore when you write data to a Hive table is optimized.</li> </ul>
HDFS	The data compression algorithm Zstandard is supported.
Flink	Flink is updated to 1.12-vvr-3.0.2.  Note Flink is removed from Hadoop clusters.
Hudi	<ul><li>Hudi is updated to 0.8.0.</li><li>Hudi can be integrated with Spark SQL.</li></ul>

Service	Description
Spark	<ul> <li>In the EMR console, the parameter names on the spark-defaults tab of the Configure tab for the Spark service are optimized.</li> <li>The performance of log collection is optimized.</li> <li>The data compression algorithm Zstandard is supported.</li> </ul>
Impala	The core dump error reported when you use HDFS is fixed.
Tez	The default parameters of Tez are optimized to improve the performance of Tez jobs.
Knox	<ul> <li>Knox is adapted to Kudu.</li> <li>Knox is adapted to Impala.</li> <li>Knox is adapted to HBase.</li> </ul>
Phoenix	The issue that no JDBC driver is found when Hive or Spark SQL is used to access Phoenix tables is fixed.
ClickHouse	The application performance management (APM) feature is available. This feature is used to support monitoring and alerting.

⊙ Hadoop cluster ⊙ Druid cluster ⊙ Flink cluster ⊙ Kafka cluster ⊙ ClickHouse cluster

Service	Version
HDFS	2.8.5
YARN	2.8.5
Hive	2.3.8
Spark	2.4.7
Knox	1.1.0
Tez	0.9.2
Ganglia	3.7.2
Sqoop	1.4.7
SmartData	3.6.1
Bigboot	3.6.1
Hudi	0.8.0

101

Service	Version
OpenLDAP	2.4.44
Hue	4.9.0
HBase	1.4.9
ZooKeeper	3.6.2
Presto	338
Impala	3.4.0
Zeppelin	0.9.0
Flume	1.9.0
Livy	0.7.1
Superset	0.36.0
Ranger	1.2.0
Storm	1.2.2
Phoenix	4.14.1
ESS	1.0.0
Alluxio	2.5.0
Kudu	1.10.0
Oozie	5.2.1

Service	Version
HDFS	2.8.5
Druid	0.20.1
ZooKeeper	3.6.2
Knox	1.1.0
Ganglia	3.7.2
SmartData	3.6.1
Bigboot	3.6.1
OpenLDAP	2.4.44

Service	Version
YARN	2.8.5
Superset	0.36.0

Service	Version
HDFS	2.8.5
YARN	2.8.5
ZooKeeper	3.6.2
Knox	1.1.0
Flink	1.12-vvr-3.0.2
OpenLDAP	2.4.44

Service	Version
ZooKeeper	3.6.2
Ganglia	3.7.2
Kafka	1.1.1
Kafka-Manager	1.3.3.16
OpenLDAP	2.4.44
Knox	1.1.0
Ranger	1.2.0

Service	Version
ZooKeeper	3.6.2
Ganglia	3.7.2
ClickHouse	20.8.12.2

# 7.4.6. Release notes for EMR V3.35.X

This topic describes the release notes for E-MapReduce (EMR) V3.35.X, including the release date and updates.

#### Release date

#### April 21, 2021 for EMR V3.35.0

#### **Updates**

Component	Description
SmartData	SmartData is updated to 3.5.0.  For more information, see SmartData 3.5.X.
Spark	<ul> <li>The issue that adaptive execution does not take effect in some scenarios is fixed.</li> <li>The issue that statistical aggregate functions are used in different manners in Spark and Hive is fixed.</li> <li>The issue that Spark cannot read valid data of the CHAR type from a Hive ORC table is fixed.</li> </ul>
HDFS	The SM4 encryption algorithm is supported.
Hue	Hue is updated to 4.9.0.
Alluxio	Alluxio is updated to 2.5.0.
Druid	<ul><li>Druid is updated to 0.20.1.</li><li>Security is enhanced.</li></ul>
Livy	Livy is updated to 0.7.1.

# 7.4.7. Release notes for EMR V3.34.X

This topic describes the release notes for E-MapReduce (EMR) V3.34.X, including the release date and updates.

#### Release date

March 15, 2021 for EMR V3.34.0

Component	Description
SmartData	SmartData is updated to 3.4.0.  For more information, see SmartData 3.4.X.

Component	Description
Spark	<ul> <li>Some default configurations are optimized.</li> <li>Performance is optimized. Window-based top-k queries can be pushed down.</li> <li>The capability of reading data from and writing data to Hive tables in the CSV or JSON format is enhanced.</li> <li>All the column names of a table can be omitted in the ANALYZE statement.</li> <li>LDAP authentication can be enabled or disabled with a click.</li> <li>For more information, see Manage LDAP authentication.</li> <li>Spark Beeline is easier to use.</li> </ul>
Hive	<ul> <li>Some default configurations are optimized.</li> <li>Performance is optimized. The cost-based optimization (CBO) feature is enhanced.</li> <li>LDAP authentication can be enabled or disabled with a click.</li> <li>For more information, see Manage LDAP authentication.</li> <li>Calcite is updated to 1.12.0.</li> <li>The hive.security.authorization.sqlstd.confwhit elist.append parameter is added.</li> </ul>
Presto	LDAP authentication can be enabled or disabled with a click.  For more information, see Manage LDAP authentication.
YARN	The risk caused by unauthorized access from a Hadoop cluster to the YARN web UI is fixed. If you access the YARN web UI by using SSH Tunnel, you no longer need to explicitly specify user.name in the URL.
ZooKeeper	ZooKeeper is updated to 3.6.2.
Flink	The <i>config.sh</i> file is updated during initialization to fix the issue related to the HADOOP_CLASSPATH parameter.

Component	Description
Impala	<ul> <li>Impala is updated to 3.4.0.</li> <li>Shiro is updated to 1.7.0.</li> <li>Data Lake Formation (DLF) metadata is supported.</li> <li>Data in the Delta format can be queried.</li> <li>LDAP authentication can be enabled or disabled with a click.</li> <li>For more information, see Manage LDAP authentication.</li> </ul>
Tez	Some default configurations are optimized.
HAS	The issue that prevents the admin.keytab file from being initiated again if a HAS installation error is reported is fixed.
Ranger	<ul> <li>The issue caused by filter pushdown in Spark is fixed.</li> <li>The issue that prevents Presto from being enabled after you disable Presto in Ranger is fixed.</li> <li>LDAP authentication can be enabled or disabled with a click.</li> <li>For more information, see Manage LDAP authentication.</li> </ul>
Knox	The link to access Druid 0.20.0 by using Knox is fixed.
Hue	LDAP authentication can be enabled or disabled with a click.  For more information, see Manage LDAP authentication.
Hudi	<ul> <li>SQL statements can be executed to query data in Hudi tables.</li> <li>The issue that causes the query results on some data to be inaccurate is fixed.</li> <li>Partition pruning is supported if you query data in a Copy on Write table of Hudi by using Spark.</li> <li>The bucket-based index mechanism is supported to improve write performance.</li> </ul>

Component	Description
Delta Lake	<ul> <li>The issue that prevents metadata from being synchronized to a Hive metastore based on existing Delta tables is fixed.</li> <li>The issue that prevents the MERGE statement from parsing asterisks (* ) in data is fixed.</li> <li>The issue that causes an error to be reported when you convert data in the Parquet format into a Delta table and create table metadata is fixed.</li> <li>The issue that causes the OPTIMIZE command to fail if no files need to be compacted is fixed.</li> <li>A subquery can be used as the source in the MERGE statement.</li> <li>Data can be cached if you use Presto to query data in a Delta table. This improves query efficiency.</li> <li>Impala can be used to query data in Delta tables.</li> </ul>
Superset	<ul> <li>The issue that prevents the admin user from logging on to the web UI is fixed.</li> <li>Datasets are compatible with Druid clusters.</li> <li>Spark SQL datasets are no longer supported.</li> </ul>
Sqoop	Files in the Parquet format can be imported to OSS.
Alluxio	Alluxio is updated to 2.4.1.
Phoenix	Field settings are supported by Hive on Phoenix.
Pig	Pig is removed.

# 7.4.8. Release notes for EMR V3.33.X

This topic describes the release notes for E-MapReduce (EMR) V3.33.X, including the release date and updates.

#### Release date

January 15, 2021 for EMR V3.33.0

Component	Description
SmartData	SmartData is updated to 3.2.0.  For more information, see SmartData 3.2.X.

Component	Description
Spark	<ul> <li>Spark is updated to 2.4.7.</li> <li>jQuery is updated to 3.5.1.</li> <li>Spark is compatible with Hive to automatically update table and partition sizes.</li> <li>Spark metadata and job running information can be sent to DataWorks.</li> </ul>
Hive	<ul> <li>Hive is updated to 2.3.7.</li> <li>Metadata from Alibaba Cloud Data Lake Formation (DLF) in an HCatalog table is supported.</li> <li>Hive metadata and job running information can be sent to DataWorks.</li> </ul>
Metastore	<ul> <li>The statistics feature for Hive is added.</li> <li>Metadata from Alibaba Cloud DLF in an HCatalog table is supported.</li> <li>Methods to obtain an STS token are optimized.</li> </ul>
HDFS	jQuery is updated to 3.5.1.
YARN	<ul> <li>jQuery is updated to 3.5.1.</li> <li>Configurations of Fair Scheduler are adjusted.</li> <li>Timeline Server is optimized.</li> </ul>
Zeppelin	Zeppelin is updated to 0.9.0.
Ranger	<ul><li>Audit log configurations of Hive are added.</li><li>Configurations of Log4j Audit are added.</li></ul>
OpenLDAP	<ul> <li>The audit feature is added.</li> <li>By default, the SSL port 10636 is enabled.</li> <li>Presto can be enabled with one click.</li> </ul>
Knox	<ul> <li>Vulnerabilities of the Spring framework are fixed.</li> <li>The issues caused when you view the Executors page on the web UI of Spark are fixed.</li> <li>The issues occurred on the job status page of the web UI of Oozie are fixed.</li> </ul>
Hue	Presto is supported.
Druid	Druid is updated to 0.20.0.

Component	Description
EmrHook	<ul> <li>Software services are added.</li> <li>hive-hook: sends Hive metadata and job running information to DataWorks.</li> <li>spark-hook: sends Spark metadata and job running information to DataWorks.</li> </ul>

# 7.4.9. Release notes of EMR V3.32.X

This topic describes the release notes of E-MapReduce (EMR) V3.32.X, including the release date and updates.

#### Release date

November 23, 2020 for EMR V3.32.0

Component	Description
Smart Dat a	SmartData is updated to 3.1.0.  For more information, see SmartData 3.1.x.
Alluxio	<ul> <li>Alluxio 2.4.0 is supported.</li> <li>The default configurations can be adjusted based on the sizes of cluster nodes.</li> <li>By default, EMR HDFS is used as UnderFS to provide out-of-box services.</li> <li>Alluxio OSS UnderFS is enhanced to support the new features of different OSS versions.</li> <li>Engines such as Hadoop, Hive, Spark, and Presto are supported.</li> </ul>
Hudi	Hudi 0.6.0 is supported.
Spark	The data collection feature of JindoTable can be enabled or disabled.

Component	Description
Hive	<ul> <li>The connection leak issue of the HiveServer connection pool is fixed.</li> <li>The data collection feature of JindoTable can be enabled or disabled.</li> <li>The performance of ADD COLUMN is optimized.</li> <li>The issue that causes data read from Hudi tables to be invalid is fixed.</li> <li>The default configurations can be adjusted based on the sizes of cluster nodes.</li> </ul>
HDFS	The number of supported snapshots is increased.
YARN	The default configurations can be adjusted based on the sizes of cluster nodes.
Tez	The default configurations can be adjusted based on the sizes of cluster nodes.
Sqoop	The issue that files in the Avro format cannot be imported is fixed.

# 7.4.10. Release notes of EMR V3.30.X

This topic describes the release notes of E-MapReduce (EMR) V3.30.X, including the release date and updates.

#### Release date

October 26, 2020 for EMR V3.30.0

Component	Description
SmartData	SmartData is updated to 3.0.0.  For more information, see SmartData 3.0.x.
Spark	<ul> <li>Metadata from Alibaba Cloud Data Lake Formation (DLF) is supported.</li> <li>Has dependencies are updated to 2.0.1.</li> <li>Issues caused by backticks (`) in Streaming SQL are fixed.</li> <li>The JAR package of Delta is removed. Delta is separately deployed.</li> <li>Logs are stored in HDFS directories.</li> </ul>

Component	Description
Hive	<ul> <li>Metadata from Alibaba Cloud DLF is supported.</li> <li>The issue caused when you read an empty Delta table directory and write data into a dummy file is fixed.</li> <li>Has dependencies are updated to 2.0.1.</li> </ul>
Presto	<ul> <li>Metadata from Alibaba Cloud DLF is supported.</li> <li>The limits on reading a Delta table are eliminated.</li> <li>The issue of no Java Virtual Machine (JVM) configurations in high security mode is fixed.</li> <li>Has dependencies are updated to 2.0.1.</li> </ul>
HDFS	<ul> <li>The Hot-swap Disk Module mode is supported.</li> <li>Has dependencies are updated to 2.0.1.</li> </ul>
YARN	<ul> <li>Issues related to YARN RMZKStateStore are fixed.</li> <li>Snappy files generated from Log Service are supported.</li> <li>Directory configurations in Local mode of MapReduce are modified to resolve the problems that occur in directory permission checks.</li> <li>The Hot-swap Disk Module mode is supported.</li> <li>Logs are stored in HDFS directories.</li> <li>Has dependencies are updated to 2.0.1.</li> </ul>
Zookeeper	<ul> <li>Internal IP addresses can be bound to Elastic Compute Service (ECS) instances to start service ports.</li> <li>Has dependencies are updated to 2.0.1.</li> </ul>
Flink-Vvp	<ul> <li>Flink-VVP is updated to a version in the range of 1.11 to 2.2.2.</li> <li>SQL and Autopilot features are supported.</li> </ul>
	Note Only Dataflow clusters support Flink-VVP.

Component	Description
Flink	<ul> <li>Data can be written into Object Storage Service (OSS) in cache mode. Exactly-once semantics are supported based on the Flink checkpoint and the data sources that support data retransmission.</li> <li>Features of open source Flink 1.11.1 are supported. An SQL statement can insert data into different destination tables or partitions and generate multiple channels of output.</li> <li>Has dependencies are updated to 2.0.1.</li> </ul>
Impala	<ul> <li>Custom configuration of parameters on the catal ogd.flgs, impalad.flgs, and statestored.flgs tabs is supported.</li> <li>Shiro is updated to 1.6.0.</li> <li>Has dependencies are updated to 2.0.1.</li> </ul>
Tez	<ul> <li>Default settings of AM memory parameters are optimized.</li> <li>Has dependencies are updated to 2.0.1.</li> </ul>
HAS	
Storm	
Zeppelin	
Ranger	
OpenLDAP	
Oozie	Has dependencies are updated to 2.0.1.
Knox	
Kafka	
HUE	
HBase	
Druid	

# 7.4.11. Release notes of EMR V3.29.X

This topic describes the release notes of E-MapReduce (EMR) V3.29.X, including the release date and updates.

#### Release date

July 29, 2020 for EMR V3.29.0

## **Updates**

Component	Description
Bigboot	<ul> <li>Bigboot is updated to 2.7.301.</li> <li>JindoDistCp allows you to write data into OSS in IA or Archive mode.</li> <li>Features of Jindo Fuse are enhanced to support multiple namespaces.</li> <li>The metadata caching feature of JindoFS in cache mode is enhanced.</li> </ul>
Spark	<ul> <li>Spark is updated to 2.4.5.2.0.</li> <li>A third-party metastore is supported.</li> <li>The datalake metastore-client parameter is added.</li> </ul>
Hive	<ul> <li>Hive is updated to 2.3.5.6.0.</li> <li>A third-party metastore is supported.</li> <li>The datalake metastore-client parameter is added.</li> </ul>
Presto	Presto is updated to 338.
Ranger	<ul> <li>The Ranger software package is updated to 1.2.0-1.5.0.</li> <li>Presto 338 is supported.</li> <li>The Description parameter is added to the configuration file.</li> </ul>
HDFS	The space reserved for DataNode is configured as needed.
Knox	Knox is compatible with Impala, later Flink versions, and PAI.
Druid	Druid is updated to 0.18.1.
SmartData	SmartData is updated to 2.7.301.

# 7.4.12. Release notes of EMR V3.28.X

This topic describes the release notes of E-MapReduce (EMR) V3.28.X, including the release date, new features, and updates.

#### Release date

## June 12, 2020 for EMR V3.28.0

## **New features**

Component	Description
Bigboot	<ul> <li>JindoTable is released for data statistics based on the popularity of tables or partitions.</li> <li>Complete storage policies in block storage mode are supported. Tiered storage policies, including the IA and Archive storage classes, are supported.</li> <li>The data migration tool Jindo DistCp is used.</li> <li>The performance of Jindo Fuse is improved and issues related to Jindo Fuse are fixed.</li> <li>The integration of JFS Scheme on Jindo Job Committer and a Hive engine is optimized in cache mode.</li> <li>The weight for reading OSS data in block storage mode is configured to alleviate the pressure on reading locally cached data.</li> <li>JindoFS is divided into the following software modules: Bigboot (management layer), SmartData (distributed service), and JindoFS SDK. These modules are independently updated and maintained.</li> </ul>

Component	Description
Flink	Open source Flink is upgraded to the enterprise- level Ververica Platform. Ververica Platform is customized based on open source Flink 1.10 and provides value-added features such as self- developed storage engine Gemini.
Bigboot	Bigboot is updated to 2.7.0.
Delta	<ul><li>Delta is updated to 0.6.0.</li><li>Delta code is decoupled from Spark code.</li></ul>
Spark	<ul> <li>Spark is updated to 2.4.5.</li> <li>Spark is compatible with the Streaming SQL scripts of DataFactory.</li> <li>Delta 0.6.0 is supported.</li> </ul>
Hive	Delta 0.6.0 is supported.

Component	Description
Ranger	<ul> <li>Custom deployment of HDFS, Hive, and Spark is supported.</li> <li>You can configure parameters on the ranger-admin-site and ranger-ugsync-site tabs for the Ranger service in the EMR console.</li> </ul>
HDFS	If no DataNode is available when you write data to HDFS, the relevant DataNode exception information (HDFS-9023) is displayed.
Hue	<ul> <li>Hue can be deployed on Gateway clusters.</li> <li>Multiple Hue instances can be deployed on a single node.</li> </ul>
DataFactory	Delta 0.6.0 is supported.
Druid	Druid is updated to 0.18.0.
Knox	<ul> <li>Knox is updated to a version that ranges from 1.0.7 to 1.1.0.</li> <li>Information on the web UI of HBase can be viewed.</li> </ul>

# 7.4.13. Release notes of EMR V3.27.X

This topic describes the release notes of E-MapReduce (EMR) V3.27.X, including release dates, new features, and updates.

#### Release dates

Version	Release date
EMR V3.27.0	April 29, 2020
EMR V3.27.1	May 8, 2020
EMR V3.27.2	May 20, 2020

## **New features**

Feature	Description
---------	-------------

Feature	Description
Custom component deployment	Components on the master node can be customized. The following components are supported:  Hadoop  Spark  Hive  Zookeeper  Presto
Graceful release during auto scaling	Graceful release during auto scaling is supported. After you enable graceful release, the nodes that are removed during a scale-in are not immediately released. Instead, they are released after tasks are completed within the specified period of time.

# Updates

Component	Description
Spark	<ul> <li>Partition fields of the date type are supported in the cube.</li> <li>The stack depth in the spark-submit scripts is increased.</li> </ul>
Delta	<ul> <li>DDL statements, including CREATE, SHOW, and DESCRIBE, are enhanced.</li> <li>The OPTIMIZE statement for Z-order is supported.</li> </ul>
Knox	<ul> <li>Information on the web UI of Druid can be viewed.</li> <li>Deployment on multiple master nodes is supported.</li> </ul>
Hive	<ul> <li>The magic committer in an HCatalog table is supported.</li> <li>Some outdated default configurations are removed.</li> </ul>
Bigboot	<ul> <li>Bigboot is updated to 2.6.3.</li> <li>Deployment on multiple master nodes is supported.</li> </ul>
SmartData	<ul> <li>SmartData is updated to 2.6.3.</li> <li>Deployment on multiple master nodes is supported.</li> </ul>

Component	Description
Ranger	<ul><li>Solr is supported.</li><li>PrestoSQL 311 is supported.</li></ul>
Tez	The ScratchDir option can be used to specify a temporary directory in OSS.
Presto	Presto is updated to 331.
Druid	Druid is updated to 0.17.1.
Superset	Superset is updated to 0.35.2.
Sqoop	<ul> <li>The MySQL JDBC JAR package is updated to 5.1.48.</li> <li>Custom encoding is supported by usingmysq l-charset in MySQL direct export mode.</li> </ul>

# 7.4.14. Release notes of EMR V3.26.X

This topic describes the release notes of E-MapReduce (EMR) V3.26.X, including release dates and updates.

#### Release dates

Version	Release date
EMR-3.26.3	April 16, 2020

Component	Description
Bigboot	Bigboot is updated to 2.6.3.
Smart Dat a	Tablestore metadata and high availability for namespaces are supported.
Hive	The direct committer in an HCatalog table is supported.
YARN	JindoOssCommitter is used as the default committer.
HDFS	JindoFS is updated.
Spark	JindoOssCommitter is used as the default committer.

## 7.4.15. Release notes of EMR V3.25.X

This topic describes the release notes of E-MapReduce (EMR) V3.25.X, including the release date, new feature, and updates.

#### Release date

January 13, 2020 for EMR V3.25.0

#### New feature

Operations related to Ranger Presto are supported.

## **Updates**

Component	Description
Ranger	<ul> <li>The database of Ranger Admin is initialized in a high-availability cluster.</li> <li>The security issues that are caused when you run scripts on Ranger UserSync are fixed.</li> </ul>
Spark	<ul> <li>Parameters related to Delta, such as spark.sql</li> <li>.extensions , can be configured in the EMR console.</li> <li>Data from a Delta table can be read by using Hive</li> </ul>
Delta	<ul> <li>to avoid manual configuration of InputFormat.</li> <li>The ALTER TABLE SET TBLPROPERTIES and UNSET TBLPROPERTIES statements are supported.</li> </ul>
Hive	The issue that causes execution failures of MapReduce tasks in automatic local mode is fixed.
Presto	<ul> <li>Presto is updated to 310.</li> <li>Joda-Time is updated to 2.10.5.</li> </ul>
Tez	<ul> <li>Tez is updated to 0.9.2.</li> <li>The issue that the application running progress cannot be properly displayed on the web UI of Tez is fixed.</li> <li>The issue that you cannot view the application running history on the web UI of Tez is fixed.</li> </ul>
Impala	The issue that an LZO table cannot be accessed by using Impala is fixed.
HDFS	JAR packages related to Mongo-Hadoop are removed.

Component	Description
ZooKeeper	ZooKeeper is updated to 3.5.6.
YARN	Information on the web UI of Tez can be viewed. You can set yarn.resourcemanager.system- metrics-publisher.enabled to true on the yarn- site tab.
Bigboot	Bigboot is updated to 2.2.3.
Smart Dat a	• The renaming operation is supported in OSS cache mode.
Knox	Dependencies are updated.
Oozie	Dependencies are updated.

# 7.4.16. Release notes of EMR V3.24.X

This topic describes the release notes of E-MapReduce (EMR) V3.24.X, including the release date, new features, and updates.

#### Release date

November 18, 2019 for EMR V3.24.0

#### **New features**

Component	Description
Delt <i>a</i>	<ul> <li>The following SQL statements are supported:         ALTER, CONVERT, CREATE, CTAS, DELETE, DESC,         INSERT, MERGE, OPTIMIZE, UPDATE, and VACUUM.</li> <li>The syntax of OPTIMIZE is optimized and built in         Delta.</li> <li>The Hive connector is supported.</li> <li>The features of open source Delta are supported.</li> </ul>
Grafana	Grafana 6.4.2 is supported in independent Flink clusters.
Prometheus	Prometheus 2.13.0 is supported in independent Flink clusters.
Alertmanager	Alertmanager 0.19.0 is supported in independent Flink clusters.

Component	Description
Tensorflow on Spark	<ul> <li>The TensorFlow framework can be deployed on Spark. TensorFlow and Spark are integrated to optimize task scheduling and data exchange, and support the entire deep learning process from the pre-processing of data to deep learning training.</li> <li>Streaming tasks are supported.</li> </ul>

# Updates

Component	Description
	<ul> <li>The cache mode of JindoFS is optimized. The usage of OssFileSystem is integrated into the cache mode to cache data and metadata. Data caching and metadata caching are controlled by using separate switches and are disabled by default.</li> </ul>
	<ul> <li>The read/write performance of the block storage mode and cache mode is optimized.</li> </ul>
SmartData	<ul> <li>Hot data cached in local disks is measured more accurately and cleared in a timely manner to ensure that disk usage does not exceed the upper limit.</li> </ul>
	<ul> <li>Both the block storage mode and cache mode are supported in gateway clusters.</li> </ul>
	<ul> <li>The separated deployment of a single JindoFS cluster from multiple computing clusters is supported.</li> </ul>
	Parameters related to Delta are supported.
Spark	<ul><li>The Spark plug-in can be configured in Ranger.</li><li>JindoCube is updated to 0.3.0.</li></ul>
	SQL statement compatibility can be checked.
	<ul> <li>Hive 2.3.5 and Hadoop 2.8.5 are released as a combination.</li> </ul>
Hive	<ul> <li>When Hive is restarted, the content in hiveserver2- site.xml is not synchronized to hive-site.xml in the spark-conf folder.</li> </ul>
	<ul> <li>The MSCK command can be used to add incremental directories.</li> </ul>
	• The bug triggered by the reuse of a Tez container in Hive is fixed.
	The MSCK command can be used to optimize column directories.

Component	Description
Bigboot	Bigboot is updated to 2.2.1. Native code can be executed on all server models.
Ranger	<ul> <li>The deployment mode of the Spark plug-in is modified.</li> <li>The bug that the emr-header-2 node of a high-availability cluster cannot obtain keytab files is fixed.</li> </ul>
Kudu	The startup logic is repaired.
ZooKeeper	Four-letter-word commands can be configured. By default, all four-letter-word commands are enabled.
HDFS	HDFS is compatible with JindoFS.
YARN	<ul> <li>The default value of yarn.scheduler.capacity.node-locality-delay is changed to -1.</li> <li>YARN is compatible with JindoFS.</li> </ul>
Has	Has is interconnected with OpenLDAP.
OpenLDAP	OpenLDAP is interconnected with Has.
Presto	Presto is updated to 0.228.
Kafka	Bad disks (d1 instance family) are automatically removed.
Druid	Druid is updated to 0.16.0.
Flume	Flume is updated to 1.9.0.
Flink	<ul> <li>Flink is updated to 1.9.1.</li> <li>Independent Flink clusters are supported. An IP address whitelist must be configured for an independent Flink cluster.</li> </ul>

# 7.4.17. Release notes of EMR V3.23.X

This topic describes the release notes of E-MapReduce V3.23.X, including the release date and updates.

## Release date

September 18, 2019

Component	Description
Druid	<ul> <li>Upgraded Druid to version 0.15.1.</li> <li>Added the Router component.</li> <li>Upgraded Fastjson to version 1.2.60.</li> </ul>
Spark	<ul> <li>Updated the code for Spark SQL Thrift Server to fix the issue where IsolatedClassLoader cannot load classes in certain cases.</li> <li>Refactored the code related to Spark transactions to improve stability.</li> <li>Fixed the issue where optimized row columnar (ORC) files cannot be read or written after the built-in Hive is upgraded to version 2.3.</li> <li>Supports the MERGE INTO syntax.</li> <li>Supports the SCAN and STREAM syntaxes.</li> <li>Supports exactly-once semantics (EOS) for the Structured Streaming Kafka sink.</li> <li>Upgraded Delta to version 0.4.0.</li> </ul>
Hive	<ul> <li>Removed Hive hooks configured in earlier versions of Hive.</li> <li>Supports using multiple COUNT(DISTINCT) for hive.groupby.skew in data optimization.</li> <li>Fixed the issue of data loss when joining tables with different bucket versions.</li> </ul>
Flink	Upgraded Flink to version 1.8.2.
Bigboot	<ul> <li>Upgraded the HDFS FsImage analysis tool.</li> <li>Updated the JAR package for Object Storage Service (OSS) to fix the issue where the non- Daemon thread keeps alive after the main thread is terminated.</li> </ul>
Kafka	<ul> <li>Supports the DepolymentSet-Aware replication distribution strategy in Kafka.</li> <li>Removed the Fastjson dependency.</li> </ul>
HDFS	<ul> <li>Optimized the deployment logic of the JAR package for SmartData OSS.</li> <li>Updated the JAR package for SmartData OSS.</li> </ul>
Flume	Upgraded Fastjson to version 1.2.60.
Tensorflow on Spark	This is a new component.

Component	Description
Has	Upgraded Fastjson to version 1.2.60.
Livy	Upgraded Fastjson to version 1.2.60.

# 7.4.18. Release notes of EMR V3.22.X

This topic describes the release notes of E-MapReduce V3.22.X, including the release date, new features, and updates.

#### Release date

July 28, 2019

#### **New features**

Component	Description
Kudu	<ul> <li>The new component, Kudu, is added. Kudu fills in the gaps in the Hadoop ecosystem. Similar to HBase, you can use Kudu to quickly read, write, and modify data. In addition, you can use Kudu to query and analyze tremendous Parquet data or larger amount of data stored in HDFS.</li> <li>Kudu provides the C++ API and Java API for secondary development.</li> <li>Kudu integrates Impala, Spark, and Hive Metastore.</li> <li>The Kudu component of E-MapReduce V3.22.0 is based on the open-source Apache Kudu 1.10.0.</li> </ul>
OpenLDAP	<ul> <li>The new component, OpenLDAP, is added to replace the previous component ApacheDS.</li> <li>High availability is provided.</li> </ul>

ponent		Description
--------	--	-------------

Component	Description
JindoFS	<ul> <li>Provides two storage modes:</li> <li>Block mode: JindoFS uses Object Storage Service (OSS) as the backend storage. In the block storage mode, JindoFS stores data as blocks in OSS and uses Namespace Service to maintain metadata. The block mode provides better performance than the cache mode when you read and write data or query metadata. The block mode supports multiple storage policies, including WARM, COLD, HOT, TEMP, and ALL_HDD. The WARM policy stores one local backup and one backup in OSS. The COLD policy stores only one backup in OSS. The HOT policy stores multiple local backups and one backup in OSS. The TEMP policy stores only one local backups. The default storage policy is WARM. You can set different storage policies for directories as required.</li> <li>Cache mode: This mode is compatible with OSS. In the cache mode, JindoFS stores data files as objects in OSS. When you access OSS objects stored in JindoFS, JindoFS can cache data and metadata of these OSS objects in the local cluster so that you can quickly access them next time. The cache mode provides multiple policies for you to synchronize metadata as required.</li> <li>Offers an external client:</li> <li>JindoFS provides an external client so that you can access JindoFS from the outside of an E-MapReduce cluster. You can use the external client to access the data cached by JindoFS in E-MapReduce clusters. In addition, the performance of accessing data through the external client is worse than that of data access within E-MapReduce clusters.</li> <li>The cache mode is compatible with the existing OSS semantics. JindoFS accelerates data caching in E-MapReduce clusters. Therefore, you can use the OSS client to directly access JindoFS from the outside of E-MapReduce clusters. For example, you can use the OSS SDK or OssFileSystem of E-MapReduce clusters. For example, you can use the OSS SDK or OssFileSystem of E-MapReduce, such as Spark, Flink, Hive, MapReduce, Impala, and Presto.</li> <li>If you need to separate d</li></ul>
OssFileSystem	<ul> <li>Added the feature of automatically detecting bad disks. This feature can resolve cache write failures caused by bad disks when you write data to OSS.</li> <li>Completed the configurations of OssFileSystem.</li> </ul>

Component	Description
Bigboot	<ul> <li>Upgraded Bigboot to version 2.0.0.</li> <li>Supports multiple namespaces, OSS, multiple storage modes, and access from external clients.</li> <li>Fixed the issue where the Bigboot monitor status is abnormal during server restart.</li> <li>Improved the service specifications of Kudu.</li> <li>Supports a validity check for the specification of each component.</li> </ul>
Hadoop	<ul> <li>HDFS</li> <li>Adapted to HDFS Federation. You can create an HDFS Federation cluster through custom configurations or API operations to avoid formatting namenodes more than one time.</li> <li>Optimized the logic for detecting bad disks. If you use a local disk, the system detects the bad disk when the datanode blockreport is triggered by the dfsadmin command.</li> <li>YARN</li> <li>Fixed the issue where the list of history MapReduce jobs is not updated when YARN container logs are stored in JindoFS or JindoFS-based OssFileSystem.</li> </ul>

Component	Description
	Relational cache
	Supports using a relational cache to accelerate data queries through precomputing. You can create a relational cache to pre-compute data. During a data query, Spark Optimizer automatically detects an appropriate relational cache, optimizes the SQL execution plan, and continues data computing based on the relational cache. This accelerates data queries. For example, you can use relational caches to implement multidimensional online analytical processing (MOLAP), generate data reports, create data dashboards, and synchronize data across clusters.
	<ul> <li>Supports using DDL to perform operations such as CACHE, UNCACHE, ALTER, and SHOW. A relational cache supports all data sources and data formats of Spark.</li> </ul>
	<ul> <li>Supports updating caches automatically or by using the REFRESH command.</li> <li>Supports incremental caching based on the specified partitions.</li> </ul>
	<ul> <li>Supports optimizing the SQL execution plan based on a relational cache.</li> </ul>
	Streaming SQL
	Normalized the parameter settings of Stream Query Writer.
	Optimized the schema compatibility check of Kafka data tables.
	<ul> <li>Supports automatically registering a schema with Schema Registry for a Kafka data table that does not have a schema.</li> </ul>
Spark	o Optimized log information recorded when a Kafka schema is incompatible.
	<ul> <li>Fixed the issue where the column name must be explicitly specified when the query result is written to a Kafka data table.</li> </ul>
	<ul> <li>Removed the restriction that streaming SQL queries only support the Kafka and LogHub data sources.</li> </ul>
	• Delta
	Added the Delta component. You can use Spark to create a Delta data source to perform streaming data writing, transactional reading and writing, data verification and data backtracking. For more information, see Delta details.
	<ul> <li>You can call the DataFrame API to read data from or write data to Delta.</li> </ul>
	<ul> <li>You can call the Structured Streaming API to read or write data by using Delta as the data source or sink.</li> </ul>
	<ul> <li>You can call the Delta API to update, delete, merge, vacuum, and optimize data.</li> </ul>
	<ul> <li>You can use SQL statements to create Delta tables, import data to Delta, and read data from Delta tables.</li> </ul>
	Others
	<ul> <li>Supports primary keys and foreign keys. This is a constraint feature.</li> </ul>
	<ul> <li>Resolved JAR conflicts such as the servlet conflict.</li> </ul>
Flink	Supports the log rollback feature of Log4j.
Kafka	Supports the log rollback feature of Log4j.
	Upgraded Fastjson.

Component	Description
Zeppelin	Upgraded the dependent commons-lang3 package to version 3.7 to fix the issue where PySpark cannot write data to OSS. For more information, see Spark 2.4 incompatibility with commons-lang3 in Zeppelin.
Ranger	Supports the SHOW GRANTS command.
Analytics-Zoo	Fixed the NumPy installation error.
Impala	Supports Apache Kudu 1.10.0.
Presto	Upgraded Presto to version 0.221.
ZooKeeper	Upgraded ZooKeeper to version 3.5.5.

# 7.4.19. Release notes of versions earlier than E-MapReduce V3.22.X

This topic describes the release notes of versions earlier than E-MapReduce V3.22.X.

#### E-MapReduce V3.1.1

- Upgraded the operating system to CentOS 7.2.
- Upgraded Spark to version 2.1.1.
- Upgraded emr-core to version 1.2.6.
- Fixed several bugs of AccessKey-free operations in Object Storage Service (OSS).

#### E-MapReduce V3.0.2

- Upgraded emr-core to version 1.2.5.
- Supports AccessKey-free OSS operations in more regions.
- Adjusted the AccessKey replacement strategy of Resource Access Management (RAM) roles.
- Fixed several bugs related to Hive and Hadoop.

#### E-MapReduce V3.0.1

- Supports table management in an interactive and unified manner, and stores Hive metadata in an external unified database. All clusters that use external Hive metadata share the same metadata.
- Upgraded emr-core to version 1.2.4 to improve the read and write performance of OSS.
- Upgraded Spark to version 2.0.2.
  - **Note** E-MapReduce V3.0.1 is fully compatible with E-MapReduce V3.0.0.

#### E-MapReduce V3.0.0

Released the initial version of E-MapReduce.