

Elastic MapReduce

Release Notes and Announcements

Product Documentation



Copyright Notice

©2013-2023 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

Release Notes and Announcements

- Release Notes

- Announcements

 - Alarm Policy Migration

- Security Announcements

 - Notice for Apache Log4j 2 RCE Vulnerability

Release Notes and Announcements

Release Notes

Last updated : 2023-12-26 17:33:42

May 2023

Update	Description	Release Date	Documentation
Added support for federation monitoring.	Query of monitoring metrics for the Router component and alerts are supported for HDFS federations.	2023-05-18	HDFS Monitoring Metrics
Added support for scaling based on the load of subqueues.	Auto-scaling based on the load of subqueues is supported.	2023-05-18	Custom Scaling Configuration

April 2023

Update	Description	Release Date	Documentation
Released the StarRocks cluster version: StarRocks v1.4.0.	The StarRocks cluster version StarRocks v1.4.0 is released, corresponding to the open-source StarRocks v2.5.3.	2023-04-18	Product Releases and Component Versions

March 2023

Update	Description	Release Date	Documentation
Added table analysis in Kudu.	Supported viewing information by table, Tablet, and TabletServer. Supported analyzing Tablet-level read and write QPS, storage size and other information.	2023-03-24	Kudu Table Analysis

Added statistics by Region in HBase.	Added Region-level read QPS, write QPS and other information to help locate hot spots.	2023-03-24	HBase Table Analysis
Added more monitoring metrics and alarms.	For HiveMetaStore, more than 14 monitoring metrics and alarms (including the number of opened and connected processes, Get Table requests, and currently active requests) were added. For HiveServer2, more than 20 monitoring metrics and alarms (including the current requests and query submission time/average execution time) were added. For YARN, alarms for the queue CPU and memory usage in percentages were added.	2023-03-24	Hive Monitoring Metrics YARN Monitoring Metrics
Added more health status events.	Two events were added: exception in service role health status and timeout of service role health status.	2023-03-24	Cluster Events
Optimized log accuracy.	The log search accuracy was optimized, with search by keyword and by phrase supported.	2023-03-24	Log Search
Released the Kafka cluster version: Kafka v2.0.0.	Released the Kafka cluster version Kafka v2.0.0, corresponding to the open-source Kafka v2.4.1.	2023-03-23	Versions of Components
Released the StarRocks cluster version: StarRocks v1.3.0.	Released the StarRocks cluster version StarRocks v1.3.0, corresponding to the open-source Kafka v2.4.3.	2023-03-10	Versions of Components
Released the Hadoop cluster version: JDK11-Beta-based EMR v4.0.0.	Released the Hadoop cluster version JDK11-Beta-based EMR v4.0.0, with all of its components running in the JDK 11 environment.	2023-03-10	Versions of Components
Suspended the provision of the container-based EMR service.	The container-based EMR service was unavailable for purchase for feature update in progress.	2023-03-10	Overview

February 2023

Update	Description	Release Date	Documentation
Added more metrics	Statistics of jobs submitted by users, resources	2023-	YARN Job

for job query in YARN.	consumed, and other information were added for job query in YARN.	02-09	Query
Added table analysis in Hive (based on Hive MetaStore).	Supported viewing global metrics at key dimensions such as total tables, total storage size, table count by access time, and trends. Supported viewing multi-dimensional data of tables and their partitions.	2023-02-09	Hive Table Analysis
Added the file-level analysis feature in HDFS.	Supported viewing HDFS data distribution at key dimensions such as file count and file size. Supported viewing and downloading the lists of top large files and top small files.	2023-02-09	HDFS File Storage Analysis

January 2023

Update	Description	Release Date	Documentation
Supported automatic node replacement	Task and router nodes can be monitored, and abnormal nodes can be replaced automatically, mitigating the risk of service unavailability due to node exceptions.	2023-01-16	
Supported local disk repair	Local disk replacement events can be automatically monitored. After disk replacement, you can initialize the new disk in the console on your own.	2023-01-16	

December 2022

Update	Description	Release Date	Documentation
Supported Java-GC online analysis.	Java-GC online log analysis is supported to help you troubleshoot process exceptions due to GC, collect and record GC logs in real time, and analyze the logs.	2022-12-06	
Supported role-level monitoring of Kafka.	Kafka adds more than 70 role-level monitoring metrics and supports configuring alarms for core monitoring metrics.	2022-12-06	
Added more default event-	More event-related metrics are added, such as the UTC time and NTP time difference of the server, Kerberos	2022-12-06	Cluster Event

related metrics.	response time, number of HDFS MissingBlocks, and safe mode for the HDFS NameNode.		
Added more default alarm policies for Kudu.	Kudu adds more alarm policies, such as the number of teservers, number of failed data directories, and number of full data directories.	2022-12-06	-
Upgraded the configuration management feature.	The configuration status is added, and you can view the details of expired configuration and failed configuration. More configuration files, configuration items, configuration item descriptions, and configuration filters are added.	2022-12-06	Configuration Management
Upgraded the YARN resource scheduling feature.	Label-based scheduling, queue deletion, and scheduling history viewing are supported.	2022-12-06	Capacity Scheduler
Added client management.	The client management entry is added, and you can view the client information.	2022-12-06	
Supported cloud disk adjustment.	You can adjust the cloud data disk models on nodes and expand and mount disks in the EMR console. Heterogeneous capacity is supported by the heterogeneous capabilities of cloud disks.	2022-12-06	
Supported batch configuration of automatic renewal	You can batch enable/disable automatic renewal for nodes in the EMR console.	2022-12-06	
Supported batch configuration adjustment.	The configurations of nodes in the same billing mode and the same region can be batch adjusted.	2022-12-06	Adjusting Configuration

November 2022

Update	Description	Release Date	Documentation
Added more StarRocks monitoring metrics.	More basic monitoring metrics are added for various StarRocks roles.	2022-11-03	

Supported Remote Shuffle Service (RSS) clusters for container-based EMR.	Container-based EMR supports RSS deployment in EKS. Spark clusters can be associated with RSS clusters.	2022-11-01	Overview
Supported managed scaling for automatic scaling.	The managed scaling feature continuously monitors the load system of the YARN cluster and automatically adds or removes compute nodes.	2022-11-01	

October 2022

Update	Description	Release Date	Documentation
Released a Hadoop cluster version of EMR v3.5.0	The Hadoop cluster version of EMR v3.5.0 is released to support Hive v3.1.3, Spark v3.2.2, HBase v2.4.5, Flink v1.14.5, Trino v389, Iceberg v0.13.1, Delta Lake v2.0.0, Hudi v0.12.0, Zeppelin v0.10.1, Superset v1.5.1, and Ranger v2.3.0.	2022-10-27	Component Version

September 2022

Update	Description	Release Date	Documentation
Supported Impala percentile distribution statistics.	Impala query supports statistics collection for execution status distribution and multidimensional percentile distribution. More metrics are added for Impala queries, such as total/HDFS/Kudu scanned rows, memory peak, read/sent bytes, and internal total scanned/sent bytes. Overview and structured profile features are provided for queries that take more than 3 seconds to execute.	2022-09-21	Impala Query Management
Added more alarm policies for monitoring events.	100+ monitoring metrics are added for Druid in the "Elastic MapReduce (dynamic)" category in CM, for which dynamic alarm policies can be configured.	2022-09-21	-

	20 default monitoring metrics for eight services are added for default alarm policies. Four Impala query events support alarm configuration through EventBridge.		
Launched application comparison and application insights features for YARN job management.	You can view task information for TEZ jobs. Application-level insights and comparison are supported for Tez/Spark/MapReduce jobs. Task-level host monitoring metric comparison is supported for Tez/MapReduce jobs. The experience to use YARN task information and filter task time is optimized.	2022-09-09	YARN Job Query

August 2022

Update	Description	Release Date	Documentation
Released a StarRocks cluster version: StarRocks v1.1.0.	Added support for StarRocks v2.2.2.	2022-08-04	Component Version
Released a Druid cluster version: Druid v1.1.0.	Added support for Druid v0.23.0.	2022-08-04	Component Version

July 2022

Update	Description	Release Date	Documentation
Launched the overview page in the console.	The cluster overview page displays various information, such as cluster status, health status of running services, and event statistics.	2022-07-22	-
Upgraded the cluster overview feature.	The cluster directory navigation is optimized and upgraded. The health status information of deployed services is refined to display real-time statistics of cluster events.	2022-07-22	Cluster Overview
Added new Impala query dimension and inspection items.	Dimensions of session ID, coordinator, and single-point memory peak are added. Four inspection items are added for Impala: query timeouts, total query failures, total query commits, and	2022-07-22	Impala Query Management

	query execution failure rate.		
Added the Impala daemon role.	Over 70 visual Impala daemon metrics are added, which can trigger alarms in CM.	2022-07-22	Impala Monitoring Metrics
Added new application APIs.	The API for collecting YARN application queue statistics is added.	2022-07-22	Query Application Statistics
Released a Hadoop cluster version—EMR v2.7.0.	Added support for Hive 2.3.9, Spark 3.2.1, HBase 2.4.5, Flink 1.14.3, Trino 385, ZooKeeper 3.6.3, Iceberg 0.13.0, Hudi 0.11.0, Alluxio 2.8.0, Zeppelin 0.10.1, Superset 1.4.1, and Ranger 2.1.0.	2022-07-18	Component Version
Improved the availability of the automatic scaling feature.	Scale-out rules support the hybrid deployment of multiple resource specifications as well as pay-as-you-go/spot instances.	2022-07-07	Auto Scaling
Added support for the RIT fixing tool for HBase.	Regions in the RIT status can be fixed.	2022-07-07	HBase RIT Fixing
Started the open beta test for container-based EMR clusters.	Container-based EMR clusters become available in Beijing, Shanghai, and Guangzhou regions.	2022-07-07	Container-Based EMR

June 2022

Update	Description	Release Date	Documentation
Optimized the configuration management feature.	Configuration management is upgraded to support configuration filtering, categorization, and comparison. The delivery priority rules for configurations at different levels are optimized to support delivery at the minimum level first. A configuration group can be specified for the selected components.	2022-06-20	Configuration ManagementCluster Scale-Out
Added support for Hive query	The Hive service supports query management.	2022-06-07	Hive Query Management

management.			
Optimized the YARN job query.	YARN job query supports displaying task information and querying task logs.	2022-06-07	YARN Job Query
Connected to Pricing Center.	The pricing of EMR is connected to the Pricing Center, where you can query EMR node fees.	2022-06-07	Pricing

May 2022

Update	Description	Release Date	Documentation
Optimized the scaling feature.	New subnets can be added to expand cluster node capacity. The Kyuubi service supports scaling.	2022-05-17	Cluster Scale-Out
Enhanced role management and user management.	User management supports batch import and export of users through APIs. Resource management supports displaying client information. The `DescribeClusterPhysicalMetaInfo` API is added for resource-level authentication.	2022-05-17	-

April 2022

Update	Description	Release Date	Documentation
Released the EMR v3.4.0 for the Hadoop cluster.	Added support for Spark 3.2.1, Trino 372, Flink 1.14.3, HBase 2.4.5, ZooKeeper 3.6.3, Impala 4.0.0, Zeppelin 0.10.1, Alluxio 2.8.0, Hudi 0.10.1, Iceberg 0.13.1, and GooseFS 1.2.0. Renamed PrestoSQL as Trino in the new version and added support for Ranger. Monitoring metric alarms need to be configured according to Trino in the new version.	2022-04-22	Component Version Trino Monitoring Metrics
Enriched monitoring.	Launched the new version of the interactive features of the event list. Added event-related metrics such as node single-disk inode utilization, single-disk space utilization, single-disk IO	2022-04-22	Cluster Event HDFS Monitoring Metrics

	utilization, all HBase HMaster standby, and all HDFS NameNodes standby. Added monitoring metrics of the HDFS-Data node, including the block count, used disk space, available disk space, and reserved disk space.		
Added support for container-based EMR clusters.	Added support for Spark deployment based on EKS in Beijing, Shanghai, Guangzhou, and Nanjing regions.	2022-04-12	Overview

March 2022

Update	Description	Release Date	Documentation
Added the StarRocks cluster type.	Added the StarRocks cluster type and integrated open-source StarRocks.	2022-03-11	StarRocks Development Guide
Added support for adding nodes without assigning ApplicationMasters by default.	Added support for adding nodes without assigning ApplicationMasters by default, ensuring that the running ApplicationMaster will not be terminated and the job can be performed normally.	2022-03-08	Automatically Adding Task Nodes Without Assigning ApplicationMasters
Added load-based scaling metrics.	Added load-based scaling metrics such as ContainerPending, AppPending, and AppRunning.	2022-03-08	Auto Scaling
Added support for setting node labels for added host resources.	Added support for setting YARN node labels for added host resources to increase the resource utilization efficiency.	2022-03-08	-

January 2022

Update	Description	Release Date	Documentation
Released a Doris cluster	Released Doris v0.15.	2022-01-21	Component Version

version— Doris v1.2.0.			
Added support for sharing component configurations.	Added support for sharing the components of an existing cluster with other clusters without deploying such components again. This makes it easier to manage multiple clusters with the same component configurations.	January 14, 2022	Component Configuration Sharing
Added monitoring metrics.	Added Kylin, Zeppelin, Oozie, Storm, Livy, and Kyuubi monitoring metrics. Added new CosRanger monitoring metrics.	-	Kylin Monitoring Metrics Zeppelin Monitoring Metrics Oozie Monitoring Metrics Storm Monitoring Metrics Livy Monitoring Metrics Kyuubi Monitoring Metrics CosRanger Monitoring Metrics
Optimized log search.	Added support for more components and global regions and optimized log search interactions.	-	Log Search

December 2021

Update	Description	Release Date	Documentation
Added four use cases for the Hadoop cluster type.	Added the ZooKeeper, HBase, Presto, and Kudu use cases for the Hadoop cluster type to help you quickly create a cluster based on your business scenario.	2021-12-17	Cluster Types
Launched a new purchase page.	Optimized the existing cluster creation process to support the placement group node dimension for unified resource management.	2021-12-17	Creating EMR Cluster
Added support for system disk selection.	Added support for selecting the system disk type and size as needed.	2021-12-17	-

Optimized the task center.	Added support for viewing the details of all tasks of the current cluster inside the cluster, and optimized the interaction path of the task center for easier viewing and use.	2021-12-17	-
Launched diversified component operation features.	Added support for visually setting instruction-level operations of components such as HDFS and YARN.	2021-12-07	Service Operation
Added monitoring metrics.	Added Doris monitoring metrics.Adjusted the original HBase table load to HBase table-level monitoring.	2021-12-07	Doris Monitoring MetricsHBase Table-Level Monitoring
Upgraded the operating system.	Upgraded the operating systems of all EMR versions of the Hadoop type to TencentOS Server to provide a more stable, secure, and high-performance cloud native runtime environment.If you encounter any security issues related to the operating system of an existing cluster, contact us for assistance.	2021-12-03	Product Overview of TencentOS Server

September 2021

Update	Description	Release Date	Documentation
Released a Hadoop cluster version —EMR v3.3.0.	Added support for HBase v2.3.5, Phoenix v5.1.2, Kudu v1.15.0, Hue v4.10.0, Hudi v0.8.0, Tez v0.10.1, Livy v0.8.0, Ganglia v3.7.2, Kyuubi v1.1.0, and LDAP v2.4.44.	September 28, 2021	Component Version
Added monitoring metrics.	Enriched Alluxio service monitoring metrics and added COSRanger monitoring metrics.	September 15, 2021	Alluxio Monitoring MetricsCOSRanger Monitoring Metrics
Connected to Cloud Monitor's default alarms	Cloud Monitor supports default alarm policies for Elastic MapReduce metrics/events.	September 15, 2021	-

August 2021

Update	Description	Release Date	Documentation
Added support for YARN resource scheduling.	Added support for UI-based YARN resource scheduling configuration.	August 31, 2021	YARN Resource Scheduling
Added support for user management.	Added support for visually managing Kerberos and Ranger users.	August 31, 2021	Managing Users

July 2021

Update	Description	Release Date	Documentation
Released a Hadoop cluster version—EMR v2.6.0.	Added support for Spark v3.0.2, Impala v3.4.0, GooseFS v1.0.0, OpenLDAP v2.4.44, Kyuubi v1.1.0, etc.	July 15, 2021	Component Version

June 2021

Update	Description	Release Date	Documentation
Added support for YARN job query.	Added detailed metric statistics for YARN jobs.	June 24, 2021	YARN Job Query
Added support for Impala query management.	Added detailed metric statistics for Impala query jobs.	June 24, 2021	Impala Query Management
Added support for automatic scaling for the spot instance mode.	Added support for automatically adding compute nodes of the spot instance mode.	June 3, 2021	Auto Scaling
Added support for graceful scale-in.	When a scale-in action is triggered, nodes that are executing tasks will not be released immediately. Instead, they will be released after completing the tasks.	June 3, 2021	Graceful Scale-In
Added support for	Added support for viewing Tez tasks via UI.	June 3,	Integrating

integrating Knox with Tez.		2021	Knox with Tez
Added availability zones.	Added support for Singapore Zone 3 and Tokyo Zone 2.	June 3, 2021	-
Added regions.	Added support for regions such as Moscow, Frankfurt, Virginia, Bangkok, Mumbai, Tokyo, Toronto, and Seoul.	June 3, 2021	-

May 2021

Update	Description	Release Date	Documentation
Released a Kafka cluster version—Kafka v1.0.0.	Added support for Kafka v1.1.1, KafkaManager v2.0.0.2, Knox v1.2.0, and ZooKeeper v3.6.1.	May 13, 2021	Kafka Development Guide
Released a Doris cluster version—Doris v1.0.0.	Added support for Doris v0.13.0 and Knox v1.2.0.	May 13, 2021	Doris Development Guide

April 2021

Update	Description	Release Date	Documentation
Released a Hadoop cluster version—EMR v3.2.0.	Added support for Hadoop v3.2.2, Spark v3.0.2, Presto v350, Flink v1.12.1, Alluxio v2.5.0, Ranger v2.1.0, Iceberg v0.11.0, and Hudi v0.7.0.	April 30, 2021	Component Version
Added the **Start/Stop Services** feature.	Added support for starting or stopping all services on a node.	April 27, 2021	Starting/Stopping Services
Released a Hadoop cluster version—EMR v2.5.1.	Added support for Alluxio v2.5.0 and Hudi v0.5.1. Alluxio supports COS Transparent-URI.	April 10, 2021	Support for COS Transparent-URI

March 2021

Update	Description	Release Date	Documentation
Supported cluster restoration	You can renew monthly subscribed clusters in isolated status to restore them.	2021-03-17	

February 2021

Update	Description	Release Date	Documentation
Released the EMR service in new regions on Tencent Cloud Chinese.	Added support for the Bangkok, Seoul, Tokyo, Virginia, and Toronto regions on Tencent Cloud Chinese.	February 2, 2021	-

January 2021

Update	Description	Release Date	Documentation
Added the cluster script feature.	Added support for running a specified script on multiple nodes at a time for higher efficiency.	January 26, 2021	Cluster Scripts
Added support for cross-model configuration change.	You can select other models to change the node configuration when the current model is sold out.	January 26, 2021	Adjusting Configuration
Added the auto scaling feature.	The auto scaling feature automatically adds or removes task nodes, helping you save costs while meeting the requirement for cluster computing power.	January 26, 2021	Auto Scaling
Added monitoring metrics.	Added new monitoring metrics for the HDFS, Yarn, HBase, Hive, and Druid services. Added support for monitoring on the Kudu and Alluxio services.	January 13, 2021	Node Monitoring Metrics
Released the EMR	Added support for Singapore Zone 1 and Singapore	January	-

service in Singapore on the international version.	Zone 2 on the international version.	13, 2021	
--	--------------------------------------	----------	--

December 2020

Update	Description	Release Date	Documentation
Released a Hadoop cluster version—EMR v3.1.0.	Added support for Impala v3.4.0, Flink v1.10.0, HBase v2.3.3, Phoenix v5.0.0, Alluxio v2.3.0, Kudu v1.13.0, and ZooKeeper v3.6.1.	December 14, 2020	Component Version

November 2020

Update	Description	Release Date	Documentation
Released the EMR TianQiong edition.	Added support for Spark materialized views.	November 27, 2020	Component VersionEMR TianQiong Introduction
Optimized monitoring and alarming.	Added support for setting alarms for disk monitoring metrics in Cloud Monitor (CM). Added support for customizing EMR monitoring dashboard in CM dashboard. Added support for subscribing to EMR event monitoring alarms in CM.	November 20, 2020	Dashboard

October 2020

Update	Description	Release Date	Documentation
Added support for Cloud Object Storage (COS) password-free access.	Added support for authorizing access to COS through roles. Authorized users do not need to enter an authorization key and private key.	2020-10-26	-
Added new service	Added the HDFS, Yarn, HBase, Hive, and	2020-	Cluster

inspection items to the inspection system.	ZooKeeper service inspection items, enhancing the health check of metrics.	10-13	Inspection
Added support for customizing event policies.	Added support for customizing event policies so you can identify cluster events more efficiently.	2020-10-13	Cluster Event

September 2020

Update	Description	Release Date	Documentation
Released a Hadoop cluster version—EMR v2.5.0.	Upgraded the Alluxio component to the stable version 2.3.0, which delivers better performance in the use cases where compute and storage are separated.	September 17, 2020	Component Version

August 2020

Update	Description	Release Date	Documentation
Released a Hadoop cluster version—EMR v2.4.0.	Added the Kudu component and upgraded some existing components.	August 27, 2020	Component Version
Added support for data migration in ClickHouse clusters.	Improved cluster resource utilization for data security.	August 17, 2020	ClickHouse Data Migration Guide

July 2020

Update	Description	Release Date	Documentation
Added support for custom configuration files.	Added support for custom configuration files, which improves the flexibility of configuration management.	July 29, 2020	Configuration Management

Optimized configuration item editing.	Added descriptions for configuration items and support for restoring them to defaults.	2020-07-29	-
---------------------------------------	--	------------	---

June 2020

Update	Description	Release Date	Documentation
Added support for Druid service monitoring and HBase Thrift role monitoring.	Improved troubleshooting efficiency.	June 15, 2020	Druid Monitoring Metrics
Released a Hadoop cluster version—EMR v2.3.0.	Added support for the TensorFlow and Jupyter Notebook components.	June 15, 2020	TensorFlow Overview Jupyter Notebook Overview
Released ClickHouse v1.1.0 for ClickHouse clusters.	Upgraded the ClickHouse component version to the LTS version, and added support for the Superset component.	June 15, 2020	-
Added support for cluster events and cluster inspections.	Added support for recording major change events and exception events in clusters in the console for improved efficiency of cluster troubleshooting. Added support for one-time inspections and periodical inspections, making it easier to keep track of the health of your clusters and to deal with exceptions and risks in time.	June 4, 2020	Cluster Event Cluster Inspection

May 2020

Update	Description	Release Date	Documentation
Optimized console interactions,	Improved operation path hierarchy for	May 15,	-

UI, and operation paths.	easier service management.	2020	
Added support for service monitoring for ClickHouse clusters.	Made it easier to keep track of the operation status of ClickHouse clusters.	May 15, 2020	ClickHouse Monitoring Metrics
Added support for monitoring status for role management.	Made it easier to keep track of the real-time status of role processes.	May 15, 2020	-

April 2020

Update	Description	Release Date	Documentation
Added new cluster types.	Added support for the ClickHouse and Druid clusters to provide more real-time query solutions.	April 15, 2020	ClickHouse Overview Druid Overview
Optimized bootstrap actions.	Added support for the running occasions of bootstrap actions before a cluster starts. Added support for adjusting bootstrap actions after a cluster is created, so you can change your bootstrap actions based on business changes.	April 15, 2020	-
Optimized service restart policies.	Added support for setting service restart policies, so you can select a proper restart method and a policy for handling restart exceptions.	2020-04-15	-

March 2020

Update	Description	Release Date	Documentation
Optimized the cluster monitoring page.	Optimized the cluster overview, service monitoring, and node monitoring pages for improved readability. Monitoring metrics now support customization and heat maps about nodes.	March 2020	-
Added support for log search.	Added support for filtering by log file, node IP, and time range, and viewing cluster log content. (Currently, this feature is supported only in Guangzhou region.)	2020-03	Log Search

Released EMR v2.2.0.	Upgraded Hadoop to v2.8.5. Upgraded common components such as Spark, Hive, and HBase. Added the Hudi, Superset, Livy, Impala, Zeppelin, and Kylin components.	2020-03	Component Version
----------------------	---	---------	-----------------------------------

February 2020

Update	Description	Release Date	Documentation
Added support for sharing Hive metadatabase.	Added support for associating an external Hive metadatabase to store the Hive metadata of a new cluster, so multiple Hive clusters can share metadata.	February 2020	Hive Metadata Management
Optimized cluster operation logs.	Added display items for cluster operation logs to increase the readability of operation logs.	February 2020	-
Released the international version.	Released the international version in Beijing, Shanghai, Guangzhou, and Mumbai regions.	2020-02	-

January 2020

Update	Description	Release Date	Documentation
Added support for cloud disk encryption.	Added support for cloud disk encryption for Cloud Block Storage (CBS) users on the allowlist when creating an EMR cluster.	January 2020	Cloud Disk Encryption
Added support for mounting cloud disks to a task node.	Added support for mounting cloud disks to a task node when creating or scaling out a cluster.	January 2020	-
Released EMR in Nanjing region.	Released EMR in Nanjing region.	2020-01	-

December 2019

Update	Description	Release	Documentation
--------	-------------	---------	---------------

		Date	
Released the feature of delivering group configuration.	Added support for delivering group configuration. If the configuration of a node is different from that of the configuration group to which the node belongs, the configuration of the configuration group can be delivered to the node.	December 2019	-
Released the feature of resource-level CAM authorization.	Added support for CAM authorization at the resource level.	December 2019	-
Released the feature of associating CHDFS.	Added support for associating CHDFS and reading/writing data on CHDFS.	December 2019	Mounting CHDFS
Released the AMD model.	Released the AMD Standard SA2 model in Beijing, Shanghai, and Guangzhou regions.	2019-12	Instance Types
Added the cluster monitoring overview page.	Added the cluster monitoring overview page, which provides the views about the cluster, node, and service status.	December 2019	Cluster Overview
Added support for selecting metric display granularity as needed on the service monitoring page.	Optimized the service monitoring page, so you can select a metric display granularity as needed.	December 2019	-
Added support for node service deployment and load status views on the node monitoring page.	Optimized the node monitoring page, which provides the views about node service deployment and load status.	December 2019	-

November 2019

Update	Description	Release Date	Documentation
Added support for node specification management.	Added support for setting the default specification of nodes based on their billing mode. The default specification is pay-as-you-go.	2019-11	Node Specification Management

Added support for the configuration of the ZooKeeper, Alluxio, and Flink components.	Added support for the configuration of the ZooKeeper, Alluxio, and Flink components.	November 2019	-
Released the feature of tagging clusters and nodes.	Added support for tagging clusters and nodes.	November 2019	Setting Tag
Released new models in Beijing, Shanghai, and Guangzhou regions.	Released the S5, M5, C3, and CN3 models in Beijing, Shanghai, and Guangzhou regions.	2019-11	-

October 2019

Update	Description	Release Date	Documentation
Released EMR v3.0.0.	Released EMR v3.0.0 with updated versions of major components.	October 2019	Component Version

September 2019

Update	Description	Release Date	Documentation
Added full support for TencentCloud API 3.0.	Added full support for TencentCloud API 3.0. Standardized the output/input parameters of certain existing v3.0 APIs. Added full support for v2.0 APIs in v3.0.	September 2019	-
Modified console configuration.	Disused the escape feature for special characters.	September 2019	-
Added support for CM alarm policy configuration.	Added support for configuring alarm policies in CM (in the Elastic MapReduce product category) for key metrics for node and service monitoring.	September 2019	-

August 2019

--	--	--	--

Update	Description	Release Date	Documentation
Released the software configuration feature.	Added support for software configuration, which enables you to create a cluster with custom component parameters. The feature of accessing external clusters was added as well.	August 2019	Software Configuration
Added support for configuring the remote login port when you purchase a cluster.	Added support for enabling or disabling the remote login port when you purchase a cluster.	August 2019	-
Added support for mounting multiple cloud disks to a new cluster.	Added support for mounting multiple cloud disks to a new cluster.	August 2019	-
Released the feature of scaling specified components.	Added support for scaling specified components.	August 2019	-
Added monitoring metrics.	Added new monitoring metrics for the Spark, Hive, Presto, and ZooKeeper components.	August 2019	-

July 2019

Update	Description	Release Date	Documentation
Released the feature of HBase table-level monitoring.	Added support for HBase table-level monitoring, which covers the number of read and write requests and storage of each table in HBase.	July 2019	Service Monitoring
Released EMR in new regions.	Added support for the Singapore, Mumbai, and Chengdu regions.	July 2019	-
Released new models.	Released the Standard S4 and Standard Network-Optimized SN3ne models in Beijing, Shanghai, and Guangzhou regions.	2019-07	-
Optimized the interactions of monitoring pages and added new monitoring metrics.	Optimized the interactions of monitoring pages and added new monitoring metrics.	July 2019	-

Optimized WebUI proxy address.	Optimized WebUI proxy address.	July 2019	-
--------------------------------	--------------------------------	-----------	---

June 2019

Update	Description	Release Date	Documentation
Released the feature of accessing external clusters.	Added support for accessing external clusters.	June 2019	-
Released the feature of setting bootstrap actions.	Added support for setting bootstrap actions.	June 2019	Bootstrap Actions
Released the feature of delivering component parameter configuration to nodes.	Added support for delivering component parameter configuration to nodes.	June 2019	-
Released the feature of rolling back parameter configuration.	Added support for rolling back parameter configuration.	June 2019	Configuration Rollback

May 2019

Update	Description	Release Date	Documentation
Released EMR v2.1.0.	Released EMR v2.1.0 with updated versions of main components.	May 2019	Component Version
Added support for Kerberos security clusters.	Added support for creating security clusters. The open source components in clusters are launched in Kerberos security mode. In this security environment, only authenticated clients can access the services (such as HDFS) of the clusters.	May 2019	Kerberos Introduction
Optimized monitoring metrics.	Optimized the Node, HDFS, YARN, and HBase monitoring metrics.	May 2019	Node Monitoring Metrics
Optimized the style of console.	Improved the style of console for a better interactive experience.	May 2019	-

Optimized CVM and TencentDB naming.	Optimized CVM and TencentDB naming with the EMR cluster serial number for easier locating of cluster information.	May 2019	-
Changed the setting of public IPs of master nodes to optional.	Changed the setting of public IPs of master nodes to optional.	May 2019	-
Released the feature of adjusting the number of common nodes.	Added support for adjusting the number of common nodes as needed.	May 2019	-

March 2019

Update	Description	Release Date	Documentation
Released new models.	Released the I3 model in Beijing, Shanghai, and Guangzhou regions. This model is included in the CVM allowlist, and you can purchase it only if you are in the I3 allowlist.	March 2019	-
Released EMR in a new region.	Made the purchase of EMR clusters available in the Silicon Valley region.	March 2019	-
Added support for router nodes.	Added support for router nodes, which are mainly used to relieve the load of master nodes and as task submitters.	March 2019	Node Type Description
Released the feature of adjusting node configuration.	Added support for adjusting node configuration. Nodes can be upgraded to a higher configuration.	March 2019	Adjusting Configuration

January 2019

Update	Description	Release Date	Documentation

Added new components to existing clusters.	Added support for adding new components to existing clusters.	January 2019	Adding Components
--	---	--------------	-----------------------------------

Announcements

Alarm Policy Migration

Last updated : 2022-11-25 16:06:46

Background

Elastic MapReduce upgraded its server and component service monitoring items on September 10, 2019. A new policy type "Elastic MapReduce" has been added, which covers hundreds of monitoring metrics. You can configure alarm policies under the [Elastic MapReduce](#) policy type in Cloud Monitor. **The original policy type "EMR" was deprecated at 23:00 on March 30, 2021, and all the configured "EMR" alarm policies were invalidated.** To add new alarm policies, configure under the "Elastic MapReduce" policy type.

Comparison of "EMR" and "Elastic MapReduce" policy types:

Policy Type	Metric Coverage	Support and Maintenance
EMR	Cluster alarming (12 metrics) Node alarming (8 metrics)	It was deprecated at 23:00 on April 9, 2021 and no longer maintained.
Elastic MapReduce	Server monitoring Service monitoring Cluster monitoring	It was released on September 10, 2019 but is still maintained.

Note:

The "Elastic MapReduce" policy type covers all the metrics of the "EMR" policy type. For details, see [Comparison of new and original metrics](#).

Alarm policy transfer

After the "EMR" policy type was deprecated, the system automatically migrated existing alarm policies under the "EMR" policy type to the "Elastic MapReduce" policy type. See further notifications for specific rules and verification methods.

Note:

A few users might need to migrate them manually.

Steps for manual migration are as follows:

1. Sort out exiting alarm metrics and policies.

Log in to the [Cloud Monitor console](#), select **Alarm Management > Alarm Configuration > Alarm Policy** on the left

sidebar, click **Advanced Filter**, select any sub-type of the "EMR" policy type in **Policy Type**, search for related alarm policies, and download them. Repeat this step to download the alarm policies under other sub-types of the "EMR" policy type.

Alarm Policy
[New Version Guide](#)

Starting November 15, 2022, the alarm policy type "TencentDB/Redis/Memory Edition (1-Minute Granularity)" will be disused and automatic Edition (5-Second Granularity)" in a gradual manner, so that you can enjoy more accurate and timely monitoring and alarming services. You can view the details in the Message Center.

Create
Delete
More ▾
Advanced Filter
Separate keywords with "|"

Monitoring Type	Select monitoring type ▾	Policy Type	Select a monitoring type first ▾	Alarm
User Group	User ▾ Please select ▾	Default Policy Filter	Please select ▾	Detec

Query

2. Configure alarm policies.

Click **Alarm Configuration** in the cluster list to go to the **Alarm Policy** page, click **Create**, select **Elastic MapReduce** in **Policy Type**, and configure an alarm policy according to one of the policies sorted out and downloaded in step 1. Repeat this step to configure other alarm policies. For more information on the configuration method, see [Alarm Configurations](#).

3. Verify the new alarm policies.

Verify that the "Elastic MapReduce" alarm policies are activated and can successfully trigger alarms. Set trigger thresholds in **Metric Alarm**, set recipient groups or recipients, and select receiving channels (email, SMS, or WeChat) for verification. Take `memory zone percentage_50` as an example: when the statistical period is five minutes, and the percentage is greater than or equal to 1% and lasts for five minutes, an alarm will be triggered once every five minutes.

4. Delete original alarm policies.

After verifying the new alarm policies, delete the original alarm policies configured under the "EMR" policy type. Select any sub-type of the "EMR" policy type in **Policy Type**, search for related alarm policies, and delete them. Repeat this step to delete the alarm policies under other sub-types of the "EMR" policy type.

If you encounter any issues during the transfer, [contact us](#) for assistance.

Comparison of new and original metrics

Original Policy Type	Metric/Event Alarm	Original Metric/Event Name	New Policy Type	New Metric/Event Name
EMR-cluster alarm	Metric alarm	Used HDFS storage space	Elastic MapReduce-HDFS-overview	Cluster storage capacity_CapacityUsed
	Metric alarm	HDFS storage utilization	Elastic MapReduce-HDFS-overview	HDFS storage space utilization_capacityused
	Metric alarm	YARN app blocks	Elastic MapReduce-YARN-overview	Applications_pending
	Metric alarm	Failed YARN apps	Elastic MapReduce-YARN-overview	Applications_failed
	Metric alarm	Assigned cluster CPU cores	Elastic MapReduce-YARN-overview	Cores_allocatedVirtualCores
	Metric alarm	Cluster CPU utilization	Elastic MapReduce-YARN-overview	CPU utilization_usageRatio
	Metric alarm	Available cluster memory	Elastic MapReduce-YARN-overview	Memory_availableMB
	Metric alarm	Cluster memory utilization	Elastic MapReduce-YARN-overview	Memory utilization_usageRatio
	Metric alarm	Cluster container blocks	Elastic MapReduce-YARN-overview	Containers_containersPending
	Metric alarm	HBase requests	Elastic MapReduce-HBASE-overview	Total cluster requests_clusterRequests
	Metric alarm	Time taken for HBase sync	No longer maintained	-
	Metric alarm	HBase sync log length	No longer maintained	-
	Event alarm	Node monitoring heartbeat missing	Elastic MapReduce-server monitoring-network	Node monitoring heartbeat missing
	Event alarm	Process restart	Elastic MapReduce-server monitoring-process	Process OOM

EMR-node alarm	Metric alarm	Disk utilization	Elastic MapReduce-server monitoring-disk	Disk capacity utilization_used_all
	Metric alarm	Memory usage	Elastic MapReduce-server monitoring-memory	Memory usage_MemFree
	Metric alarm	Server restart	No longer maintained	Server restart
	Metric alarm	Memory utilization	Elastic MapReduce-server monitoring-memory	Memory utilization_used_percent
	Metric alarm	CPU usage	Elastic MapReduce-server monitoring-CPU	CPU utilization_idle
	Metric alarm	Inbound packets over private networks	No longer maintained	-
	Metric alarm	Outbound packets over private networks	No longer maintained	-
	Metric alarm	TCP connections	Elastic MapReduce-server monitoring-network	TCP connections

Security Announcements

Notice for Apache Log4j 2 RCE Vulnerability

Last updated : 2022-05-16 12:15:40

Vulnerability Description

Recently, Tencent Cloud has noticed that the remote code execution vulnerability in Apache Log4j 2 has been disclosed. There is a JNDI injection vulnerability in Log4j 2, which can be triggered when the program logs the user-input data. It can be exploited to run any code on the target server. For more information on the vulnerability, see [here](#).

Impact

Components in EMR such as Flink, Hive, Ranger, Presto, Oozie, Knox, Storm, and Druid are affected by this vulnerability. If you are affected, fix it as instructed below.

Solution

Replace the Log4j 2 package with a safe version.

Affected versions: Apache Log4j2 2.0–2.15.0-rc1.

Safe versions: Apache Log4j 2.17.1.

Fix Command

1. Run the fix command in the standard EMR directory:




```
wget https://image-repo-gz-1259353343.cos.ap-guangzhou.myqcloud.com/user-patches/log4j2.sh -O fix-log4j2.sh && bash -x fix-log4j2.sh /usr/local/service
```

2. Fix the JAR packages in the cache directory when running the task.

Make sure that there are no problematic JAR packages in the submitted task; otherwise, they will be cached again in the task submitted next time.


Directly delete the problematic JAR packages in the directory.



```
/data/emr/yarn/local/filecache/  
/data/emr/yarn/local/usercache/  
/data1/emr/yarn/local/filecache/  
/data1/emr/yarn/local/usercache/  
/data2/emr/yarn/local/filecache/  
/data2/emr/yarn/local/usercache/
```

The above lists the information of only three data disks, where the number following `/data` is the data disk index. You need to clear the corresponding files in the `/data` directory of all data disks.

3. Run the following command to fix non-standard directories (directories other than `/usr/local/service`).



```
EXTRA_DISRUPTOR_DIR=/path/to/other bash fix-log4j2.sh /path/to/other
```

4. Fix in other scenarios.

Upgrade the six JAR packages related to the vulnerability: log4j-api, log4j-core, log4j-jul, log4j-slf4j-impl, log4j-web, and disruptor. If such a package doesn't exist, you don't need to replace it.

Service Restart and Grayscale Fix

1. Perform a fix on a node in the cluster.

Restart the Flink, Spark, Hive, Ranger, Presto, Oozie, Storm, Impala, Knox, and Druid services on the node.

Restart all resident tasks, including Flink, Storm, and Spark tasks.

2. After confirming that everything is OK by restarting the services on the node, perform the fix on other nodes.

Fix Process

1. Place the six fixed JAR packages in the `fix-log4j` directory of the execution directory.

2. Search for problematic packages in the target directory and replace the found ones with the six fixed JAR packages. The problematic JAR packages in the `tar.gz` and `war` packages, as well as the cache package in the `/user/hadoop/share` path on HDFS will be replaced at the same time.

Replace `log4j-api, log4j-core, log4j-jul, log4j-slf4j-impl, log4j-web 2.0-2.17.1` with the 2.17.1 version.

Replace the versions earlier than `disruptor-3.4.2.jar` with the 3.4.2 version. Note that the disruptor needs to be replaced only for certain components.

Rollback Steps for Service Problems

You need to copy the problematic JAR packages back and delete the added latest JAR packages.

1. Decompress the backup file.



```
cd fix-log4j2
tar zxvf rm_if_no_need_to_rollback.tar.gz.1639576622
```

2. Copy the backup file back.



```
cp -r ./root/fix-log4j2/emr_fix_log4j_bak_10812_1639576622/usr/local/service/* /us
```

3. Delete the added latest Log4j JAR packages.



```
find /usr/local/service/ -name log4j-api-2.17.1.jar | xargs -n1 -I{} rm -f {}  
find /usr/local/service/ -name log4j-web-2.17.1.jar | xargs -n1 -I{} rm -f {}  
find /usr/local/service/ -name log4j-jul-2.17.1.jar | xargs -n1 -I{} rm -f {}  
find /usr/local/service/ -name log4j-slf4j-impl-2.17.1.jar | xargs -n1 -I{} rm -f {}  
find /usr/local/service/ -name log4j-core-2.17.1.jar | xargs -n1 -I{} rm -f {}
```

4. To roll back in the `/path/to/other` directory, replace `/usr/local/service/` with `/path/to/other`.