

Elastic MapReduce

Release Notes and Announcements

Product Documentation



Copyright Notice

©2013-2022 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

 - Notice for Apache Log4j 2 RCE Vulnerability

 - Alarm Policy Migration

Release Notes and Announcements

Release Notes

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

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 | StarRocks Monitoring Metrics |

October 2022

| Update | Description | Release Date | Documentation |
|--|---|--------------|---|
| Released a Hadoop cluster version: EMR v3.5.0. | Added support for 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 | Product Releases and Component Versions |

September 2022

| Update | Description | Release Date | Documentation |
|--------|-------------|--------------|---------------|
|--------|-------------|--------------|---------------|

| Update | Description | Release Date | Documentation |
|--|---|--------------|---|
| Supported Impala percentile distribution statistics. | <ul style="list-style-type: none"> 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. | <ul style="list-style-type: none"> 100+ monitoring metrics are added for Druid in the "Elastic MapReduce (dynamic)" category in CM, for which dynamic alarm policies can be configured. 20 default monitoring metrics for eight services are added for default alarm policies. Four Impala query events support alarm configuration through EventBridge. | 2022-09-21 | - |
| Launched application comparison and application insights features for YARN job management. | <ul style="list-style-type: none"> 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 | Product Releases and Component Versions |

| Update | Description | Release Date | Documentation |
|---|----------------------------------|--------------|---|
| Released a Druid cluster version: Druid v1.1.0. | Added support for Druid v0.23.0. | 2022-08-04 | Product Releases and Component Versions |

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. | <ul style="list-style-type: none"> 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. | <ul style="list-style-type: none"> 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 query execution failure rate. | 2022-07-22 | Impala Query Management |
| 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 | DescribeInstancesList |
| 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 | Product Releases and Component Versions |

| Update | Description | Release Date | Documentation |
|--|--|--------------|----------------------------------|
| 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 | Overview |

June 2022

| Update | Description | Release Date | Documentation |
|---|---|--------------|--|
| Optimized the configuration management feature. | <ul style="list-style-type: none"> 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 Management Cluster Scale-Out |
| Added support for Hive query management. | The Hive service supports query management. | 2022-06-07 | Hive Query 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 Elastic MapReduce |

May 2022

| Update | Description | Release Date | Documentation |
|---|--|--------------|-----------------------------------|
| Optimized the scaling feature. | <ul style="list-style-type: none"> 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. | <ul style="list-style-type: none"> User management supports batch import and export of users through APIs. Resource management supports displaying client information. The `DescribeClusterPhysicalMetalInfo` 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. | <ul style="list-style-type: none"> 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 | Product Releases and Component Versions Trino Monitoring Metrics |
| Enriched monitoring. | <ul style="list-style-type: none"> 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 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. | 2022-04-22 | Cluster Event HDFS Monitoring Metrics |

| Update | Description | Release Date | Documentation |
|---|---|--------------|--------------------------|
| 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 Overview |
| 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 version: Doris v1.2.0. | Released Doris v0.15. | 2022-01-21 | Product Releases and Component Versions |

| Update | Description | Release Date | Documentation |
|---|---|--------------|---|
| 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. | 2022-01-14 | 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 |

| Update | Description | Release Date | Documentation |
|--|--|--------------|--|
| 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 Metrics HBase 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 | TencentOS Server Introduction |

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. | 2021-09-28 | Version Overview |
| Added monitoring metrics. | Enriched Alluxio service monitoring metrics and added COSRanger monitoring metrics. | 2021-09-15 | Alluxio Monitoring Metrics COSRanger Monitoring Metrics |

| Update | Description | Release Date | Documentation |
|---|---|--------------|---------------|
| Connected to Cloud Monitor's default alarms | Cloud Monitor supports default alarm policies for Elastic MapReduce metrics/events. | 2021-09-15 | - |

August 2021

| Update | Description | Release Date | Documentation |
|---|--|--------------|--|
| Added support for YARN resource scheduling. | Added support for UI-based YARN resource scheduling configuration. | 2021-08-31 | YARN Resource Scheduling |
| Added support for user management. | Added support for visually managing Kerberos and Ranger users. | 2021-08-31 | 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. | 2021-07-15 | Version Overview |

June 2021

| Update | Description | Release Date | Documentation |
|-----------------------------------|---|--------------|--------------------------------|
| Added support for YARN job query. | Added detailed metric statistics for YARN jobs. | 2021-06-24 | YARN Job Query |

| Update | Description | Release Date | Documentation |
|---|--|--------------|---|
| Added support for Impala query management. | Added detailed metric statistics for Impala query jobs. | 2021-06-24 | 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. | 2021-06-03 | 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. | 2021-06-03 | Graceful Scale-In |
| Added support for integrating Knox with Tez. | Added support for viewing Tez tasks via UI. | 2021-06-03 | Integrating Knox with Tez |
| Added availability zones. | Added support for Singapore Zone 3 and Tokyo Zone 2. | 2021-06-03 | - |
| Added regions. | Added support for regions such as Moscow, Frankfurt, Virginia, Bangkok, Mumbai, Tokyo, Toronto, and Seoul. | 2021-06-03 | - |

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. | 2021-05-13 | Kafka Overview |
| Released a Doris cluster version: Doris v1.0.0. | Added support for Doris v0.13.0 and Knox v1.2.0. | 2021-05-13 | Doris Overview |

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. | 2021-04-30 | Version Overview |
| Added the **Start/Stop Services** feature. | Added support for starting or stopping all services on a node. | 2021-04-27 | Service Start/Stop |
| 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. | 2021-04-10 | Support for COS Transparent-URI |

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. | 2021-02-02 | - |

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. | 2021-01-26 | 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. | 2021-01-26 | Configuration Change |

| Update | Description | Release Date | Documentation |
|---|---|--------------|---|
| 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. | 2021-01-26 | Auto Scaling |
| Added monitoring metrics. | <ol style="list-style-type: none"> Added new monitoring metrics for the HDFS, Yarn, HBase, Hive, and Druid services. Added support for monitoring on the Kudu and Alluxio services. | 2021-01-13 | Node Monitoring Metrics |
| Released the EMR service in Singapore on the international version. | Added support for Singapore Zone 1 and Singapore Zone 2 on the international version. | 2021-01-13 | - |

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. | 2020-12-14 | Version Overview |

November 2020

| Update | Description | Release Date | Documentation |
|-------------------------------------|---|--------------|--|
| Released the EMR TianQiong edition. | Added support for Spark materialized views. | 2020-11-27 | Version Overview EMR TianQiong Introduction |

| Update | Description | Release Date | Documentation |
|------------------------------------|---|--------------|--------------------------|
| Optimized monitoring and alarming. | <ul style="list-style-type: none"> 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. | 2020-11-20 | Overview |

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 inspection items to the inspection system. | Added the HDFS, Yarn, HBase, Hive, and ZooKeeper service inspection items, enhancing the health check of metrics. | 2020-10-13 | Cluster 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. | 2020-09-17 | Version Overview |

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. | 2020-08-27 | Version Overview |
| Added support for data migration in ClickHouse clusters. | Improved cluster resource utilization for data security. | 2020-08-17 | 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. | 2020-07-29 | 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. | 2020-06-15 | Druid Monitoring Metrics |

| Update | Description | Release Date | Documentation |
|---|---|--------------|--|
| Released a Hadoop cluster version: EMR v2.3.0. | Added support for the TensorFlow and Jupyter Notebook components. | 2020-06-15 | 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. | 2020-06-15 | - |
| 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. | 2020-06-04 | Cluster Event Cluster Inspection |

May 2020

| Update | Description | Release Date | Documentation |
|---|--|--------------|---|
| Optimized console interactions, UI, and operation paths. | Improved operation path hierarchy for easier service management. | 2020-05-15 | - |
| Added support for service monitoring for ClickHouse clusters. | Made it easier to keep track of the operation status of ClickHouse clusters. | 2020-05-15 | 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. | 2020-05-15 | - |

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. | 2020-04-15 | 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. | 2020-04-15 | - |
| 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. | 2020-03 | - |
| 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 | Version Overview |

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. | 2020-02 | Hive Metadata Management |
| Optimized cluster operation logs. | Added display items for cluster operation logs to increase the readability of operation logs. | 2020-02 | - |
| 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. | 2020-01 | 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. | 2020-01 | - |
| Released EMR in Nanjing region. | Released EMR in Nanjing region. | 2020-01 | - |

December 2019

| Update | Description | Release Date | Documentation |
|--------|-------------|--------------|---------------|
|--------|-------------|--------------|---------------|

| Update | Description | Release Date | Documentation |
|--|--|--------------|----------------------------------|
| 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. | 2019-12 | - |
| Released the feature of resource-level CAM authorization. | Added support for CAM authorization at the resource level. | 2019-12 | - |
| Released the feature of associating CHDFS. | Added support for associating CHDFS and reading/writing data on CHDFS. | 2019-12 | 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. | 2019-12 | 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. | 2019-12 | - |
| 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. | 2019-12 | - |

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. | 2019-11 | - |
| Released the feature of tagging clusters and nodes. | Added support for tagging clusters and nodes. | 2019-11 | 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. | 2019-10 | Version Overview |

September 2019

| Update | Description | Release Date | Documentation |
|--------|-------------|--------------|---------------|
|--------|-------------|--------------|---------------|

| 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. | 2019-09 | - |
| Modified console configuration. | Disused the escape feature for special characters. | 2019-09 | - |
| 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. | 2019-09 | - |

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. | 2019-08 | 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. | 2019-08 | - |
| Added support for mounting multiple cloud disks to a new cluster. | Added support for mounting multiple cloud disks to a new cluster. | 2019-08 | - |
| Released the feature of scaling specified components. | Added support for scaling specified components. | 2019-08 | - |
| Added monitoring metrics. | Added new monitoring metrics for the Spark, Hive, Presto, and ZooKeeper components. | 2019-08 | - |

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. | 2019-07 | Service Monitoring |
| Released EMR in new regions. | Added support for the Singapore, Mumbai, and Chengdu regions. | 2019-07 | - |
| 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. | 2019-07 | - |
| Optimized WebUI proxy address. | Optimized WebUI proxy address. | 2019-07 | - |

June 2019

| Update | Description | Release Date | Documentation |
|--|--|--------------|-----------------------------------|
| Released the feature of accessing external clusters. | Added support for accessing external clusters. | 2019-06 | - |
| Released the feature of setting bootstrap actions. | Added support for setting bootstrap actions. | 2019-06 | Bootstrap Actions |

| Update | Description | Release Date | Documentation |
|--|--|--------------|--|
| Released the feature of delivering component parameter configuration to nodes. | Added support for delivering component parameter configuration to nodes. | 2019-06 | - |
| Released the feature of rolling back parameter configuration. | Added support for rolling back parameter configuration. | 2019-06 | 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. | 2019-05 | Version Overview |
| 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. | 2019-05 | Kerberos Introduction |
| Optimized monitoring metrics. | Optimized the Node, HDFS, YARN, and HBase monitoring metrics. | 2019-05 | Node Monitoring Metrics |
| Optimized the style of console. | Improved the style of console for a better interactive experience. | 2019-05 | - |
| Optimized CVM and TencentDB naming. | Optimized CVM and TencentDB naming with the EMR cluster serial number for easier locating of cluster information. | 2019-05 | - |

| Update | Description | Release Date | Documentation |
|--|---|--------------|---------------|
| Changed the setting of public IPs of master nodes to optional. | Changed the setting of public IPs of master nodes to optional. | 2019-05 | - |
| Released the feature of adjusting the number of common nodes. | Added support for adjusting the number of common nodes as needed. | 2019-05 | - |

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. | 2019-03 | - |
| Released EMR in a new region. | Made the purchase of EMR clusters available in the Silicon Valley region. | 2019-03 | - |
| 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. | 2019-03 | Cluster Types |
| Released the feature of adjusting node configuration. | Added support for adjusting node configuration. Nodes can be upgraded to a higher configuration. | 2019-03 | Configuration Change |

January 2019

| Update | Description | Release Date | Documentation |
|--------|-------------|--------------|---------------|
|--------|-------------|--------------|---------------|

| Update | Description | Release Date | Documentation |
|--|---|--------------|-----------------------------------|
| Added new components to existing clusters. | Added support for adding new components to existing clusters. | 2019-01 | Adding Components |

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/common/fix-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
```

Here, `1639576622` is a temporarily generated timestamp, and you need to find the corresponding file for decompression.

2. Copy the backup file back.

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

Here, `10812_1639576622` is the number temporarily generated during execution, and you need to find the corresponding file for copying.

3. Delete the added latest Log4j JAR packages.

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

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

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 | <ul style="list-style-type: none"> Cluster alarming (12 metrics) Node alarming (8 metrics) | It was deprecated at 23:00 on April 9, 2021 and no longer maintained. |
| Elastic MapReduce | <ul style="list-style-type: none"> 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.

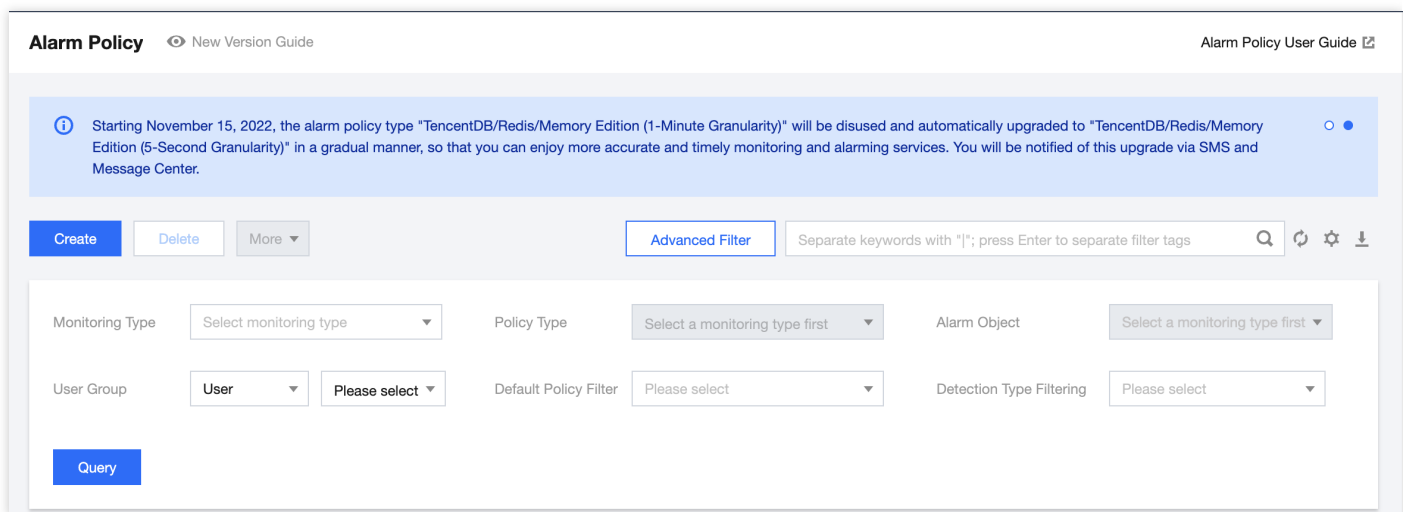
Note :

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.



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 |
| | | | | |

| Original Policy Type | Metric/Event Alarm | Original Metric/Event Name | New Policy Type | New Metric/Event Name |
|----------------------|--------------------|--|---|------------------------------------|
| | 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 |