

Elastic MapReduce

Purchase Guide

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

Purchase Guide

Server-Based EMR

Billing Overview

Purchase Instructions

Billing Mode

Payment Overdue

Container-Based EMR Billing

Overview

Purchase Guide

Server-Based EMR

Billing Overview

Last updated : 2023-05-08 15:22:21

Billing Mode

EMR offers one billing mode for clusters: pay-as-you-go. The cost of a cluster is the sum of the costs of all nodes in the cluster and of associated cloud products. Elastic nodes (task nodes) can be billed on a spot instance basis.

A pay-as-you-go cluster can have pay-as-you-go and spot instance nodes.

The following table lists the differences between these two billing modes:

Billing Mode	Pay-As-You-Go	Spot Instance
Payment method	Postpaid; amount freezing upon purchase, billed hourly	Postpaid; amount freezing upon purchase, billed hourly
Billing unit	USD/second	USD/second
Unit price	Relatively higher	The price fluctuates. In most cases, the price is about 10-20% of the price of a pay-as-you-go instance with the same specifications.
Minimum use time	Charged by second and billed by hour. Purchase and release at any time.	Charged by second and billed by hour. Purchase and release at any time. May be repossessed by the system.
Configuration change	No limit. Change (node CPU and memory only) at any time.	Not supported
Use case	Suitable for a cluster to exist for a short period or periodically.	Suitable for a cluster using elastic compute resources to get more computing power.
Billing mode change	Not supported	Not supported

The pay-as-you-go mode offers 3-tiered pricing, except for model S5 and new models launched after November 2019. A longer usage period means a bigger discount.

- Tier 1: $0 < T1 \leq 96$

- Tier 2: $96 < T2 \leq 360$
- Tier 3: $T3 > 360$

For more information, see [Billing Overview](#).

You can view node specifications and pricing and estimate your resource costs on the [Pricing | Elastic MapReduce](#) page.

EMR bills only show the costs of nodes, namely, the costs of CPU, memory, system disk, and local data disk. The bills of associated cloud products are available in their respective consoles. For EMR bill details, see [Viewing Bills](#) or [Cost Allocation by Tag](#).

Note :

- **Select the shutdown mode with caution when shutting down a pay-as-you-go EMR cluster node in the CVM console, because EMR nodes do not support the "no charges when shut down" mode.**
- **The prices of a model given on the [Pricing | Elastic MapReduce](#) page cover only the configuration of its CPU, memory, and local data disk**, excluding the costs of images, cloud system disks, cloud data disks, and associated cloud products.
- For the pricing of different types of disks, see [Price Overview](#). For charges of associated cloud products, see "Billable Items > Charges of associated cloud products" below.
- The prices are subject to changes as appropriate. Please visit our official website for the latest prices.

Billable Items

The cost of a cluster is the sum of the costs of all nodes in the cluster and of associated cloud products. Cloud products such as Elastic IP, CDB, CBS, CHDFS, COS, and Virtual Private Cloud (VPC) may be used when you use EMR. You will be charged for such products in their respective billing mode. For details, see their respective billing document.

Cost Category	Billable Item	Resource Use in EMR	Billing Description	Pricing
Costs of EMR nodes	Node cost	You can select the model specifications as needed. The nodes are subject to separate pricing by EMR. The price of a node covers its resources such as CPU, memory, system disk, and local data disk. Two billing modes are available: pay-as-you-go and spot instance.	The EMR cost is only the costs of all nodes.	Pricing of Elastic MapReduce

Cost Category	Billable Item	Resource Use in EMR	Billing Description	Pricing
Costs of associated cloud products	Elastic IP (EIP)	Public network access is enabled for the Master.1 node in a cluster by default so that you can access the WebUI pages of various Hadoop components from outside the cluster. Traffic fees are incurred for data interactions when you visit these pages, but a little traffic is generated in most cases. Therefore, the bill-by-traffic mode is used by default, which costs less than the bill-by-bandwidth mode.	Elastic IP Billing	Elastic IP Billing
	TencentDB for MySQL (CDB)	If you plan to deploy one or more components among Hive (local), Hue, Ranger, Oozie, Druid, and Superset in an EMR cluster, you need to purchase a CDB instance for storing metadata.	TencentDB for MySQL Billing Overview	TencentDB for MySQL Pricing
	Cloud Block Storage (CBS)	If you plan to deploy cloud disks as the data disks in a node, you need to purchase at least one cloud disk. The cost of the cloud disk will be paid together with the EMR node purchased, with the corresponding CBS order generated. The cloud disk will be subject to the same billing mode as the node.	CBS Billing Overview	CBS Price Overview
	Cloud Object Storage (COS)	If you use COS to separate the compute and storage resources of a cluster, you will be charged for data storage and requests when the cluster requests to pull data stored in COS for computing. Meanwhile, result data storage, backup, or other operations in the computing will also generate new data in COS.	COS Billing Overview	COS Pricing
	Cloud HDFS (CHDFS)	If you use CHDFS to separate the compute and storage resources of a cluster, you will be charged for data storage and requests when the cluster requests to pull data stored in CHDFS for computing. Meanwhile, result data storage, backup, or other operations in the computing will also generate new data in CHDFS.	CHDFS Purchase Guide	CHDFS Purchase Guide

Purchase Instructions

Last updated : 2023-05-30 11:24:00

Pay-As-You-Go

- You are charged by usage duration of a cluster. This billing mode requires identity verification and will freeze an amount of 2-hour usage fee when the cluster is purchased (vouchers cannot be used here). After this cluster is terminated, the frozen amount will be refunded. Before creating a cluster, check your Tencent Cloud account balance. If your balance is less than the service cost, top up your account first.
- When you purchase an EMR cluster, the price will be listed as an hourly fee. However, you will be billed by the actual seconds of usage and the charge will be rounded to two decimal places. Billing starts from the second the cluster is created and stops the second the cluster is terminated.
- When you purchase a pay-as-you-go cluster, the fee for 2-hour usage under the current configuration will be frozen in your account balance as a deposit. You will then be billed by the hour for your usage over the past hour. When you change the node configuration, the frozen amount will be unfrozen and a new 2-hour deposit will be frozen based on the unit price of the new configuration. Your deposit will be released back to your account when the cluster is terminated.
- For information about what you can do when your account balance is insufficient, see [Payment Overdue](#).
- For information about the use and limits of Tencent Cloud vouchers, see [Promo Vouchers](#).

Spot Instance

- Similar to pay-as-you-go instances, spot instances are charged by second and billed by hour. The prices of spot instances fluctuate according to market demand, which provide you with a substantial discount (about 80-90% off the prices of pay-as-you-go instances with the same specifications). However, spot instances may be repossessed automatically by the system as a result of inventory shortages or higher bids from other users.
- For more information about spot instance policies, use cases, and limitations, see [Spot Instance](#).

Note :

Spot instances only support auto scaling to replace compute nodes. The system may repossess spot instances due to higher bids from other users. Therefore, please use them with caution.

Billing Example

Pay-as-you-go

Assuming you create a Hadoop cluster in Guangzhou Zone 7, with a high-availability cluster of EMR 3.5.0 for the "default scenario". Services deployed include hdfs-3.2.2, yarn-3.2.2, zookeeper-3.6.3, openLDAP-2.4.44, Knox-1.6.1, and hive-3.1.3 (with the cluster default mode selected for the Hive metadata storage mode). The pay-as-you-go mode is adopted, with a discount of 15% for EMR, 10% for CDB, and 5% for CBS.

Note :

The above discounts are provided as an example only. Actual discounts may differ.

EMR cluster fees are shown below:

Node Type	Model Specification	System Disk	Data Disk	Model Specification Fee	System Disk Fee	Data Disk Fee	Number of Nodes
Master	Standard SA2: 4-core CPU, 16 GB memory	SSD 50 GB × 1	SSD 200 GB × 1	0.77	0.0025 × 50 GB × 1	0.0025 × 200 GB × 1	2
Core	Standard SA2: 4-core CPU, 8 GB memory	SSD 50 GB × 1	SSD 200 GB × 1	0.52	0.0025 × 50 GB × 1	0.0025 × 200 GB × 1	3
Common	Standard SA2: 2-core CPU, 4 GB memory	SSD 50 GB × 1	SSD 200 GB × 1	0.26	0.0025 × 50 GB × 1	0.0025 × 200 GB × 1	3
MetaDB	High-IO TencentDB - memory: 4,000 MB, hard disk size: 100 GB, 1 instance					USD 0.48/instance	1

Master node fee = $((0.77 + 0.125) \times 0.85 + 200 \times 0.0025 \times 0.95) \times 2 = 2.4715$ (USD/hour)

Core node fee = $((0.52 + 0.125) \times 0.85 + 200 \times 0.0025 \times 0.95) \times 3 = 3.06975$ (USD/hour)

Common node fee = $((0.26 + 0.125) \times 0.85 + 200 \times 0.0025 \times 0.95) \times 3 = 2.40675$ (USD/hour)

MetaDB node fee = $0.48 \times 0.9 = 0.432$ (USD/hour)

Total fee = $2.4715 + 3.06975 + 1035.324 + 2.40675 = 8.38$ (USD/hour)

Note :

- The prices provided in this example are for reference only. Refer to the purchase page for actual fees.
- For information about model specifications and pricing, see [Pricing | Elastic MapReduce](#).
- The above example only involves the following associated products: CBS and CDB.
- EMR bills only show the costs of nodes, namely, the costs of CPU, memory, system disk, and local data disk. **The bills of associated cloud products are available in their respective consoles.** For EMR bill details, see [Viewing Bills](#) or [Cost Allocation by Tag](#).

Purchase

Evaluating your business

Before purchasing a cluster, you need to evaluate your business according to the actual situation to ensure that the created cluster meets your actual needs. For more information, see [Business Evaluation](#).

Purchasing EMR clusters

Before using EMR services, you need to register a Tencent Cloud account and go to the [EMR purchase page](#) to purchase an EMR cluster. For more information, see [Creating EMR Cluster](#).

Billing Mode

Last updated : 2022-07-04 11:23:36

EMR offers one billing mode for clusters:

- **Pay-as-you-go:** All nodes in a cluster are charged on a [pay-as-you-go](#) basis. This is suitable for clusters that exist for a short time or periodically.

Note :

Select the shutdown mode with caution when shutting down a pay-as-you-go EMR cluster node in the CVM console, because EMR nodes do not support the "no charges when shut down" mode.

For more information on node types, see [Cluster Types](#).

When you purchase an EMR cluster, the price will be listed as an hourly fee. However, you will be billed by the actual seconds of usage and the charge will be rounded to two decimal places. Billing starts from the second the cluster is created and stops the second the cluster is terminated.

When you purchase a pay-as-you-go cluster, the fees for 2-hour usage under the current configuration will be frozen in your account balance as a deposit. You will then be billed by the hour (Beijing time) for your usage over the past hour. When you change the node configurations, the frozen amount will be unfrozen and a new 2-hour deposit will be frozen based on the unit price of the new configuration. Your deposit will be released back to your account when the cluster is terminated.

EMR provides elastic computing cluster capabilities, so that you can select and combine multiple EMR specifications in a customized manner. EMR fees are charged based on all nodes in each cluster. You can view node specification prices in [Pricing | Elastic MapReduce](#) to estimate resource fees.

Note :

- Published prices are subject to change.
- For the prices of disks, see [CBS Price Overview](#).
- The following prices are for configuration fees of CPU and memory, excluding fees of images, system cloud disks, local data disks, data cloud disks, and bandwidth.

Other items that may incur costs include network traffic, metadata storage, and COS.

- **Network traffic**

Public network access is enabled for the Master.1 node in a cluster by default so that you can access the WebUI pages of various Hadoop components from outside the cluster. Traffic fees are incurred for data interactions when you visit these pages, which is very little in most cases. Therefore, the bill-by-traffic mode is used by default, which costs less than the bill-by-bandwidth mode.

- **Metadata storage**

The metadata storage such as Hive on EMR (v2.x and earlier) uses TencentDB, the fee of which is included in the total cost on the purchase page.

- **COS**

When you separate computing and storage for a cluster via COS, data storage and request fees (see [Pricing | Cloud Object Storage](#)) will be incurred when the cluster pulls data from COS for computing. In addition, new data may be generated for the storage or backup of computing results in COS.

Payment Overdue

Last updated : 2023-05-08 15:22:21

Pay-As-You-Go

Overdue payment alert

Type	Description
Overdue payment reminder	Pay-as-you-go clusters are billed on the hour. When your account balance becomes negative, your Tencent Cloud account creator, global resource collaborators, and financial collaborators will be notified by email and SMS.
Overdue payment alert	This feature is disabled by default.

Overdue payment policy

When your account balance falls below zero, a pay-as-you-go cluster can be used for 2 more hours and the corresponding amount will be deducted. After 2 hours, the cluster will be moved to the recycle bin and become unavailable, and we will also stop billing you for service.

Time After Service Suspension	Description
≤ 15 days	If your account is topped up to a positive balance, the charging will continue, and the cluster will be automatically recovered.
	If your account balance remains negative, the cluster cannot be recovered.
> 15 days	If your account is not topped up to a positive balance, your pay-as-you-go cluster resources will be repossessed and released. All data will be cleared and cannot be recovered. When your cluster resources are repossessed, your Tencent Cloud account creator and all collaborators will be notified by email and SMS.

Note :

- If you no longer need to use pay-as-you-go resources, terminate them as soon as possible to avoid further fee deductions.
- After a cluster is terminated, the data in it will be cleared and cannot be recovered.
- Since your actual resource consumption may change over time, some balance alerts may be inaccurate.
- You will continue to be charged for associated cloud products such as EIP and CBS during cluster isolation due to overdue payments until your account balance becomes negative.

Overdue Payments of Associated Cloud Products

Product	Description
TencentDB for MySQL (CDB)	Payment Overdue
CBS	Payment Overdue
COS	Payment Overdue
CHDFS	Payment Overdue
VPC	Overdue Payment Alert Payment Overdue

Container-Based EMR Billing Overview

Last updated : 2022-11-25 14:41:41

Billing Overview

Container-based EMR clusters are deployed in EKS and free you from managing cluster nodes. However, to reasonably allocate resources and accurately calculate fees, you need to specify resource specifications for Pods when configuring a job. Tencent Cloud allocates computing resources to the workload and calculates the fees based on the specifications.

Currently, container-based EMR is pay-as-you-go by the amount of EKS CPU and memory resources required to run services.

Note :

1. The following prices are only for the configurations of CPU and memory, which don't include cloud data disk and bandwidth fees.
2. For cloud disk prices, see [Billing Overview](#).

Billable Items

Resource Type	EKS Resource (Intel)	EKS Resource (AMD)
Billable items	CPU	CPU
	Memory	Memory

Container-based EMR provides Intel resources by default. If you need AMD resources, configure the job accordingly. The service calculates applicable fees based on the container resource type you select.

Billing formula: Fees = billable item configuration * resource unit price * execution time

For more information on configurations currently supported for billable items, see [Resource Specifications](#).

For more information on how to select configurations for billable items, see [Specifying Resource Specifications](#).

Pricing

Pay-as-you-go (Intel)

Billable Item	Price (Per Second)	Price (Per Hour)
CPU	0.0003438 USD/core	0.020628 USD/core/hour
Memory	0.0001434 USD/GB	0.008604 USD/GB/hour

Pay-as-you-go (AMD)

Billable Item	Price (Per Second)	Price (Per Hour)
CPU	0.0001518 USD/core	0.009108 USD/core/hour
Memory	0.0000882 USD/GB	0.005292 USD/GB/hour

Overdue Payment Policy

For more information, see [Payment Overdue](#).