

Elastic MapReduce

Purchase guide

Product Documentation



Copyright Notice

©2013-2019 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

 Tencent Cloud

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

 Billing Overview

 Billing Mode

 About Overdue Payment

Purchase guide

Server-Based EMR

Billing Overview

Billing Overview

Last updated : 2021-09-23 09:24:43

Billing Mode

Pay-as-you-go: the billing mode of all nodes in a cluster is pay-as-you-go. For more information on node types, please see [Node Type Description](#).

When you purchase an EMR cluster, the unit price of the cluster will be displayed by hour. During bill settlement, the price will be calculated based on the number of seconds of actual usage, and the fees will be rounded to 2 decimal places. The billing starts at the time point when the cluster is created and ends at the time point when termination of the cluster is completed.

When you purchase a pay-as-you-go cluster, the amount of 2 hours' usage fees for the current configuration will be frozen in advance. Bills will be settled on the hour (Beijing time), and fees will be deducted based on the actual usage time of the cluster in the previous hour.

When the configuration of a pay-as-you-go node is adjusted, the amount frozen during purchase will be unfrozen, and a new amount of 2 hour's usage fees will be frozen at the latest unit price. When you terminate a pay-as-you-go cluster, the system will unfreeze the frozen amount.

Price Description

EMR provides elastic computing cluster capabilities, so that you can select and combine multiple EMR specifications in a customized manner. EMR fees will be charged for all nodes in each cluster.

The prices of nodes in Beijing, Shanghai, and Guangzhou regions are as listed below:

1. The prices on our official website are subject to change. Please check periodically for the latest prices.
2. For the prices of disks, please see [CBS Price Overview](#).
3. The following prices are for configuration fees of CPU and memory, excluding fees of images, system cloud disks, data cloud disks, and bandwidth.

4. Some models may be sold out in the actual situation. For more information, please see the model list on the purchase page.

Tier-1 pay-as-you-go fees (USD/hour)

Model	CPU	Memory (GB)	Beijing, Shanghai, and Guangzhou
Standard S2	2	4	0.12
Standard S2	4	8	0.23
Standard S2	4	16	0.34
Standard S2	8	16	0.45
Standard S2	8	32	0.67
Standard S2	12	24	0.67
Standard S2	16	32	0.89
Standard S2	16	64	1.33
Standard S2	24	48	1.33
Standard S2	24	96	1.99
Standard S2	32	64	1.77
Standard S2	32	128	2.65
Standard S2	48	125	3.05
High-IO I2	8	16	0.45
High-IO I2	8	32	0.67
High-IO I2	12	24	0.67
High-IO I2	16	32	0.89
High-IO I2	16	64	1.33
High-IO I2	24	96	1.99
High-IO I2	32	64	1.77
High-IO I2	32	128	2.65

Model	CPU	Memory (GB)	Beijing, Shanghai, and Guangzhou
MEM Optimized M2	2	16	0.23
MEM Optimized M2	8	64	0.91
MEM Optimized M2	12	96	1.36
MEM Optimized M2	16	128	1.81
MEM Optimized M2	24	192	2.71
MEM Optimized M2	32	256	3.61
Computing C2	4	8	0.33
Computing C2	4	16	0.5
Computing C2	16	32	1.31
Computing C2	16	60	1.89
Computing C2	32	120	3.78
Big Data D1	8	32	0.72
Big Data D1	16	64	1.44
Big Data D1	24	96	2.16
Big Data D1	32	128	2.88
Big Data D1	56	224	5.03
Big Data D2	8	32	0.68
Big Data D2	16	64	1.36
Big Data D2	24	96	2.04
Big Data D2	32	128	2.72
Big Data D2	64	256	5.44
Big Data D2	76	320	6.68
MEM Optimized M4	2	4	0.11
MEM Optimized M4	4	32	0.46

Model	CPU	Memory (GB)	Beijing, Shanghai, and Guangzhou
MEM Optimized M4	8	64	0.91
MEM Optimized M4	12	96	1.36
MEM Optimized M4	12	144	1.85
MEM Optimized M4	16	128	1.81
MEM Optimized M4	16	192	2.47
MEM Optimized M4	32	256	3.61
MEM Optimized M4	32	384	4.94
MEM Optimized M4	64	512	7.22
MEM Optimized M4	72	648	8.86
Standard S3	2	4	0.12
Standard S3	2	8	0.17
Standard S3	4	8	0.23
Standard S3	4	16	0.34
Standard S3	8	16	0.45
Standard S3	8	32	0.67
Standard S3	16	32	0.89
Standard S3	16	64	1.33
Standard S3	24	48	1.33
Standard S3	24	96	1.99
Standard S3	32	64	1.77
Standard S3	32	128	2.65
Standard S3	48	96	2.65
Standard S3	48	192	3.98
Standard S3	64	128	3.54

Model	CPU	Memory (GB)	Beijing, Shanghai, and Guangzhou
Standard S3	64	256	5.3
High-IO I3	8	32	0.68
High-IO I3	16	64	1.35
High-IO I3	24	96	2.02
High-IO I3	32	128	2.69
High-IO I3	48	192	4.03
High-IO I3	64	256	5.38
High-IO I3	80	320	6.72
Standard S4	1	4	0.09
Standard S4	2	4	0.12
Standard S4	2	8	0.17
Standard S4	4	8	0.23
Standard S4	4	16	0.34
Standard S4	8	16	0.46
Standard S4	8	32	0.68
Standard S4	16	32	0.91
Standard S4	16	64	1.35
Standard S4	24	48	1.36
Standard S4	24	96	2.02
Standard S4	32	64	1.81
Standard S4	32	128	2.69
Standard S4	48	96	2.71
Standard S4	48	192	4.03
Standard S4	64	128	3.61

Model	CPU	Memory (GB)	Beijing, Shanghai, and Guangzhou
Standard S4	64	256	5.38
Standard S4	72	288	6.05
Big Data DS2	8	32	0.68
Big Data DS2	16	64	1.36
Big Data DS2	24	96	2.04
Big Data DS2	32	128	2.72
Standard Network Optimized SN3ne	2	4	0.12
Standard Network Optimized SN3ne	4	8	0.23
Standard Network Optimized SN3ne	4	16	0.34
Standard Network Optimized SN3ne	8	16	0.46
Standard Network Optimized SN3ne	8	32	0.68
Standard Network Optimized SN3ne	12	24	0.68
Standard Network Optimized SN3ne	16	32	0.91
Standard Network Optimized SN3ne	16	64	1.35
Standard Network Optimized SN3ne	24	48	1.36
Standard Network Optimized SN3ne	24	96	2.02
Standard Network Optimized SN3ne	32	64	1.81
Standard Network Optimized SN3ne	32	128	2.69
Standard Network Optimized SN3ne	48	96	2.71
Standard Network Optimized SN3ne	64	128	3.61
Standard Network Optimized SN3ne	64	192	4.49
Standard Network Optimized SN3ne	64	256	5.38
Standard Network Optimized SN3ne	72	288	6.05
Standard Network Optimized SN3ne	76	192	4.84

Model	CPU	Memory (GB)	Beijing, Shanghai, and Guangzhou
Standard Network Optimized SN3ne	76	256	5.72
MEM Optimized M5	2	16	0.23
MEM Optimized M5	4	32	0.46
MEM Optimized M5	8	64	0.91
MEM Optimized M5	12	96	1.36
MEM Optimized M5	32	256	3.61
Standard S5	2	4	0.07
Standard S5	2	8	0.11
Standard S5	4	8	0.14
Standard S5	4	16	0.21
Standard S5	8	16	0.28
Standard S5	8	32	0.41
Standard S5	16	32	0.55
Standard S5	16	64	0.82
Standard S5	24	48	0.82
Standard S5	24	64	0.95
Standard S5	24	96	1.23
Standard S5	32	64	1.09
Standard S5	32	128	1.63
Standard S5	48	96	1.63
Standard S5	48	192	2.45
Standard S5	64	192	2.72
Standard S5	64	256	3.26
Standard S5	76	256	3.46

Model	CPU	Memory (GB)	Beijing, Shanghai, and Guangzhou
High-IO IT3	16	64	1.44
High-IO IT3	32	128	2.88
High-IO IT3	64	256	5.75
High-IO IT3	76	160	4.67
Computing Network Enhanced CN3	4	8	0.28
Computing Network Enhanced CN3	4	16	0.42
Computing Network Enhanced CN3	8	16	0.56
Computing Network Enhanced CN3	8	32	0.83
Computing Network Enhanced CN3	16	32	1.11
Computing Network Enhanced CN3	16	64	1.66
Computing Network Enhanced CN3	32	64	2.21
Computing Network Enhanced CN3	32	128	3.32
Standard SA2	2	4	0.14
Standard SA2	4	4	0.07
Standard SA2	4	8	0.09
Standard SA2	8	8	0.14
Standard SA2	8	16	0.18
Standard SA2	16	32	0.36
Standard SA2	16	64	0.53
Standard SA2	32	64	0.71
Standard SA2	64	128	1.41
Standard SA2	64	192	1.77
Standard SA2	80	160	1.77
Standard SA2	128	256	2.82

Model	CPU	Memory (GB)	Beijing, Shanghai, and Guangzhou
Standard SA2	160	320	3.53
Computing C3	4	8	0.35
Computing C3	4	16	0.52
Computing C3	8	16	0.69
Computing C3	8	32	1.04
Computing C3	16	32	1.38
Computing C3	16	64	2.08
Computing C3	32	64	2.76
Computing C3	32	128	4.16

Other billable items that may incur fees when you use an EMR cluster include network traffic, metadata storage, and COS.

- **Network traffic**

Public network access is enabled for the Master.1 node of the cluster by default, so that WebUIs of various Hadoop components can be accessed from outside the cluster. Traffic fees will be incurred by the data interaction when these pages are accessed, which are very low in most cases; therefore, bill-by-traffic is used by default, which is more cost-effective than bill-by-bandwidth.

- **Metadata storage**

TencentDB for MySQL is used for metadata storage such as Hive in EMR 2.x and below. The storage fees are included in the total fees on the purchase page.

- **COS**

If you use COS to isolate storage and computation of your cluster, data storage and request fees will be incurred (for more information, see [COS Product Pricing](#)) when your cluster makes a computation request and pulls data from COS. In addition, new data may be generated in COS during computation for the purposes of result data storage or backup.

Billing Mode

Last updated : 2021-12-27 12:12:16

Billing mode for EMR 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.

For information on node types, see [Node Type Description](#).

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.

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 lower) uses TencentDB, the fee of which is included in the total cost on the purchase page.

- **COS**

When you separate compute and storage for a cluster via COS, data storage and request fees (see [COS Product Pricing](#)) 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 compute results in COS.

About Overdue Payment

Last updated : 2021-03-01 11:36:10

Pay-As-You-Go

Overdue payment reminder

Reminder 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 through email and SMS.
Overdue payment alert	This feature is disabled by default.

Overdue payment policy

When your account balance becomes negative, the pay-as-you-go cluster can be used and fees will be deducted for 2 more hours. 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 billing 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 erased and cannot be recovered. When your cluster resources are repossessed, your Tencent Cloud account creator and all collaborators will be notified through email and SMS.

 **Note :**

- When you no longer need to use pay-as-you-go clusters, please 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.