

Alibaba Cloud ApsaraDB for HBase

Quick start

Issue: 20200228

Legal disclaimer

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

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequent

ial, exemplary, incidental, special, or punitive damages, including lost profits arising from the use or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
6. Please contact Alibaba Cloud directly if you discover any errors in this document

.

Document conventions

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

Style	Description	Example
<code>{}</code> or <code>{a b}</code>	This format is used for a required value, where only one item can be selected.	<code>switch {active stand}</code>

Contents

Legal disclaimer.....	I
Document conventions.....	I
1 Cluster sizing.....	1
2 Select an ApsaraDB for HBase edition.....	7
3 Product editions (single-node/cluster edition).....	10

1 Cluster sizing

ApsaraDB for HBase allows you to choose different specifications, numbers, and disk types for master and core nodes. You can choose specifications and disk types based on your workload requirements, such as the queries per second (QPS), storage capacity, response latency, and stability. You can size your cluster as follows:

- Determine the specification of the master nodes.
- Determine the specification and number of core nodes.
- Determine the size and type of disks.

For more information about how to select an ApsaraDB for HBase version, see [ApsaraDB for HBase versions](#).

Node specification families

- **Dedicated nodes:** Dedicated resources are allocated to nodes to ensure their stability. If your workloads require a low response latency, use dedicated nodes and standard SSDs.
- **General nodes:** General nodes compete for resources. All general nodes have less than 8 vCPUs. If your cluster is used for production purposes, we recommend that you do not use general nodes.

Master nodes

No storage is attached to master nodes. By default, two master nodes (primary and backup) are used for disaster recovery in case of single point of failures. Master nodes are important. HBase masters, HDFS namenodes, and ZooKeeper are deployed on master nodes. If the master nodes do not have sufficient CPU or memory resources, the performance of your cluster is severely degraded.

Number of core nodes	Master node specification
< 4	4c8g
$4 \leq \text{Number of core nodes} < 8$	8c16g(recommended for small-sized clusters)
$8 \leq \text{Number of core nodes} < 16$	8c32g
> 16	16c64g or above

Note: The specification of master nodes also determines the numbers of tables and regions managed by the cluster. If the cluster manages a large number of tables or regions, select a high specification for the master nodes.

Core nodes

Core node specification: The minimum specification is 4c8g and the maximum specification is 32c128g.

Core nodes refer to RegionServers in HBase. You must select a core node specification based on the number of requests to be processed and the size of each request.

Note: The number of requests is not the only standard to size the core nodes. For example, your workload needs to process hundreds of requests per second. In most cases, you can select 4c8g core nodes to handle the requests. However, this rule is not applicable in the following cases: a row that stores kilobytes or megabytes of data is queried, a scan request contains complex filters, the cache hit rate is low, and the cluster manages a large number of tables or regions. In these cases, using 4c8g core nodes may affect the stability of your workload or increase the response latency.

The following table lists the recommended core node specifications for handling different workloads. For more information about storage sizing, join the [ApsaraDB for HBase Q&A DingTalk group](#) or submit a ticket.

TPS + QPS	Recommended number and specification of core nodes	Suggestion
0 to 1,000	2 core nodes with 4c16g	The minimum specification recommended for handling light loads. We recommend that you do not deploy more than 600 regions on each core node. The minimum specification available is 4c8g. We recommend that you do not choose 4c8g because using only 8 GB of memory may cause out of memory errors when the load or KV store significantly grows.
1,000 to 20,000	2 to 3 core nodes with 8c32g	In comparison to 8c16g, 8c32g is more cost-effective. It offers an additional 16 GB of memory to guarantee the stability of your workloads. We recommend that you choose 8c32g to handle light and medium loads.

TPS + QPS	Recommended number and specification of core nodes	Suggestion
More than 20,000	8c32g, 16c32g, 32c64g, or higher	Specify the number of core nodes based on the actual number of requests. If your workload is deployed online, we recommend that select a specification with larger memory to increase the cache hit rate. If you need to run MapReduce or Spark tasks offline with heavy loads, or when the TPS or QPS is extremely high, select a specification with more CPU resources.

Select a higher specification or add more core nodes

You can scale out your ApsaraDB for HBase cluster by adding more core nodes when the load spikes, response latency increases, or the cluster becomes unstable. However, hotspotting may occur if your workload in the cluster is not designed or served properly. The specification of a core node determines its capability to prevent hotspotting. Using low-specification core nodes and then simply adding more core nodes is not as efficient as using high-specification core nodes to guarantee the stability of your workload when the load spikes. For example, if large requests are directed to the nodes or the user traffic spikes in a region, the nodes with a low specification may be overloaded or run out of memory. As a result, the performance of your cluster is degraded.

We recommend that select a specification for your core nodes based on the requirements of your workload.

If the specification of your master or core nodes fails to meet the expected requirements, you can upgrade the nodes. For more information about how to upgrade nodes, join the [ApsaraDB for HBase Q&A DingTalk group](#) or submit a ticket.

Storage types

ApsaraDB for HBase supports three types of storage: standard disks, local disks, and cold storage.

- **Standard disks:** Standard disks are scalable and highly reliable. We recommend that you use standard disks. They are replicated to ensure redundancy and can be expanded based on your needs. Unlike physical disks, standard disks are independent of hardware specifications, which can help you prevent data loss caused by physical damages. Standard disks include SSDs and ultra disks.
- **Local disks:** Local disks are physical disks. Local disks are charged at a lower price rate than standard disks. Local disks are unscalable. The specification of the core nodes that you choose determines the sizes of the local disks that you can choose. You cannot upgrade Elastic Compute Service (ECS) instances attached with local disks to get more storage space. When two or more physical disks are damaged, your workloads will be affected. Your data will not be lost if only one physical disk is damaged. The ApsaraDB for HBase support team will replace the damaged disk for you at the earliest opportunity. The startup costs for local disks are high. Local disks are more suitable for storing large amounts of data.
- **Cold storage:** Cold storage is dedicated to ApsaraDB for HBase. Cold storage is based on Object Storage Service (OSS). It helps you reduce data archiving costs. You can use [cold storage](#) to store infrequently accessed data or use the [hot and cold data separation](#) feature to automatically archive data.

After you create a cluster, you can no longer change the storage type of the cluster. If you choose standard disks, you can increase the storage space by expanding the standard disks or adding more core nodes. If you choose local disks, you can increase the storage space only by adding more core nodes. Cold storage is an exception. You do not need to activate cold storage when you create a cluster. You can activate it and expand the storage space after the cluster is created.

Feature	Storage type	Workload type
High performance	Standard SSDs or local SSDs	For online workloads that require extremely low response latency, such as advertising, recommendation, feed streaming, and user profiling. In this case, you can use SSDs to ensure a low response latency of 1 to 2 milliseconds and reduce noises. SSDs are suitable for users that require a low P99 latency. (P99 latency: 99 % of the requests must be faster than the given latency.)
High efficiency	Ultra disks or local HDDs	For latency-sensitive online workloads. In most cases, if you use HDDs, the response latency will be around 10 milliseconds. However, it produces more noises than using SSDs.
Cold data storage	OSS (cold storage)	Near-line storage and data archiving. Using cold storage with standard/local disks can achieve almost the same write throughput as using standard/local disks. The read throughput of cold storage is low, with a latency of tens of milliseconds. For more information, see Cold storage .

2 Select an ApsaraDB for HBase edition

ApsaraDB for HBase Standard Edition is developed based on the distribution of the open source HBase. It supports custom features, and its stability and operations and maintenance capabilities are deeply optimized. These capabilities are based on the following optimization and improvements: 1. Optimization of the internal distribution (Performance-enhanced Edition). 2. Patches in the new open source edition. 3. Bug fixes and user experience optimization. The custom features supported by ApsaraDB for HBase can be easily migrated across versions. When a new community edition is released, Alibaba Cloud can quickly develop an upgraded version of ApsaraDB for HBase based on the new community edition. By using a comprehensive internal testing and regression system, Alibaba Cloud ensures that you can enjoy the latest and most comprehensive features with high stability and improved user experience. Currently, ApsaraDB for HBase Enterprise Standard Edition supports two versions: V1.1 and V2.0. These versions are based on HBase 1.1.2 and HBase 2.0.0, and are fully compatible with the open source editions.

ApsaraDB for HBase Performance-enhanced Edition

ApsaraDB for HBase Performance-enhanced Edition uses the distribution of HBase inside Alibaba, and has been tested for more than eight years to serve all business units in the Alibaba economy system. It focuses on maturity, stability, high performance, cost-effectiveness, multi-tenant, security, and other capabilities that are needed in volumetric data storage and processing scenarios. In the development process, Alibaba Cloud integrates and optimizes features from various open source distributions so that they can be used in production. Alibaba Cloud also contributes to the open source communities by sharing their technologies. Users can access ApsaraDB for HBase Performance-enhanced Edition by using the official client or standard APIs. This edition supports HBase-1.x and HBase-2.x clients. ApsaraDB for HBase is nearly 100% compatible with these native features and APIs. ApsaraDB for HBase Performance-enhanced Edition delivers seven-fold performance of the open source edition at half of the storage cost. For more information, see [Benefits](#).

ApsaraDB for HBase Serverless Edition

ApsaraDB for HBase Serverless Edition is a serverless service developed based on the kernel of the performance-enhanced edition. It is available at an ultra-low starting price of CNY 6 per month. The Serverless Edition is an extremely flexible edition that is ideal for scenarios such as training, learning, testing, intermittent loads, and small-scale production. For more information, see [Serverless documentation](#). For more information about how to choose clusters and Serverless product forms, see [Comparison document](#).

Choose an edition

- **ApsaraDB for HBase Performance-enhanced Edition is developed based on the internal HBase distribution maintained by Alibaba Group. This edition is superior in terms of maturity, stability, and cost-effectiveness. It also provides enterprise-class features and management tools, and has been tested in core businesses inside Alibaba Group. It is ideal for general enterprise customers.**
- **ApsaraDB for HBase Standard Edition is a customized open source edition that is optimized by Alibaba Cloud. It is ideal for scenarios where the performance-enhanced edition cannot meet the user needs.**
- **ApsaraDB for HBase Serverless Edition is a Serverless service, which is suitable for low-cost learning and testing scenarios.**
- **Note: ApsaraDB for HBase Performance-enhanced Edition does not support some ecosystem components such as Phoenix, Kylin, and Geomesa. Alibaba Cloud is making efforts on optimization. Other features and client APIs are complied with the open source standards.**

Description

	Cloud HBase Enhanced Edition	Cloud HBase Standard Edition	HBase Op Source Edit
Performance optimization	✓ (7 times)	◆ (30%)	✗
Lightweight SQL	✓	◆	✗
Global secondary index	✓	◆	✗
Full-text search	✓	✗	✗
Cold and hot separation	✓	✗	✗
Enterprise backup	✓	✓	✗
Active-active Dual clusters	✓	✓	✗
Strong consistency across data centers	✓	✗	✗
Smart diagnosis	✓	◆	✗

3 Product editions (single-node/cluster edition)

ApsaraDB for HBase provides the following editions.

	Cluster Edition	Single-node Edition	Serverless Edition	Cross-zone Dual-cluster Edition (consult for details)
Charges (check the actual charges on the buy page)	> RMB 2,500/month	RMB 100 to 1,000/month	> RMB 6 /month	> RMB 5,000/month
Scenarios	Any scale of production in enterprises	Testing period, preview period, and small scale production	Training, testing, periodic operation, and small scale production	Any scale of production in enterprises
SLA	Standard SLA	N/A	N/A	Higher SLA
Features	All	All	Basic data consumption (advanced features are not supported, see the following)	All
Scalability	High (nearly unlimited)	Medium (can be dynamically upgraded to Cluster Edition)	Medium (scalable within a certain scale)	High (nearly unlimited)

	Cluster Edition	Single-node Edition	Serverless Edition	Cross-zone Dual-cluster Edition (consult for details)
Billing method	Unit price × Number of nodes (subscription and pay-as-you-go on an hourly basis are supported)	Unit price × Number of nodes (subscription and pay-as-you-go on an hourly basis are supported)	Charged based on requests and storage (subscription and pay-as-you-go on an hourly basis are supported)	Unit price × Number of nodes (subscription and pay-as-you-go on an hourly basis are supported)
Architecture	Distributed deployment (multi-replica and automatic disaster recovery)	Single node deployment (disaster recovery is not supported)	Multi-tenant sharing and distributed deployment (multi-replica and automatic disaster recovery)	Dual-cluster deployment and data center-based disaster recovery

Serverless Edition features

- The native HBase API is supported.
- SQL (Phoenix API) is not supported (planning).
- Full-text search (Solr API) is not supported (planning).
- HDFS open access is not supported.
- For more information, see [ApsaraDB for HBase Serverless Edition](#)