

Cloudera Runtime 7.2.7

Phoenix Frequently Asked Questions

Date published: 2020-02-29

Date modified: 2021-02-04

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with a stylized 'E' that has a horizontal bar extending to the right.

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2025. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Frequently asked questions.....	4
--	----------

Frequently asked questions

Find answers to frequently asked questions around Apache Phoenix and its deployment.

Can Phoenix be used for ETL use cases?

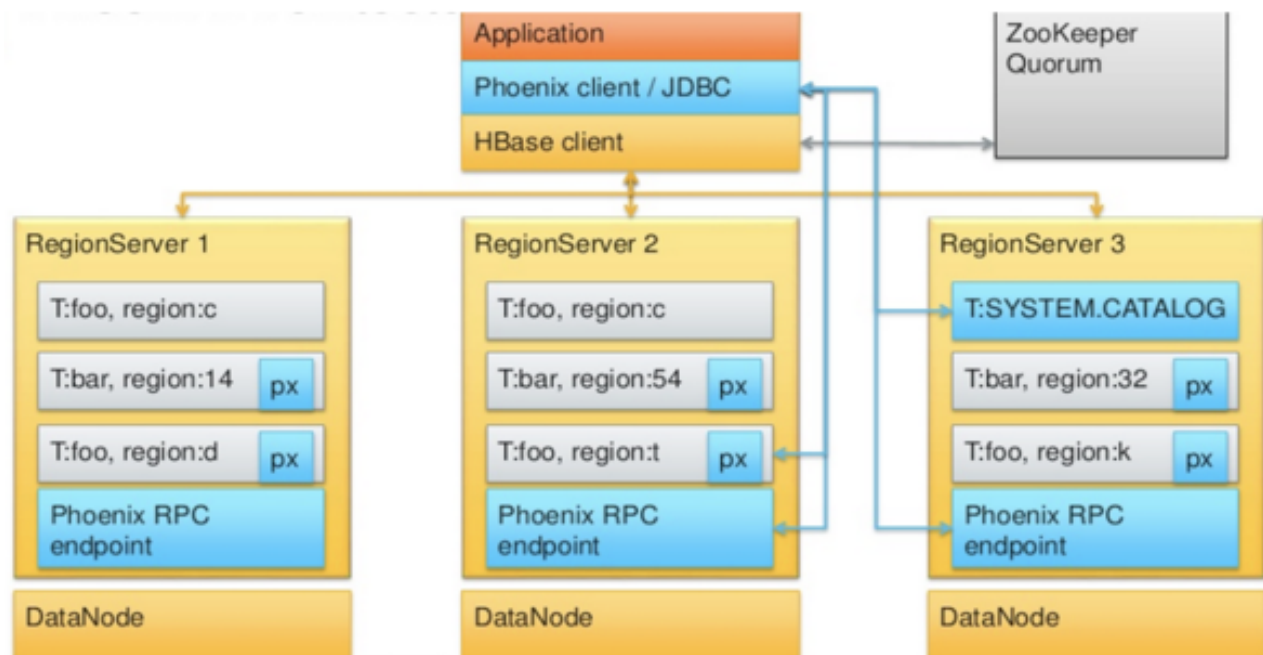
Yes. Apache Phoenix is used for OLTP (Online Transactional Processing) use cases and not OLAP (Online Analytical Processing) use cases. Although, you can use Phoenix for real-time data ingest as a primary use case.

What is the typical architecture for a Phoenix deployment?

A typical Phoenix deployment has the following:

- Application
- Phoenix Client/JDBC driver
- HBase client

A Phoenix client/JDBC driver is essentially a Java library that you should include in your Java code. Phoenix uses HBase as storage similar to how HBase uses HDFS as storage. However, the abstraction for Phoenix is not yet complete, for example, for implementing access controls, you need to set ACLs on the underlying HBase tables that contain the Phoenix data.



Are there sizing guidelines for Phoenix JDBC servers?

For Phoenix applications, you must follow the same sizing guidelines that you follow for HBase. For more information about Phoenix performance tuning, https://phoenix.apache.org/tuning_guide.html.

Can I govern access to the Phoenix servers?

Yes, you can use Kerberos for authentication. You can configure authorization using HBase authorization.

Can I see the individual cell timestamps in Phoenix tables? Is this something that's commonly used?

You can map HBase's native row timestamp to a Phoenix column. By doing this, you can take advantage of the various optimizations that HBase provides for time ranges on the store files as well as various query optimization capabilities built within Phoenix.

For more information, see <https://phoenix.apache.org/rowtimestamp.html>

What if the Phoenix index is being built asynchronously and data is added to a table during indexing?

Phoenix does local Indexing for deadlock prevention during global index maintenance.: Phoenix also atomically rebuild index partially when index update fails (PHOENIX-1112).

How do sequences work in Phoenix?

Sequences are a standard SQL feature that allows for generating monotonically increasing numbers typically used to form an ID.

For more information, see <https://phoenix.apache.org/sequences.html>.

What happens to a Phoenix write when a RegionServer fails?

Writes are durable and durability is defined by a WRITE that is committed to disk (in the Write Ahead Log). So in case of a RegionServer failure, the write is recoverable by replaying the WAL. A "complete" write is one that has been flushed from the WAL to an HFile. Any failures will be represented as exceptions.

Can I do bulk data loads in Phoenix?

Yes, you can do bulk inserts in Phoenix. For more information see https://phoenix.apache.org/bulk_dataload.html.

Can I access a table created by Phoenix using the standard HBase API?

Yes, but it is not recommended or supported. Data is encoded by Phoenix, so you have to decode the data for reading. Writing to the HBase tables directly would result in corruption in Phoenix.

Can I map a Phoenix table over an existing HBase table?

Yes, as long as Phoenix data types are used. You have to use asynchronous indexes and manually update them since Phoenix won't be aware of any updates.

What are guideposts?

For information about guideposts, see https://phoenix.apache.org/update_statistics.html.