

JaguarDB NoSQL Database

Frequently Asked Questions

1. Is Jaguar relational or NoSQL database?
Jaguar is NoSQL, meaning scalable with some standard SQL features.
2. What types of data does JaguarDB support?
Time series data (data that is associated with timestamps), geolocation or geospatial data, and any types of data that are supported by relational databases and nonrelational databases.
3. Does Jaguar support indexes?
Yes, you can create as many indexes as you want for each table. In each index, you can duplicate some columns in the original table for fast data access.
4. Is Jaguar distributed?
Yes, data is distributed among multiple servers in the cluster.
5. How many nodes can Jaguar support?
There is no limit. You can run it on as many nodes as you can.
6. Is Jaguar a in-memory database?
No, Jaguar is a persistent storage database. Memory is used for temporary caching and computing in Jaguar. Of course, Jaguar runs faster if there is more memory on each node.
7. Is JaguarDB optimized for SSD?
Yes, Jaguar runs on HDD as well as on SSD. On SSD, it may use less DRAM.
8. Is Jaguar developed in Java?
No, Jaguar server is written with C++. Client API is in Java, Scala, NodeJS, PHP, Python, etc.
9. I have started Jaguar server, why jag has connection failures?
When Jaguar starts up, it may take a while for it to read some data and warm up. After a few seconds, jag should be able to connect OK.

10. I still get connection error on jag, what happened?
If you are running the enterprise version, make sure the license is correctly installed.
11. I copied the commands from the Web or the document, why I get errors?
The Web pages or the document may contain special characters. You should make sure the commands are in pure text format.
12. When I run Java client programs, I get error of “Unable to load library”. Why?
Make sure you have included the \$JAGUAR_HOME/jaguar/lib in LD_LIBRARY_PATH so that the libJaguarClient.so file can be found by your Java program.
13. When I start and stop Jaguar server, should I do this on each node?
Yes, you can login to each node and execute the scripts or using some tools to manage your cluster and execute the scripts from only one host.
14. I have multiple network cards on my server, how can I connect to Jaguar?
If you have multiple NIC cards and IP addresses, you must specify which IP address your Jaguar should listen to. This is specified by LISTEN_IP in conf/server.conf. When your client connects to Jaguar, it also must use the same IP address.
15. How many cores does Jaguar require?
Jaguar does not have requirement on the number of cores. However, more cores will enable Jaguar to accept more client connections and speed up its read performance.
16. What is the jagexportsql program?
You can execute this script to pull all data records from all nodes and save data in a single file. You just need to run it from one node. The data can be imported to database with the jagimportsqli program.
17. What is the jagexport program?
This program reads all data from a table and writes the data to a temp file in the export directory. The table can be dropped, recreated with a new schema. The jagimport program can be executed to reload the data from the temp file to the table.
18. How can I keep snapshot of tables?
There are local backup settings in the server.conf file which can specify the interval for taking snapshots of all tables in Jaguar.
19. Does one of Jaguar nodes goes down, will Jaguar continue working?
Yes, Jaguar cluster has fault-tolerance support in which if one node goes down or several hosts have hardware failure, the Jaguar cluster continue to work. Inserting new data and querying data still work.

20. How many replicates does Jaguar offer for each data record?
Jaguar can store a maximum of three copies for each data record. The replication factor is set by the REPLICATION parameter in conf/server.conf file.
21. What is Jaguar in terms of CAP theory?
Jaguar is an AP system. In case of hardware failure or network glitch, Jaguar favors availability over consistency. It offers eventual consistency.
22. How long it will take to install Jaguar on a cluster of 100 hosts?
A few minutes. Suppose you have a user account on all the 100 hosts and the account password is the same on the hosts, then you just need to provide a file listing all the 100 hosts and execute one shell script, which will automatically install Jaguar software on all the 100 hosts. Some ssh servers do not allow login with passwords, in such case you need root admin to enable "PasswordAuthentication yes" in the /etc/ssh/sshd_config file and restart sshd server.
23. I have Windows server, how can I run jaguar?
JaguarDB does not support Windows now.
24. I have Jaguar installed on Centos 7 hosts, but I cannot connect to the Jaguar server. What went wrong?
Probably Centos 7 firewall is blocking connections to port 8888. You need to open the port 8888 for connections (run as root or sudo):
- ```
firewall-cmd --zone=public --add-port=8888/tcp --permanent
firewall-cmd --reload
systemctl restart firewalld.service
```
25. For time series data, how many time windows can I create?  
You create as many as you want. For example, you can setup 5-minute windows, 15-minute windows, hourly windows, daily windows, 7-day windows, etc. There is no limit.
26. What is a tick?  
A tick represents a time window. For example, a 5-minute window is a tick. A 3-hour window is another tick. A tick is one of the segments along the timeline.
27. When can I query the data in the ticks?  
You can query the aggregation data in the tick as soon as the time series data is written to the base tables.
28. What timestamp is used in the tick tables?

The timestamp in the tick tables is the starting time of the window. For example, an hourly tick has the start time of the hour, i.e., the minute and second of the timestamp are always zero.

29. Can I use the window function on the base tables?

Yes, you can use the window function on the base as well as on the tick tables.

30. What are the metrics in the location data?

These are metrics data at each location. Other databases might allow you store only one metric value at each location. JaguarDB allows you to store unlimited metrics at each location. For example, at each location, you can store temperature, wind speed, air quality, population density, traffic speed, average traffic congestion index, zone info, municipal info, geopolitical info, crime data, school data, demographics info, income level, etc.

31. Can I combine location query and time query together?

Yes, JaguarDB natively supports location data and time data, both data being stored in the same table in the same file. So naturally it allows for reading both types data in one query.

32. Is time series data distributed?

Yes.

33. Is location data distributed?

Yes. Distribution of data allows high capacity and performance.