

JaguarDB

Vector Database

Server Configuration

Configuration file `$JAGUAR_HOME/jaguar/conf/server.conf` includes the following parameters (JAGUAR_HOME by default points to your HOME environment variable):

1. `PORT=8888`
Port is the port number that all JaguarDB servers listen to. Default is 8888.
2. `VEVECTOR_STORE_MAX_ELEMENTS=1000000`
This number of maximum number vectors allowed on each node, for each key definition, in each database, in each table, for each vector index store. A key definition refers to the key string in “DISTANCE_INPUT_QUANTIZATION” triplets. On a JaguarDB distributed vector database system, there can be multiple nodes, and multiple databases. Each database can have multiple tables, where each table can have multiple vector columns. Each vector column can have multiple index stores. Each index store is uniquely defined by a key definition.
3. `REPLICATION=1`
This parameter determines the number of copies saved for each data record and vector index store. All data is replicated on multiple server hosts. Default value is 1, and the maximum value is 3.
4. `LISTEN_IP=192.168.3.100`
This parameter is the IP address that the server will use if there are multiple network interfaces on the same server host. If there is only one IP address on the server host, this parameter should be commented out and ignored. Normally this configuration should be left commented out.
5. `INIT_EXTRA_THREADS=50`

This parameter is the extra initialization threads for client connections
If there are more concurrent client connections, you should consider increasing this number. There are 15 initialization threads for each host and each replica.

6. `SPARE_COLUMN_PERCENT=30`

This is the percentage value for creating extra space to store future new columns based on current record size. The extra space is pre-allocated for easy and quick addition of new columns in a table.

7. `GROUPBY_SORT_SIZE_MB=1024`

The memory buffer size in sorting of group by operations.

8. `MEMORY_MODE=high`

This parameter specifies the memory usage type on JaguarDB server. If it is high, then a small amount of extra memory is used by JaguarDB in storing block indexes for sorted key values. If it is low, then less memory is used by JaguarDB, relying on SSD or HDD for key block indexing. Default value is high.

9. `JAG_LOG_LEVEL=0`

A Lower number (min is 0) indicates to the server to generate less logging messages. A higher number (max is 9) makes the server generate more debugging information.

10. `LOCAL_BACKUP_PLAN` Specifies when and how data is backed up. There are five types of intervals when duplicate data is saved: 15MIN, HOURLY, DAILY, WEEKLY, and MONTHLY. When data is saved, it can be either SNAPSHOT or OVERWRITE mode. SNAPSHOT means each and separate copy of data is saved with a timestamp (uses more storage space as times goes on). OVERWRITE means only one copy of data is saved. The format for `REMOTE_BACKUP_PLAN` is frequency:policy|frequency:policy|... where frequency is one of 15MIN, HOURLY, DAILY, WEEKLY, and MONTHLY, and policy is one of SNAPSHOT or OVERWRITE. If no value is provided for `BACKUP_PLAN`, then no data is saved as backup.

11. `REMOTE_BACKUP_SERVER` and `REMOTE_BACKUP_INTERVAL`: These parameters specify remote backup server IP address and backup interval in

seconds. If these parameters are provided, all servers in the cluster will periodically send local data to the remote server for backup.

12. `SERVER_TOKEN=...string...`

This text string is used as authentication token between server communication. It is created automatically during the installation process. It is recommended that this string is modified on your site and copied to all servers in the cluster.

13. `DELAYED_INDEX_CREATE_RECORDS=100000`

This number specifies the maximum of records in a table that can create index records immediately after the “create index” command is applied to an existing table. This configuration is to prevent system from being swamped by the index creation operation in case there are a large number of records already in a table.

14. `AUTH_USER_PERM=no`

It tells whether system should check user permissions (grants) on data operations. Default is no. Please see "help grant;" "help revoke;" help menus.

15. `WAL_LOG=yes`

This config parameter tells the system whether to save Write Ahead Logs. It should be yes unless you use the server as pure memory cache.

16. `SEQ_READ_SPEED=200`

This number is the estimated sequential read speed in MB per second for SSD or HDD on your system. It is for optimized data storage layout plans.

17. `SEQ_WRITE_SPEED=150`

This number is the estimated sequential write speed in MB per second for SSD or HDD on your system. It is for optimized data storage layout plans.

18. `MAX_FLUSH_TIME=1`

This parameter is the longest interval that can be tolerated in data write operations, in seconds. If users can tolerate a longer interval, this number should be increased. It is a trade-off between this number and read performance. If this number is low, then data storage is less optimized for read operations. If this number is high, then read performance can be enhanced.