



СОВРЕМЕННЫЕ ИНФОРМАЦИОННЫЕ ТЕХНОЛОГИИ

*Системное и программное
обеспечение информационных технологий*

A. O. Volshtein

(EE «F. Skorina GSU», Gomel)

USING IN-MEMORY DATABASE FOR OFTEN ACCESSED UNSTRUCTURED DATA

В работе речь идет о использовании Redis для неструктурированных данных и её преимуществах относительно других средств.

Redis is an open source, in-memory data structure store. Redis can be used as a database, cache and message broker. As far as it stores data in memory instead of storing it on disk, Redis is much more efficient than on-disk databases which makes it very useful for data which require to be often accessed, written or read.

As for other database solutions to compare we have two major types of databases: SQL and NoSQL databases. If one needs a simple solution then the most popular SQL database would be SQLite. It is serverless which makes it much easier to deploy an application. But SQLite store information in a file, not memory, which gives it an overhead of limitation of filesystem. The problem of SQLite and all SQL databases is that SQL databases has limited abilities of storing unstructured data.

As for NoSQL databases, Redis is not the most popular one. The most popular one is MongoDB, which has its advantages and disadvantages comparing to Redis. MongoDB is a document-oriented database when Redis is a bit simpler Key-Value database. This means that with Mongo you can have complex data structures while with Redis only a few different databases. Plus, again, MongoDB stores data on disk which gives us an overhead and can be critical when speed is critical.

Redis as a key-value store is the best solution among listed, especially when efficiency and speed matters. Redis supports various data structures such as strings, hash tables, lists, sets, bitmaps, geospatial indexes.

Материалы XXI Республиканской научной конференции студентов и аспирантов «Новые математические методы и компьютерные технологии в проектировании, производстве и научных исследованиях», Гомель, 19–21 марта 2018 г.

These structures support atomic (uninterrupted) operations such as appending to a string, incrementing the value in a hash, pushing an element to a list, computing set intersection, union and difference or getting the member with highest ranking in a sorted set. Support of such operations means that those operations are thread-safe.

As for other advantages of Redis: it is an open source solution which has an active community, it is simple to install and has no dependencies. In order to achieve its outstanding performance, Redis works with an in-memory dataset, but depending on your use case, you can persist it either by dumping the dataset to disk every once in a while, or by appending each command to a log. Persistence can be optionally disabled, if you just need a feature-rich, networked, in-memory cache.