

Fundamentals of Database Systems

[NoSQL]

Malay Bhattacharyya

Assistant Professor

Machine Intelligence Unit
Indian Statistical Institute, Kolkata

October, 2019

1 Basics

2 Advantages and Challenges

3 Some Concepts

4 Implementations

Limitations of relational databases

Relational databases are entirely dependent on relations

- Relational databases cannot deal with data that are infeasible of fitting to relations.
- They cannot handle unpredictable, unstructured or messy data.
- They may require vertical and horizontal expansion of servers, as the data or processing requirements grow.

What are beyond relational databases?

Non-relational distributable databases depending on unstructured data. It's high performance with high availability, and offers rich query language and easy scalability.

CAP Theorem: Consistency, availability and partition tolerance.

NoSQL

Examples include - MongoDB, Hadoop, etc.

Note: Non-relational databases may not provide full ACID (atomicity, consistency, isolation, durability) guarantees, but still has a distributed and fault tolerant architecture.

Basics of NoSQL

NoSQL (often referred to as Not only SQL) is a broad class of data management systems where the data is partitioned across a set of servers, where no server plays a privileged role.

- NoSQL database management systems are useful when working with a huge quantity of data (especially big data), when the data's nature does not require a relational model
- The data can be structured, but NoSQL is used when what really matters is the ability to store and retrieve great quantities of data (e.g., millions of keyvalue in one or a few associative arrays)
- This is particularly useful for statistical or real-time analysis of growing lists of elements (e.g., huge number of Twitter posts, Internet server logs from a large group of users)
- It provides flexibility of the data model (e.g., ERP and BPM)

History

“It’s not about BASE vs. ACID”

– Emin Gün Sirer.

1998: Carlo Strozzi used the term NoSQL in 1998 to name his lightweight, open-source relational database that did not expose the standard SQL interface.

2009: Eric Evans, a Rackspace employee, reintroduced the term NoSQL to label the emergence of a growing number of non-relational, distributed data stores.

2011: Work began on UnQL (Unstructured Query Language), a specification for a query language for NoSQL databases, designed to query collections (versus tables) of documents (versus rows) with loosely defined fields (versus columns).

Note: BASE refers to basically available, soft-state and eventual consistency.

NoSQL – Advantages

- Elastic scaling – database can be distributed across multiple hosts as load increases (scale out)
- Handling big data – massive volume of data and transactions can be managed
- No DBA – the mechanism of automatic repair, data distribution, and simpler data models lead to lower administration and tuning requirements
- Low cost – clusters of cheap commodity servers are typically used to manage the exploding data and transaction volumes
- Flexible data models – schemaless relaxed (or even nonexistent) data model restrictions are used
- Caching layer – employs a useful cache layer
- Real time analysis and streaming model – distributed processing makes it faster

NoSQL – Challenges


- No normalization, referential integrity, ACID – basic criteria of relational databases are not satisfied
- Poor data security – distributed data becomes less secure and might get lost easily
- Immaturity – most NoSQL alternatives are in pre-production versions with many key features yet to be implemented
- Poor support – NoSQL systems are mostly open source projects without the global reach, support resources, or credibility
- Weakness in analytics and business intelligence – commonly used BI tools do not provide connectivity to NoSQL
- Targeting no administration – the goal of providing a zero-admin solution still has not been fulfilled
- Limited expertise – almost every NoSQL developer is in a learning mode

Taxonomy of NoSQL

- **Column:** Accumulo, Cassandra, HBase, etc.
- **Document:** Clusterpoint, Couchbase, MarkLogic, MongoDB, etc.
- **Key-value:** Dynamo, MemcacheDB, Project Voldemort, Redis, Riak, etc.
- **Graph:** Allegro, Neo4J, OrientDB, Virtuoso, etc.

JSON format

JSON (stands for JavaScript Object Notation) is a binary and typed data model that supports the data types list, map, date, Boolean, and numbers (with different precisions). It can be serialized to and deserialized from bytes as well as strings.




Some concepts in NoSQL – Versioning of datasets

To process “concurrent” modifications and versions of datasets in a distributed environment, different schemes can be employed. Some of these are listed below.

- **Timestamps:** It helps to develop a chronological order but rely on synchronized clocks and don't capture causality.
- **Optimistic Locking:** It implies that a unique counter or clock value is saved for each piece of data. When a client tries to update a dataset it has to provide the counter/clock-value of the revision it likes to update. It does not work well in a distributed and dynamic scenario where servers show up and go away often and without prior notice.
- **Multiversion Storage:** It means storing a timestamp for each table cell, thus providing optimistic concurrency control with compare-and-swap on timestamps.
- **Vector Clocks:** It is an alternative approach to capture order and allow reasoning between updates in a distributed system.

Some concepts in NoSQL – Vector Clocks

A vector clock is defined as a tuple $V_{i0}, V_{i1}, \dots, V_{in}$ of clock values (may be real timestamps, version/revision numbers or some other ordinal values) collected from each node $i = 1, \dots, n$, which represent the state of itself and the other (replica) nodes' state.



An overview of a vector clock managing six nodes

Some concepts in NoSQL – Vector Clocks


Vector clocks can be utilized “to resolve consistency between writes on multiple replicas”.

Vector clocks are updated in a way defined by the following rules:

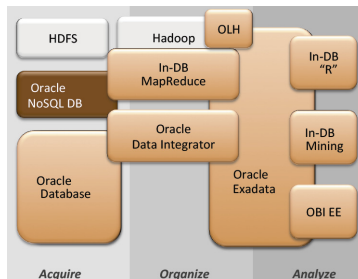
- 1** If an internal operation happens at node i , this node will increment its clock V_{ii} (immediate execution).
- 2** If node i sends a message to node k , it first advances its own clock value V_{ii} and attaches the vector clock V_i (about his internal state and his view of the other nodes) to the message sent to node k .
- 3** If node i receives a message from node j , it first advances its vector clock V_{ii} and then merges its own vector clock with the vector clock $V_{message}$ attached to the message from node j so that $V_i = \max(V_i, V_{message})$. To compare two vector clocks V_i and V_j in order to derive a partial ordering, the following rule is applied $V_i > V_j$, if $\forall k V_{ik} > V_{jk}$. If neither $V_i > V_j$ nor $V_i < V_j$ applies, a conflict caused by concurrent updates has occurred and needs to be resolved (e.g., by a client application).

Consistent hashing

Which object is mapped to which node is determined by moving clockwise around the ring. Therefore, if $N3$ leaves and $N5$ enters the system, the object 45 will get mapped to node $N5$.



Oracle NoSQL database



Oracle NoSQL database seamlessly integrates into an enterprise application architecture

Oracle-provided adapters allow the Oracle NoSQL database to integrate with a Hadoop MapReduce framework or with the Oracle database in-database MapReduce, data mining, R-based analytics, or whatever business needs demand.