ONTOTEXT ANSWER:

When interacting with the GraphDB API, you have two main options for building your application – the SPARQLRepository or the HTTPRepository. On the surface, both classes are configured in similar ways, they have the same functionality and similar performance. However, there are some differences that you should consider when choosing.

Flexibility – the SPARQLRepository works on any SPARQL endpoint – hence the name. There is no protocol check and no requirements. You can use it with a SPARQL 1.0 or a 1.1 endpoint. It is more straightforward, the repository only packages the query or update and lets the backend process it.

The HTTPRepository works only with RDF4J servers.

Functionality – the SPARQLRepository makes no assumptions about what the backend can do, it's an interface for sending queries and updates. By comparison, the HTTP Repository supports the full breadth of RDF4J API capabilities:

- Different isolation levels – how much of the transaction is visible within it? Can you view other transaction before they are committed?
- Namespace operations – adding, removing and fetching namespaces.
- Custom HTTP headers – useful for passing, for example, a request ID to audit.
- Two-phase commits – good for distributed inserts, as in SPARQL Federation.
- Advanced validation configurations – you can disable SHACL, or force it to validate the whole database.

Transactionality – both options support transactions. SPARQLRepository does not guarantee that the server itself is transactional. Isolation levels are only supported by HTTPRepository. SPARQLRepository does not support transactions for SPARQL updates, only for RDF data inserts and queries.

Performance – RDF4J has some performance optimizations, compared to other databases. With SPARQLRepository, we don't trigger these optimizations, since we don't make the assumption that the endpoint which we interact with is RDF4J-enabled. Since the two repositories optimize their queries differently, in some edge cases, SPARQLRepository may

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

have slightly better performance. However, in most situations, HTTPRepository would evaluate queries faster.

In conclusion, if you are certain you are using a RDF4J database, use HTTPRepository. If you are interacting with a database which you don't know, SPARQLRepository is a safer bet.

Did this help you solve your issue? Your opinion is important not only to us but also to your peers.

Help us get better: review GraphDB on Gartner

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com