Oxford Semantic Technologies

RDFox v7.7a, v7.6c, & v7.5d

Upgrade instructions

Patch releases v7.7a, v7.6c, and v7.5d have been issued to address a range of bugs. This guide provides upgrade instructions for the patches, including a general upgrade procedure and information on two of the fixed bugs that may have introduced faults that will or could persist even after upgrading. Descriptions of these issues and guidance for eliminating any such persistent faults are presented below.

Full release notes are available for v7.7a here, for v7.6c here, and for v7.5d here.

1. General

Each patch release is a drop-in replacement for the previous release in the same series, meaning that they are fully backwards compatible with the latter both in terms of persistence and APIs. As a result, the upgrade process involves simply restarting the RDFox server using the executable, dynamic library, or jar file from the later release.

2. Password hashing initialization bug (RD-2654)

The password hashing initialization bug is present in RDFox version v7.4 to v7.5c inclusive, v7.6 to v7.6b inclusive, and v7.7.

The bug occurs when an RDFox server is started without the `role-manager.argon2i.memory-cost`, `role-manager.argon2i.time-cost` and `role-manager.argon2i.parallelism` parameters. If these parameters a not specified at server initialization time, RDFox performs automatic optimization of these values and **should** store the resulting parameter values in the persisted server directory. However, under this issue, the optimized values are not stored, and thus subsequent RDFox instances use the default (minimum) values for these parameters.

Minimum values for these parameters weaken the security of password hashing, making brute-force attacks possible. Using such an attack, a bad actor could potentially recover plaintext passwords, thereby gaining access to those RDFox roles.

2.1. Securing an existing RDFox instance

For RDFox instances initialized before these patches, the steps provided in the RD-2654 Security Advisory should be followed to ensure that the server is not vulnerable to password cracking. Note that these steps can be followed either before or after upgrading.

3. Incremental aggregation bug (RD-2656)

RDFox versions before v7.7a, contain a bug that can lead to stale facts being leftover by incremental evaluation of rules that use aggregates.

The bug occurs when an RDFox datastore contains a rule that uses an order-sensitive aggregate function, which can cause slight variations in results based on the datatype used. For example, when a rule calculates an **xsd:double** value by applying the **SUM()** aggregate function to a list of **xsd:double** values, the result may

Copyright Oxford Semantic Technologies 2026

vary slightly depending on the order in which the values are accumulated, due to floating-point precision errors. Therefore, this potentially can lead to different facts derived on the same set of input values, causing stale facts to be left behind due to failing to retract old inferences during incremental reasoning

To address this bug, aggregate functions have been updated to support an *Aggregate mode*. This change introduces new optional keywords, RELAXED and STABLE, which allows users to specify whether the aggregate function enforces deterministic input ordering. Datalog rules will be STABLE by default.

This bug has been fixed only in the v7.7a patch release.

3.1. Removing stale facts during previous incremental reasoning

With the server running using the v7.7a executable or library, remove stale facts that were left behind due to previous incremental reasoning by rematerializing each data store that has ever contained rules that perform aggregation. This can be achieved with the remat shell command or an equivalent API call (see **recompute-materialization** here).

3.2. Performance considerations during reasoning

Within reasoning, STABLE mode is used by default to preserve correctness. This is important for order-sensitive aggregate functions, where processing values in different orders can produce different results and cause stale facts to remain after incremental reasoning. However, if an impact to performance has been observed during reasoning, datalog rules with aggregate functions can explicitly use RELAXED mode under certain constraints without violating correctness, please follow the guidance in the documentation.

4. Requesting additional support

If you encounter any problems following the guidance, or you have any questions about the server initialization or any of the other changes in v7.7a please contact support at Oxford Semantic Technologies via your usual channel.

Copyright Oxford Semantic Technologies 2026