

# Fuseki as alternate triple store

Although VIVO is configured to use SDB backed with an RDMS as the default triple store, VIVO comes with configuration examples for the following other triple stores as well:

- Jena TDB
- Virtuoso
- Remote triple store, connecting via SPARQL over HTTP

This documentation details how to configure Apache Jena's Fuseki 2 triple store, such that VIVO uses it as a remote triple store.

## Installation and Configuration

### Fuseki 2

The installation of Fuseki described here uses the standalone, executable JAR file. This is perfect for easily getting an environment up and running; however, for production, Fuseki's WAR file installation is likely more suitable.

1. [Download](#) Fuseki (using apache-jena-fuseki-3.14.0 in this documentation)
2. Extract the download package into a directory we will refer to as FUSEKI\_BASE
3. The following files/directories should be found in FUSEKI\_BASE

```
.
bin
fuseki
fuseki-backup
fuseki-server
fuseki-server.bat
fuseki-server.jar
fuseki.service
fuseki.war
LICENSE
NOTICE
README
webapp
```

4. Start fuseki:

```
$ java -jar fuseki-server.jar
```

5. By default, the fuseki service web console is available at: <http://localhost:3030>
6. In the web console, there is an option to "add a dataset"... do so, naming it 'vivo'
  - a. As the "Dataset type", select "Persistent". This uses TDB as Fuseki's backend triple store
7. Stop the fuseki-server
8. A new directory has been created within FUSEKI\_BASE, named "run" (see [Fuseki reference](#) for files contained within the "run" directory)
9. There should be a configuration file named "vivo.ttl" found within the "FUSEKI\_BASE/run/configuration" directory
10. Update "vivo.ttl" to use a union graph so that queries to fuseki do not need to specify named graphs

```
** Change **
:tdb_dataset_readwrite
    a          tdb:DatasetTDB ;
    tdb:location "/path/to/FUSEKI_BASE/run/databases/vivo" .

** To **
:tdb_dataset_readwrite
    a          tdb:DatasetTDB ;
    tdb:location "/path/to/FUSEKI_BASE/run/databases/vivo" ;
    tdb:unionDefaultGraph true .
```

11. Restart fuseki

## VIVO

These instructions detail how to configure VIVO to use the above, remote Fuseki triple store as an alternative to VIVO's default triple store configuration.

1. In VIVO's "applicationSetup.n3", make the following two changes
  - a. Uncomment and update the ":sparqlContentTripleSource"

```
:sparqlContentTripleSource
  a      vitroWebapp:triplesource.impl.sparql.ContentTripleSourceSPARQL ,
        vitroWebapp:modules.tripleSource.ContentTripleSource ;
  # The URI of the SPARQL endpoint for your triple-store.
  :hasEndpointURI "http://localhost:3030/vivo/query" ;
  # The URI to use for SPARQL UPDATE calls against your triple-store.
  :hasUpdateEndpointURI "http://localhost:3030/vivo/update" .
```

- i. Notice the "query" and "update" endpoints provided by Fuseki on the newly created "vivo" dataset
- b. Change the ":application" section to reference the ":sparqlContentTripleSource"

```
:application
...
  :hasContentTripleSource      :sparqlContentTripleSource ;
...
```

## Success

Restart VIVO... connected to Fuseki!