Upgrading VIVO

VIVO 1.10.x Documentation

Exported on 05/27/2018

Table of Contents

1 Overview

⚠ Due to the dependency updates in VIVO 1.10.x, it is important that you read the instructions here carefully, and plan your upgrade accordingly. You will need to consider:

1) The time it will take to upgrade the triple store (can not be performed against a running system)

2) SPARQL queries need to be checked for any explicit use of string datatypes

3) Any applications directly accessing the SDB triple store need to be upgraded to use Jena 3 libraries.

- This includes VIVO Harvester - if you use this, you must upgrade your Harvester to 2.x at the same time as VIVO. You must not use Harvester 1.x with VIVO 2.x, or Harvester 2.x with VIVO 1.x.

🛇 If you are planning an upgrade from anything prior to 1.9.x, please ensure that you read upgrade instructions relating to previous releases. If your version of VIVO is prior to 1.8.x, you will have some ontology changes to consider.

2 Java 8

With the upgrade to Jena 3.x, it is required that Java 8 is used. The Maven projects have been upgraded to state a dependency on version 8, and Maven will not run without it.

If you have previously been using Java 7, you will need to install Java 8. As usual, it needs to be a full JDK to support Tomcat.

⚠️ Java 9 has now been released, but has not been tested with VIVO. However, it does build and appear to run on Java 9. Note that to do so, you need to have a compatible version of Tomcat, as endorsed directory support has been removed.

3 Jena 3.x

As the core framework used throughout VIVO, as well as the implementation for the default triple stores, this is the most significant part of the upgrade (although, if you have extensive Javascript customisations in your front end, that side may take more time).

3.1 Upgrading The Triple Store

> ⚠ This needs to be done by everybody upgrading from a previous version, using the (default) Jena triple stores. It is necessary due to the changes in handling untyped literals in RDF 1.1. If you fail to perform the upgrade, or you mix Jena libraries
>
> If you are using an alternative triple store implementation (e.g. Virtuoso, Stardog), then you do not need to reload that triple store. But please remember that VIVO uses two triple stores - content and configuration - and if you are only using an alternative content triple store, you will still need to upgrade the configuration store.

Upgrading the triple store(s) (there are two - content and configuration) involves dumping the contents of your stores, and then reloading them. Whilst it is possible to do this using command line tools that are part of the Jena projects, VIVO has created and distributes tools that work with your VIVO configurations to make the process easier.

In order to upgrade your triple store, use the following steps (replace <your-settings.xml> and <vivo_home> with the appropriate values for your system:

1. Stop Tomcat - it is vital that Tomcat / VIVO, and any other applications that may access the triple stores, are not running during this process.

2. Run "mvn clean install -s <your-settings.xml>" in your VIVO 1.10.0 development area to update your web application and home directory
   This will install the tools into your <vivo_home>/bin directory.

3. Export your existing triple stores:

   > ⚠ • **To export successfully, you need to ensure no other programs are accessing your triple stores.**
   > • **Your file system must have the space available and be capable of storing files large enough to contain your entire triple store serialisation.**

   To export your triples store, use the jena2tools utility provided with VIVO 1.10.0, in <vivo home dir>/bin, specifying the export command, as shown below.

   ```
   java -jar jena2tools.jar -d <vivo home dir> -e
   ```

   Arguments:
        -d - the location of the Vitro/VIVO home directory
        -e - run in export mode
   On execution, the program will read your configuration files, find your Vitro or VIVO configuration within the vivo/vitro home directory, and get the necessary information to connect to your configuration triple store

(usually <vivo home dir>/tdbModels), and your content triple store (usually in SDB). If your triple store(s) are not SDB or TDB backed, then it will simply skip them.
jena2tools will then extract the contents of the available triple stores, and dump them to <vivo home dir>/dumps in TriG format.

> ⊘ In rare cases `jena2tools` will fail with a `java.lang.NullPointerException`. This can occur if VIVO was not properly shut down before upgrading.
> The preferred fix for this is to restore VIVO 1.9.x and see that it is properly started and shut down before upgrading. If this is not practical, another workaround is to delete the file named `tdbModels/journal.jrnl` in the Vitro/VIVO home directory. This may result in the loss of the most recent login information.

4. Check that the export has completed - you should have a <vivo_home>/dumps directory, that contains the files "configuration.trig" and "content.trig".

5. Empty your triple stores
   - Drop your database and recreate it as empty, just as you would for creating a new VIVO install. jena3tools must find an empty database (no tables) as named in your runtime.properties and will recreate your SDB triple store as tables in the named database using the triples produced by jena2tools and stored in <vivo home dir>/dumps/content.trig

   ```
   mysql> DROP DATABASE vitrodb;
   mysql> CREATE DATABASE vitrodb CHARACTER SET utf8;
   mysql> GRANT ALL ON vitrodb.* TO 'vitrodbUsername'@'localhost' IDENTIFIED BY
   'vitrodbPassword';
   ```

   - Delete all files in <vivo home dir>/tdbModels. Jena3tools will rebuild your configuration tdbModels based on the content created by jena2tools and stored in <vivo home dir>/dumps/configuration.trig
     rm -rf <vivo_home>/tdbModels

6. Reload your triple stores
   Having exported your triple stores, you can reload them using jena3tools, also available with VIVO 1.10.0, specifying the import command.

   ```
   java -jar jena3tools.jar -d <vivo home dir> -i
   ```

   Arguments:
       -d - the location of the Vitro/VIVO home directory
       -i - run in import mode
   On execution, the program will find your Vitro or VIVO configuration within the home directory, as well as the dumps that you have created with jena2tools. It will import them into the SDB and TDB triple stores, based on the configuration of your Vitro/VIVO instance.
   jena3tools will be present in <vivo home dir>/bin when you install the 1.10.0 beta. Alternatively, it can be downloaded from GitHub[1].

   Note that this can take a while. A rough guide is to expect about 600 triples per second to reload. (Roughly 1 hour per 2 million triples).

---

1 https://github.com/vivo-project/jenatools

7. Restart Tomcat

VIVO checks its filegraphs when starting. Restarting for the first time after an upgrade will take some time.

3.2 Upgrade Local Java Code Using Jena

If you have local customisations or additional applications that make use of the Jena libraries, you will need to upgrade these to work with Jena 3. Mostly this is simply a case of renaming any packages in imports for Jena classes:

import com.hp.hpl.jena.*

becomes

import org.apache.jena.*

However, some classes have been moved, or removed, and some interfaces have additional methods. So in rare cases you may find that you need to make a few small changes beyond this.

3.3 A Note on Other Dependency Changes

To remove the possibility of incompatible classes being loaded, and to remove known vulnerabilities from the code base, most of the Java dependencies in VIVO have been removed, updated or replaced.

For the most part, this will have minimal impact on local customisations.

Some libraries - such as commons-lang3 - have new package names, but mostly compatible classes, and usually just require the imports to be adjusted.

The CSV libary was outdated and unmaintainted, and has been replaced with commons-csv.

There were originally 5 JSON libaries (Jackson, Gson, Glassfish, Sourceforge.net and Json.org parsers). All code is now using Jackson, and the other parsers have been removed.

38 dependencies have been removed from the standard distribution. If you have any customisations that make use of them, that should work, but you will need to add the dependencies to your own projects.

22 dependencies will appear to have been added, but these are mostly from unbundling the owlapi OSGi dependency, and these - along with other conflicting classes - had been packaged as part of the bundle.

Whilst every effort has been made to eliminate known vulnerabilities, the Maven dependency-check plugin still reports 13 vulnerabilities across 8 libraries - for which these are currently the latest releases.

In total, 11 out of an original 113 dependencies are still at the same version as the previous release.

4 UI Changes

4.1 jQuery 1.12.4

Bootstrap based themes require a newer version of jQuery than was shipped with VIVO. In order for all of the functionality aross the UI to work, and to minimise the amount of duplication in the UI, it was necessary to upgrade all of the pages and plugins that used jQuery, so that the same upgraded version works across all of the pages and themes.

This release does require some migration of javascript that makes use of it. There is a script - jquery-migrate.js that helps with this, providing extra backwards compatibility, and logging warnings to the console whenever an old method is used.

All of the existing code that was relying on the migrate script has been updated, so that it is no longer needed.

To aid transition, VIVO still ships with the jquery-migrate.js script, allowing most existing code to work. You should monitor the Javascript console, and apply updates if you see any logged messages in your customisations - future versions of VIVO will remove this migration script in order to upgrade to later versions.

4.2 jQuery plugins

In order to support the upgrade to jQuery 1.12.4, and remove any backward compatibility messages from the Javascript console, the following plugins have been upgraded:

jQuery UI

jCrop

qTip

DataTable

In addition, the following plugins have been replaced with equivalents for compatibility and/or licensing issues:

Old	New	Notes
mb.FlipText	Jangle	Used for rotating text on the graphs (e.g. temporal graph)
isotope	wookmark	Used to render the three columns on the index page

4.3 D3 v4

This is another major upgrade that has architectural changes to improve modularity and make writing visualizations easier, as well as a selection of new features.

Where D3 was being included by VIVO (e.g. on the profile page, in co-authorship networks, etc.), these now include D3 v4, and the visualizations have been upgraded to use D3 v4.

The only visualization that has NOT been upgraded to D3 v4, is the Capability Map - as a full page visualization, that page is still including it's own D3 v3.

When upgraading, you have a few choices:

1) If you have a full page custom visualization, you can continue to specify whatever libraries and versions that you want to use, although if you are referencing the "shared" D3, you will need to adjust this to include your own D3 that is the version you require.

2) If you are embedding the visualization on a page that may have other visualizations, you should either:

a) Upgrade your custom visualization to use D3 v4.

b) Render your visualization into an iframe, so that you can specify your own javascript dependencies.

By making this change now, we are getting on to a maintained version of D3 (v3 has not had any releases for 18 months), and it prepares us for using D3.express[2] in the future, for more dynamic visualization building.

---

2 https://medium.com/@mbostock/a-better-way-to-code-2b1d2876a3a0

5 ORCiD API

It has been announced that ORCiD are looking to shut down their v1.x endpoints by the end of 2017. During development to support the new v2 api, it was recognised that the way the settings are currently configured should be simplified. As a result, in order for the ORCiD integration to work in VIVO 1.10, you will need to update your runtime.properties - regardless of which version of the api that you want to use.

The new settings look like this:

```
orcid.clientId = 0000-0000-0000-000X
orcid.clientPassword = 00000000-0000-0000-0000-000000000000
orcid.webappBaseUrl = http://localhost:8080/vivo
orcid.externalIdCommonName = VIVO Cornell Identifier
orcid.apiVersion = 2.0
orcid.api = sandbox
```

Setting	Values	Notes
orcid.clientId	<your client id>	The client ID that you obtained from ORCiD for API access
orcid.clientPassword	<your password>	The password the you obtained from ORCiD for API access
orcid.webappBaseUrl	<your VIVO website>	The URL for your VIVO site
orcid.externalIdCommonName	<site description>	The text that you want to appear in ORCiD for links back to VIVO profiles
orcid.apiVersion	2.0 (or 1.2)	The version of the API to use. It is recommended that you use 2.0.
orcid.api	release or sandbox	Use sandbox for development, release for production.

6 List View Configurations

To produce complex views of data associated with properties (e.g. a person's publication list), VIVO allows the association of externalized SPARQL queries, and associated Freemarker template links.

The configuration for this are the files called "listViewConfig-*" in the <webapp>/configs/ directory.

Due to the performance of OPTIONAL clauses for some triple stores, in VIVO 1.x the configuration files usually consist of a SELECT query (to retrieve the data for the associated Freemarker template), and one or more CONSTRUCTs (using UNIONs) to create a smaller model that the SELECT query would then be performed against.

This creates additional buffering of an intermediate model, the overhead of performing multiple queries, and makes the configurations harder to write and maintain.

For triple stores with poor OPTIONAL performance, this is a result of the OPTIONAL clause returning large amounts of data that is then joined to the rest of the query.

This can be avoided by making the OPTIONAL clauses contain the restrictions that they will be joined to, in addition to them appearing outside for the join.

As this is not necessary for all triple stores, an element <precise-subquery> has been introduced, so that these additional statements can be filtered out for triple stores where they aren't required.

So,

```
SELECT ?menuItem
       ?linkText
 HERE {
    ?subject ?property ?menuItem .
    OPTIONAL {
        ?menuItem display:linkText ?linkText .
    }
}
```

can be rewritten as

```
SELECT ?menuItem
       ?linkText
 HERE {
    ?subject ?property ?menuItem .
    OPTIONAL {
        <precise-subquery>?subject ?property ?menuItem .</precise-subquery>
        ?menuItem display:linkText ?linkText .
    }
}
```

```
When all OPTIONAL clauses are rewritten like this, the <query-construct>
elements can be removed, for approximately 10% performance improvement, and a
larger reduction in Java processing overhead. This allows the web server to
scale to handle more requests.
```

If you have customized listViewConfig-*.xml files, you do not need to rewrite them for VIVO 1.10 - they will work unmodified. However, you will have a small performance improvement, and more readable, maintainable queries, if you choose to modify them to take advantage of the new features.

7 Vocabulary Services

<< Upgraded web services >>

<< Registering for a UMLS key >>

8 Servlet 3.0 Upgrade

The web.xml that ships with VIVO has been updated to use 3.0 semantics (the required version of Tomcat already supported 3.0).

If you have customisations that introduce new servlets, then you can still add the configuration to web.xml, or you can modify the servlet to use @WebServlet annotations.

If you are replacing a servlet, then you will need to map the existing servlet to an unused url, and map the new servlet by adding configuration for these servlets to web.xml.

Alternatively, if you want to disable a servlet, then you can set the web.xml back to 2.5 spec, and add all of the servlet configuration explicitly to the web.xml.

9 Java Dependencies

9.1 HttpClient

Due to OSGi bundling of some of the core dependencies, VIVO had been shipping three different versions of HttpClient, all of which were incompatible with each other, and sometimes generated errors depending on which code paths got executed first.

VIVO 1.10 brings convergence around HttpClient 4.5.3, removing the problems caused previously. If you have any customisations that depend on HttpClient (or the fluent api - fluent-hc), please ensure that you are using the versions that are already indlueed with VIVO. This may require some minor adjustments to your client code to make it compatible.

9.2 OSGi Dependencies

The OWLAPI previously included with VIVO was an OSGi bundle, and included the classes of a number of libraries (such as HttpClient above). This causes classloading conflicts. In VIVO 1.10, we are depending directly on the libraries that were being bundled, rather than the bundle in it's entirety.

In addition, Jena has OSGi dependencies for HttpClient, leading to more duplicate classes. These have been explicitly excluded - see dependencies/pom.xml[3] for how this is done.

9.3 JSON Parsers

Four JSON parsers - GSON (com.google.gson), Glassfish (javax.json.Json, Sourceforge.net (net.sf.json) and JSON.org (org.json) have been removed, to simplify the code base, reduce vulnerabilities and improve performance.

All JSON parsing in VIVO is now handled through Jackson.

If you have local customisations that use a different JSON parser, you can add the dependency to your projects, although it is recommended that migrating to Jackson is preferable for long term maintenance.

9.4 Replaced Dependencies

The CSV parser has been replaced with commons-csv.

9.5 Removed Dependencies

The following dependencies have been removed from VIVO. If you have customisations that require any of them, you will need to add them as dependencies in your own projects.

agrovocws-3.0.jar, asm-3.1.jar, aterm-java-1.8.2.jar, axis-1.3.jar, axis-jaxrpc-1.3.jar, axis-saaj-1.3.jar, bcmail-jdk14-1.38.jar, bcprov-jdk14-1.38.jar, bctsp-jdk14-1.38.jar, c3p0-0.9.2-pre4.jar, cglib-2.2.jarcommons-beanutils-1.7.0.jar, commons-discovery-0.2.jar, cos-05Nov2002.jar, csv-1.0.jar, cxf-xjc-runtime-2.6.2.jar, cxf-xjc-ts-2.6.2.jar, dom4j-1.6.1.jar, ezmorph-1.0.4.jar, gson-2.5.jar, jai_codec-1.1.3.jar, jai_core-1.1.3.jar, JavaEWAH-0.8.6.jar, javax.json-api-1.0.jar, jjtraveler-0.6.jar, jsonld-java-jena-0.2.jar, json-lib-2.2.2-jdk15.jar, lucene-

---

3 https://github.com/vivo-project/Vitro/blob/develop/dependencies/pom.xml

analyzers-common-5.3.1.jar, lucene-core-5.3.1.jar, lucene-memory-5.3.1.jar, lucene-queries-5.3.1.jar, lucene-queryparser-5.3.1.jar, lucene-sandbox-5.3.1.jar, mail-1.4.jar, mchange-commons-java-0.2.2.jar, shared-objects-1.4.9.jar, sparqltag-1.0.jar, stax-api-1.0-2.jar, wsdl4j-1.5.1.jar

10 Location of runtime.properties

The preferred location of runtime.properties has changed from the <vivo_home> directory to <vivo_home>/config. VIVO will raise a warning on startup if runtime.properties is found in <vivo_home> or both <vivo_home> and <vivo_home>/config. Continue startup by refreshing or clicking continue. Move runtime.properties into the config directory to avoid the warning.

11 Preserving Customizations During Build

> (i) **Draft**
>
> This page is being written for 1.10 and Maven. Check back for more.

11.1 Overview

VIVO can be customized in many ways. Most sites customize the interface using themes. See Creating a custom theme[4]. Many sites create ontology extensions to introduce local terminology. See Adding Additional Ontologies to VIVO[5]. Sites can add custom configurations for the Data Distribution API (see), providing custom APIs returning data you specify in formats you specify at web addresses you specify.

In each case, files distributed with VIVO are replaced or added.

11.2 Maven Custom Installer

11.3 Example – Custom Theme

11.4 Example – Local Ontology

11.5 Example – Data Distribution API

---

4 https://wiki.duraspace.org/display/VIVODOC110x/Creating+a+custom+theme
5 https://wiki.duraspace.org/display/VIVODOC110x/Adding+Additional+Ontologies+to+VIVO