VIVO v1.8 Release Notes

VIVO v1.8 Release Notes

What is this document?

The VIVO 1.8 Release Announcement covers the main features of the release at a high level but leaves out a number of helpful specifics. This document is intended to help those planning to implement the release as either a new installation or upgrade to learn more about new features and find pointers to more complete documentation elsewhere on this wiki.

More modular architecture

The VIVO application is now partitioned into a main body and seven modules:

- content triple store
- configuration triple store
- search engine
- search indexer
- TBox reasoner
- file storage system
- image processor

The modules provide well-defined interfaces, encouraging alternative implementations from the VIVO community. Note that the main reasoner – the ABox reasoner – has not yet been partitioned into a module.

Performance improvements

The time required for a full re-inferencing has been dramatically reduced. The amount of improvement depends on several factors, but an order of magnitude improvement is not unusual.

Rebuilding the search index is faster in v1.8, due largely to reduced memory usage.

VIVO v1.8 also starts up in about half the time required by VIVO v1.7.

Interactive management of "faux" properties

VIVO v1.6 introduced the idea of faux properties, which provide greater specialization of information on profile pages. The faux properties were defined in a configuration file, and could not be easily changed.

VIVO v1.8 includes new pages for viewing and managing faux properties interactively.

Individuals represented by VCards

VIVO v1.8 supports the practice of representing some individuals by VCard information, rather than as a Persons. In this way, the data model is not overwhelmed by co-authors, co-investigators, and other persons who are incidental to the institution.

Triple store flexibility

VIVO can be configured to use several different types of triple stores:

- Jena SDB backed by MySQL (still the default configuration)
- Generic triple store that supports standard SPARQL query and update protocols.
- Jena TDB within the application
- Virtuoso via a SPARQL endpoint

VIVO has been extensively tested only with Jena SDB.

Added schema.org markup

The HTML markup in profile pages now includes "microdata" tags that are recognized by most popular search engines. These tags allow the search engine to easily extract structured data from VIVO, improving the appearance of search results.

A new customization guide

The VIVO documentation now includes a guide entitled ***Changing the appearance of VIVO***. Much of this information was already available in the VIVO wiki. It has been curated, updated, and augmented with new customization techniques.

This guide is available in the VIVO wiki, as a tree of 22 wiki pages. It is also included in the VIVO distribution, as a 118-page PDF document.

A new search index builder

The search indexer code has been re-written and re-designed, to make it more efficient, more maintainable, and more configurable. Most changes to the search index configuration can now be accomplished by changes to a runtime file.

GUI refinements

Restructure CSS and JavaScript files to reduce duplication and improve performance.

Improved "citation views" for publications and journals.

Reduce the time needed to load profile pages by delaying the QR code generation.

Improve the use of display limits and "functional" data properties.

Miscellaneous improvements

Compatibility with Java 1.8.

Security improvements to guard against cross-site scripting attacks and click-jacking.

Bug fixes.