

VIVO hardware and software requirements



Old

The information here is for older versions of VIVO and should be updated for current versions

There is no single preferred configuration for VIVO. Your needs will depend on anticipated size of your VIVO database and server load.

Server(s)

VIVO may be hosted on one or more physical servers, on virtual servers, or in the cloud. The three currently separable components are the main web application, the MySQL database (with the default Jena SDB triple store), and the Apache Solr search index.

Please add your site to this table!

Site	physical, virtual, cloud	app, solr, db, other	CPU cores, GHz, cache	memory (dedicated vs. burst)	OS	disk storage, MySQL DB size, triple count	average page hits per day /month	notes (e.g., ingest solution, whether allow editing, whether using HTTP caching)
Cornell #1 (VIVO 1.7)	virtual (local)	VIVO, HTTPD, Tomcat, Solr	2 / 2.67 / 12MB	16 GB	64b RHEL	< 150 GB disk		About 12,000 faculty/academics/staff; does allow editing but most done by VIVO staff & students
Cornell #2	virtual (local)	MySQL	2 / 2.67 / 12 MB	16 GB	64b RHEL	< 150 GB disk		
CU-Boulder (VIVO 1.7) 3 VIVO's	physical	Tomcat, MySQL, HTTPD, Solr	Intel 2 socket/8 core, 2 threads /core = 32 virtual CPUs	128 GB	64b RHEL	DB = 50MB; Disk = 600GB; 340K triples	roughly 100 page hits/day, 37K page loads in 2014	Full nightly loads / duration = 1 hour; 1700 people plus positions, degrees, and research (no publications yet). 3 VIVO instances on this server)
Scripps (VIVO 1.7)		VIVO, Tomcat, MySQL, Solr	4 cores	8 GB	64b CentOS 5.4	4.2 million triples		355 people, 31K+ publications (journal articles)
Weill Cornell (VIVO 1.7)	virtual		4 cores	8 GB				does use HTTP caching for non logged-in users – subsecond rendering of large profiles IF a page has been visited already
Memorial University #1 (VIVO 1.7)	virtual	VIVO, Tomcat, Solr, Drupal	4 cores / 2.4G / 12 MB	16 GB	RHEL 5	80 GB disk		still working on the data – 2000+ profiles, 2000-2500 projects, 500 knowledge mobilization ideas
Memorial University #2	virtual	MySQL	4 cores / 2.4G / 12 MB	16 GB	RHEL 5	80 GB disk		
Duke #1 (VIVO 1.5)	virtual	VIVO / Tomcat	2 VMs / 4 cores / 2.90GHz	8 GB	Scientific Linux release 6.6	25 GB NAS (shared space for photos)		No editing in VIVO. Faculty: 5,957
Duke #2	virtual	MySQL, Ingest Process, Solr	4 cores / 2.90GHz	12 GB	Scientific Linux release 6.6	~8 GB (MySQL data dir) ~2 GB innodb reports		
Agriperfiles (VIVO 1.9.1)	virtual	Tomcat, MySQL, Solr, Drupal	8 cores / 2 GHz	14 GB	64b OpenLogic linux /CentOS 6.3	500 GB disk		

MySQL

- MySQL 14.12 distribution 5.0.95 – very likely the RHEL standard distribution
- Database size – the largest reported in the [2014 VIVO Annual Survey \(response data\)](#) was ~5 GB (10+ million rows in the quads table)

Tomcat

- VIVO 1.7 is supported under Tomcat 7
- VIVO 1.8 is supported under Tomcat 7 or Tomcat 8
- sample server.xml
- setting memory parameters
- other Java Servlet containers in use


- Glassfish (1 reported)

Java

- VIVO 1.7 requires Oracle's Java 7 (JDK 1.7)
- VIVO 1.8 requires either Oracle's Java 7 (JDK 1.7) or Oracle's Java 8 (JDK 1.8)

2014 Survey

The following chart and environment summary shows the range of reported hardware types, memory, and related technologies used.



Environment

- MySQL database (1 AllegroGraph)
- Tomcat (1 Glassfish)
- Apache web server
 - 2 sites using Apache mod_cache, 1 Squid
- 1 site using Amazon EWS
- 1 site doing load balancing
- 2 use other technologies (Ruby, .Net) in their stack to support editing