

## NDS: Running Dataverse 4.2.3 as Docker containers

NDS is exploring deploying data management tools under Docker (and eventually Kubernetes). We have a preliminary Docker-based installation of Dataverse 4.2.3 that includes both TwoRavens and now iRODS support, based on feedback from Odum.

The docker image files have been moved to a new repo:

<https://github.com/nds-org/ndslabs-dataverse>

### Overview

We've broken the installation out into six separate containers:

- Dataverse webapp/Glassfish: CentOS 7 base with Glassfish, dataverse.war and dependencies.
- R/Rserve: CentOS 7 base with all dependent packages.
- Solr 4.6.0: CentOS 7 base with Solr 4.6.0 and Dataverse schema.xml
- TwoRavens: Above R/Rserve base with Apache, rApache, and TwoRavens
- Postgres 9.3 (official image)
- iRODS 4.1.3: iRODS iCAT base image with custom bitcurator and archive rules provided by Odum/RENCI.

The images contain basic installed software. Any configuration that requires linking to other services (e.g., Solr, Rserve, Postgres) are run during container creation through entrypoint.sh.

### Observations

Things went generally smoothly, but changes were required:

- The current install scripts for both Dataverse and TwoRavens require human input and combine installation and configuration tasks.
- Both scripts were heavily modified to run silently (relying on environment variables).
- Installation tasks were moved to the Dockerfiles, where possible.
- 

### Questions/problems

- eclipselink.ddl-generation: In the container environment, containers can easily be brought up and down. The current persistence.xml seems to default to "create-tables" when the webapp is deployed. If a container is restarted, resulting in redeployment of the webapp, the system fails during table creation. We've temporarily worked around this by including a different persistence.xml if we know the database has already been initialized. It's quite possible that I've misunderstood the ORM here and there's an easy solution.
- Is there a dependency on Solr 4.6.0, or will any Solr work?