Streamline-connecting data to interactive-apps in CyVerse Discovery Environment via iRODS CSI Driver

Illyoung Choi
CyVerse /
University of Arizona
iychoi@arizona.edu

Sarah Roberts
CyVerse /
University of Arizona
slroberts@arizona.edu

Edwin Skidmore
CyVerse /
University of Arizona
edwin@cyverse.org

Nirav Merchant
CyVerse /
University of Arizona
nirav@arizona.edu

July 7, 2021
iRODS User Group Meeting

CyVerse Discovery Environment (DE)

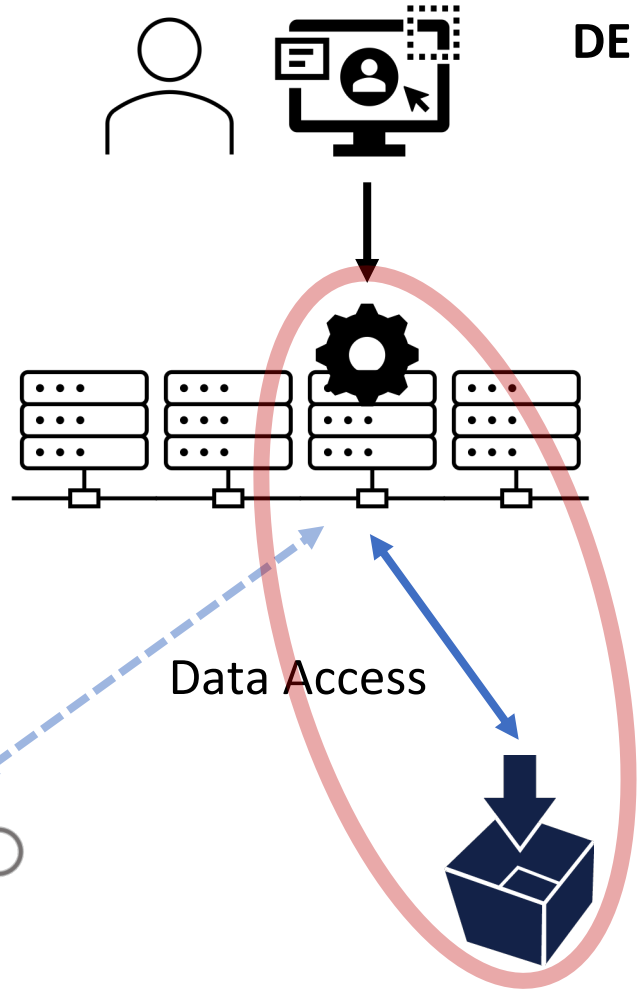

DE Apps

DE Cluster
- Kubernetes Cluster
- Computation

External Storage
(S3, FTP…)

Data Store
- iRODS Cluster
- User data & community shared data

Problems with traditional "*data staging*"

***Data Staging*: Download input to local disks and upload output to iRODS before and after running apps**

- Require a long time to launch apps with large input data
- Require disk space to store copies of large input data
- No access to data other than input

iRODS CSI Driver

***iRODS CSI Driver*: A Kubernetes Container Storage Interface Driver for iRODS**

- Create a Kubernetes *Persistent Volume (PV)* for iRODS
  - Mount iRODS data on Pods' filesystem using "***iRODS FUSE Lite***"
  - On-demand data transfer

Advantages

- Quick launching time of apps
  - About 20 seconds (Pod scheduling & initializing)
- Full access to Data Store
  - Input data of an analysis
  - Auto-created output directory for an analysis
  - User's home directory
  - Community shared datasets
- Immediate access to output data
  - Pipelining data processing

Integration to DE

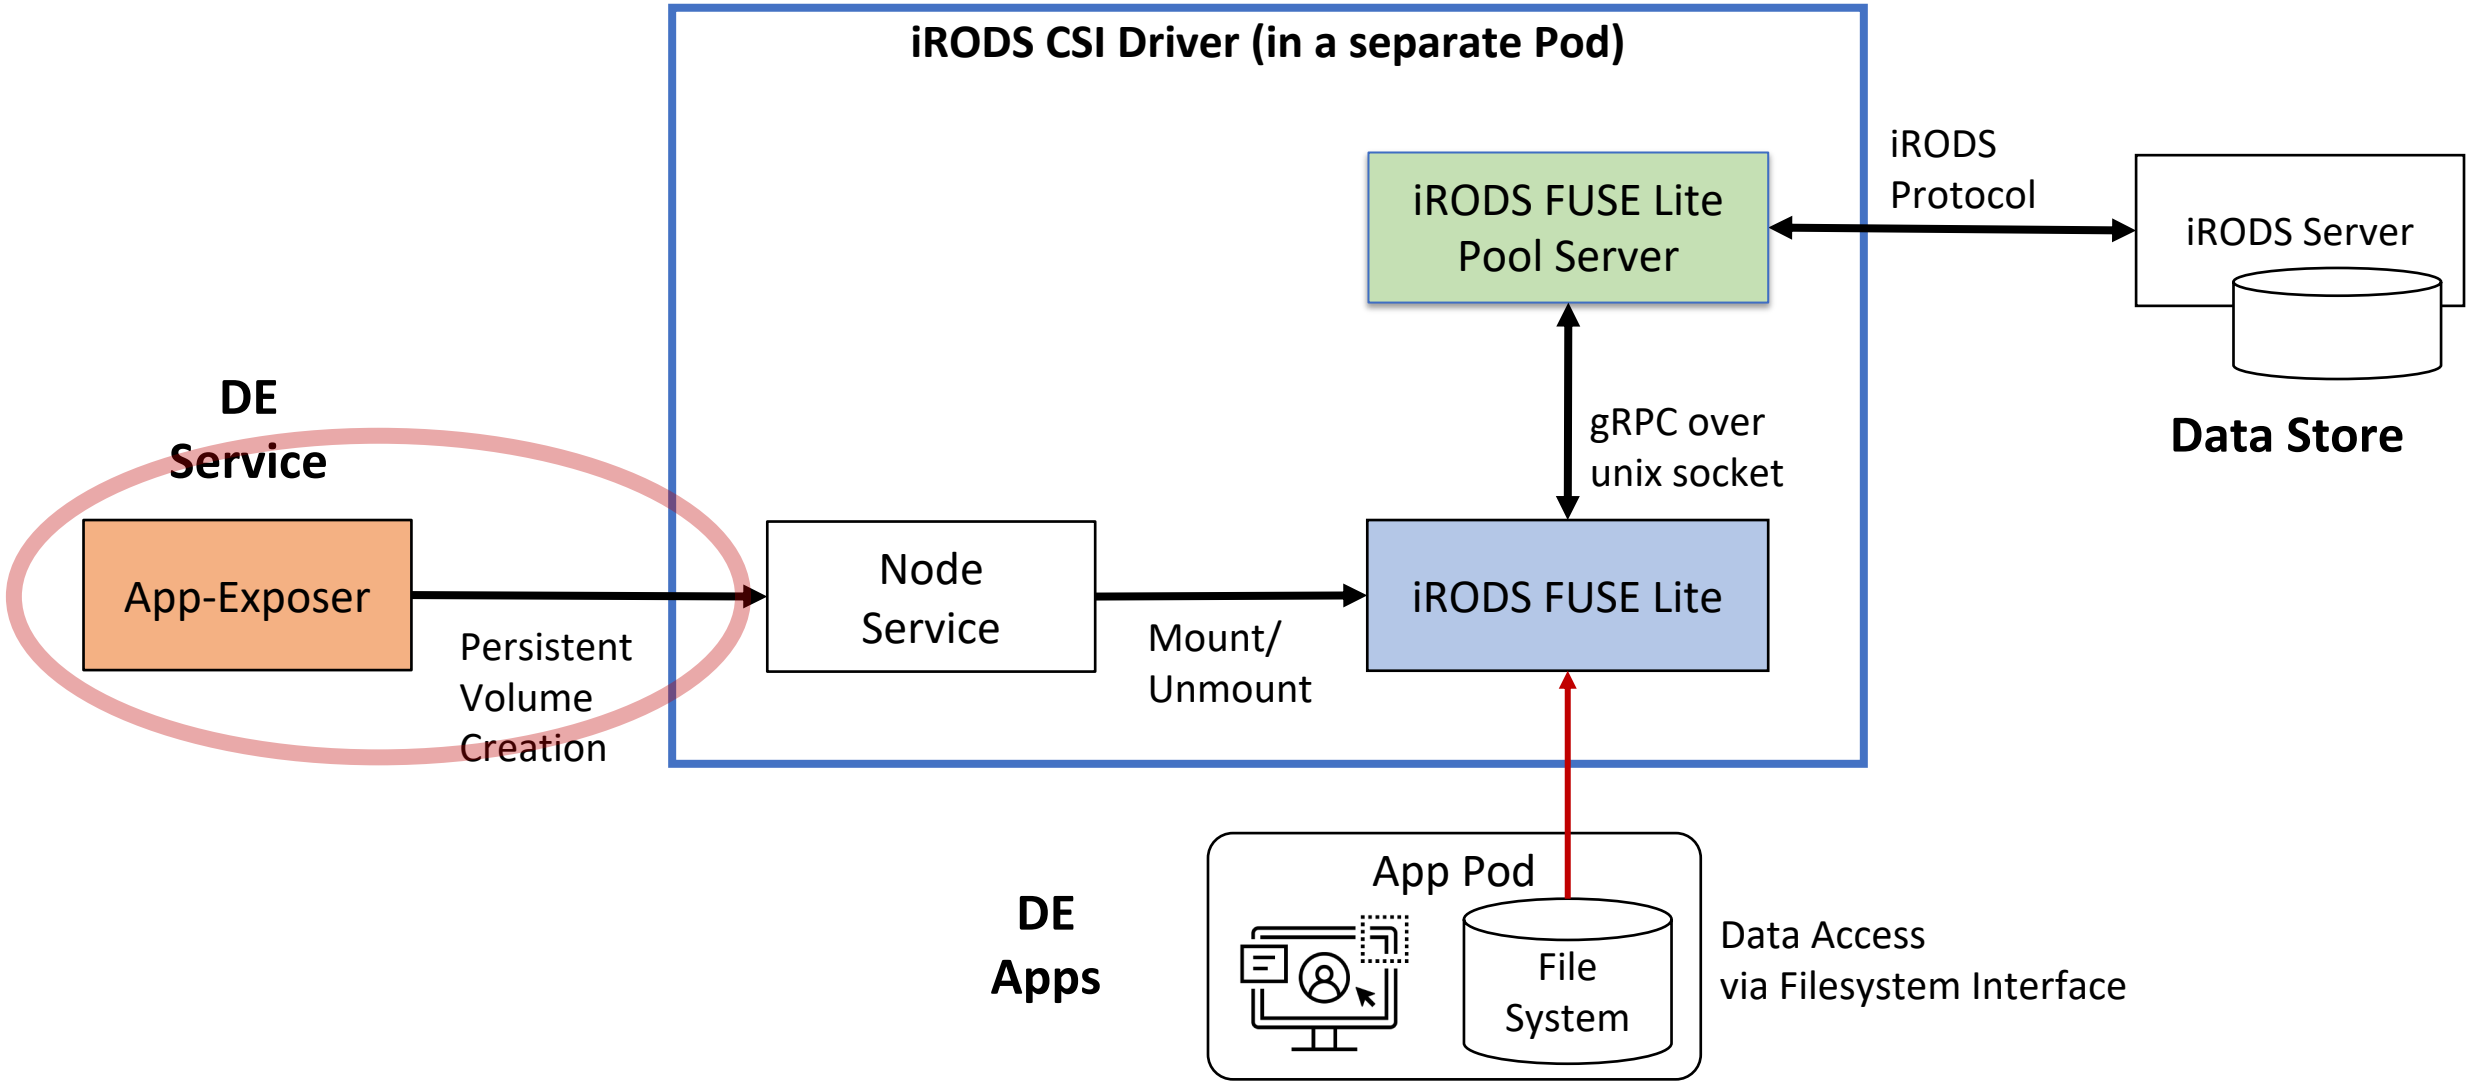

Integration to DE (Cont.)

- DE (App-exposer) creates/deletes ***PV*** *(Persistent Volume)* and ***PVC*** *(Persistent Volume Claim)*
  - Static volume provisioning
  - Proxy authentication (User password is not required)
  - 1 PV/PVC per app ➔ 1 **“iRODS FUSE Lite”** instance for reuse of iRODS connections and cache
  - Custom path mappings (+ different access permissions)

Mount Path	Permission	Description
/data-store/input	Read-Only	Input data of an analysis
/data-store/output	Read-Write	Auto-created output directory for an analysis
/data-store/iplant/home/<username>	Read-Write	User’s home directory
/data-store/iplant/home/shared	Read-Only	Community shared datasets

Troubleshooting

- Default work dir. of RStudio & JupyterLab
  - Difficult to access data out of the default work dir.
  - DE creates symbolic links under work dir.
- Create too many iRODS connections
  - Depletion of file descriptors in iRODS Server
  - **"*iRODS FUSE Lite Pool Server*"** manages iRODS connections

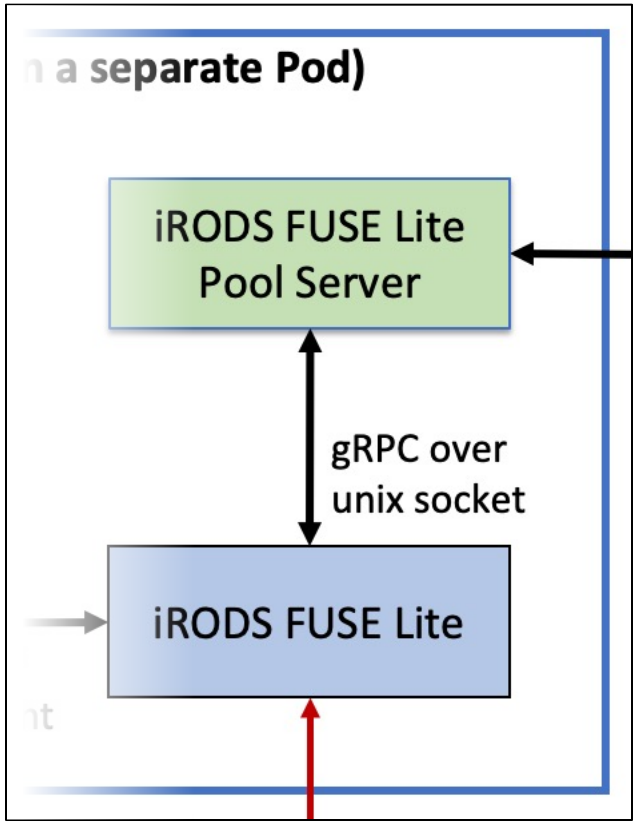

Troubleshooting (Cont.)

- “git clone” fails
  - “create - write - rename - close” pattern
  - A patch applied to **“go-irodsclient”**
    - Close the file handle when calling rename, reopen afterwards

- RStudio fails
  - Use of “ftruncate” or “O_TRUNC” to clear the file content
  - A patch applied to **“go-irodsclient”**
    - Close the file handle when calling ftruncate, reopen afterwards

Performance Optimizations

- Metadata access performance
  - Directory listing (listdir), file information retrieval (stat), ETC…
  - Adjust cache expiration time
    - User data: 5 minutes
    - Community data: 1 hour
- I/O performance
  - Asynchronous writing
  - Pre-fetching
  - Data caching

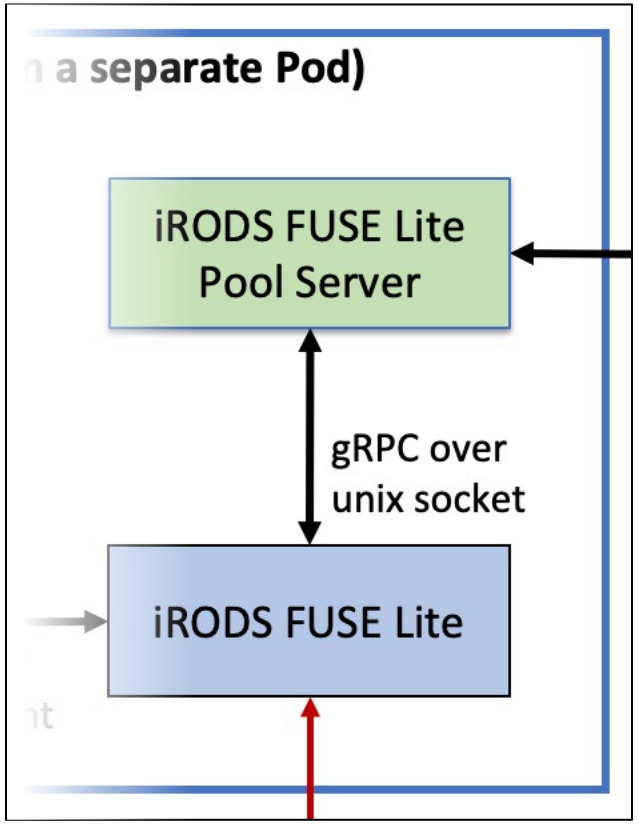

I/O Performance of iRODS Clients

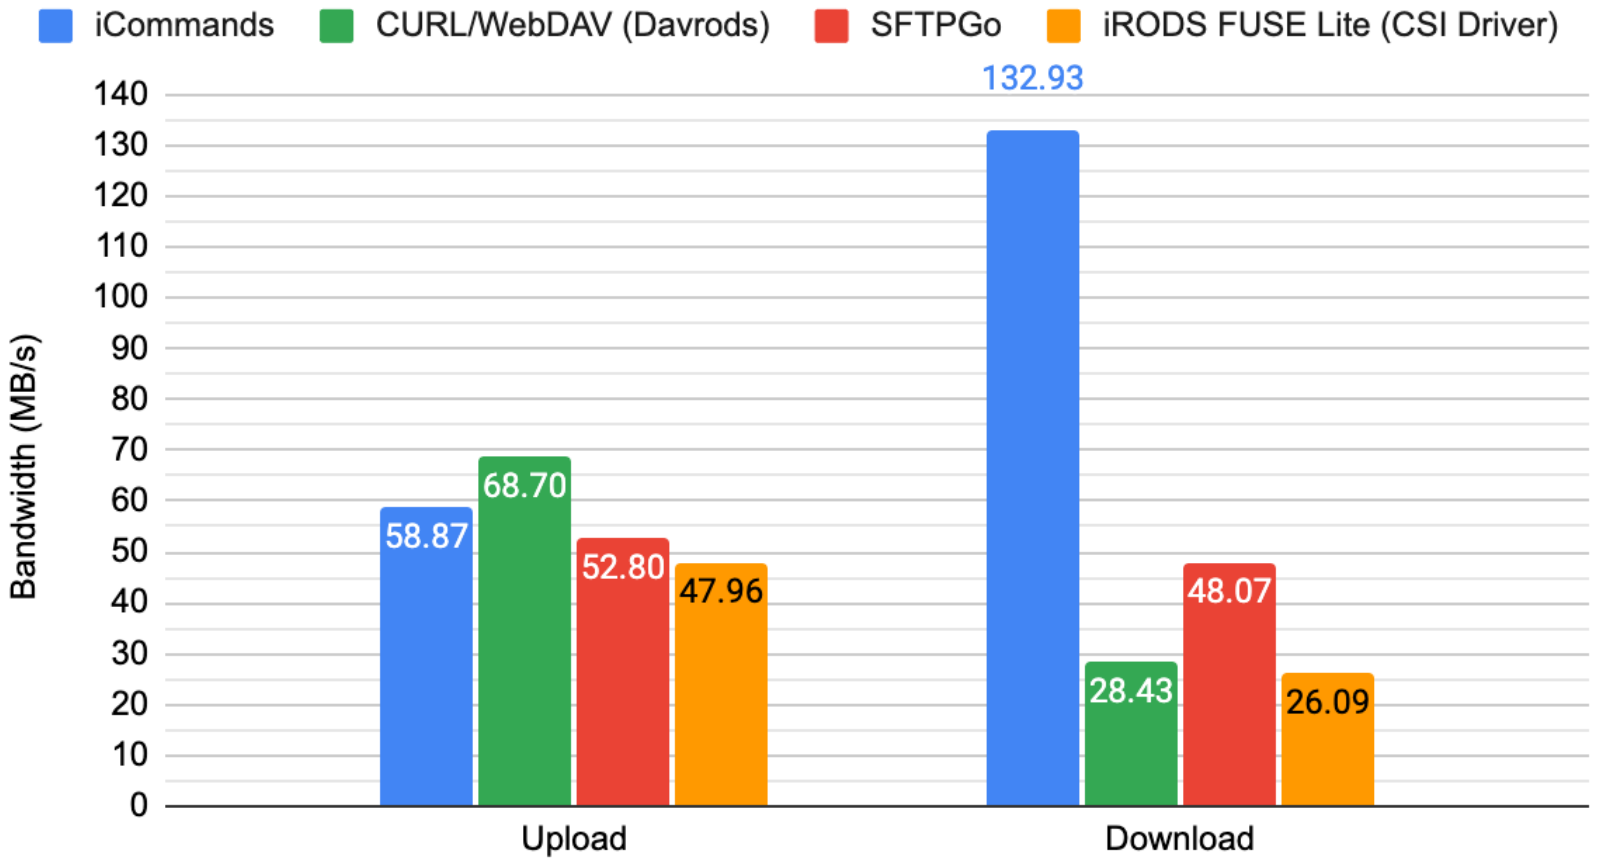

iRODS FUSE Lite (CSI Driver) transfers data on-demand.

Therefore, it's slower than other clients.

Open-source

iRODS CSI Driver

- https://github.com/cyverse/irods-csi-driver
- https://github.com/cyverse/irods-csi-driver-helm

iRODS FUSE Lite and Pool Server

- https://github.com/cyverse/irodsfs
- https://github.com/cyverse/irodsfs-pool
- File system clients using “go-irodsclient”

DE Integration (App-Exposer)

- https://github.com/cyverse-de/app-exposer
- PV/PVC creation & deletion

Conclusion

Production release

- Passed 6 months after release
- Users were satisfied with accessibility
- No critical issues reported

Troubleshooting & Optimizations

- Filesystem compatibility issues
- I/O performance issues

Questions?