



CYVERSE[®]

Keeping Pace with Science

The CyVerse Data Store in 2020 and the Future

Tony Edgin and Edwin Skidmore



THE UNIVERSITY
OF ARIZONA



Cold
Spring
Harbor
Laboratory

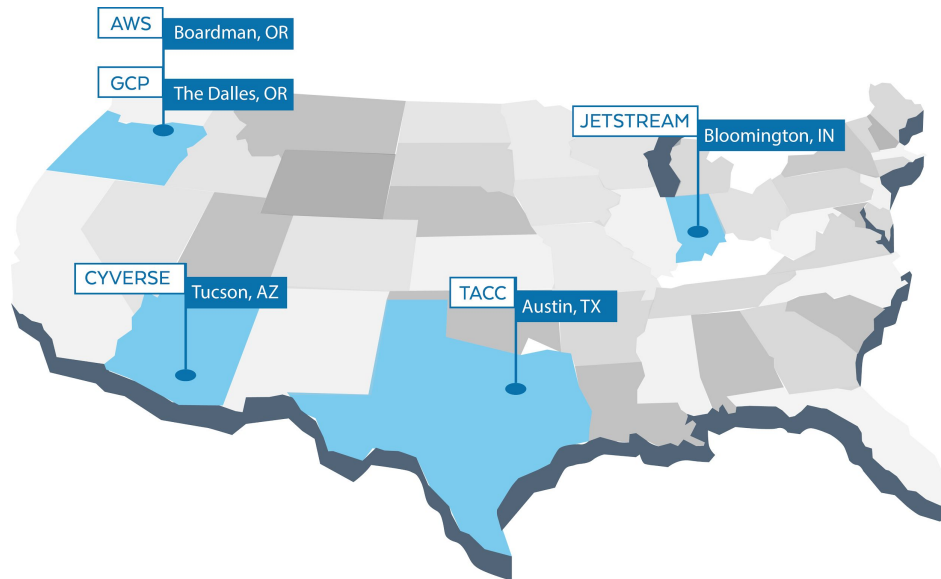
iRODS UGM 2020

Data Store Statistics

- ★ 190 million data objects (9 PiB)
- ★ 80 million files (5 PiB) transferred in 2019
- ★ 200 thousand files (14 TiB) transferred daily
- ★ 80 thousand users
- ★ 50 concurrent user connections on average

File Transfer Performance Between CyVerse and Various Compute Platforms

10 GiB File Transfer	
Computation Platform	Throughput (MiB/s)
Texas Advanced Computing Center (TACC)	170
Jetstream	330
Amazon Web Services (AWS)	240
Google Cloud Platform (GCP)	260



What is the CyVerse Data Store?

- Offsite replication
- **Optimization for accessing large sets of small files**
- Event publishing
- Customer-driven extensions
 - **Project-specific storage**
 - Service integration (see Appendix)
 - Custom application integration (see Appendix)



Optimizing Access to Large Sets of Small Files

Use Case

Datasets for genome browser, e.g., *JBrowse* or *UCSC Genome Browser*

- thousands of kilobyte-sized files
- browsers are interactive,
 - loads files as needed
 - must be responsive, i.e., cannot take 20 seconds for each user request

CyVerse Solution

Set up a WebDAV server with a file cache

- **apache** web server with
 - **davrods** for iRODS access
 - **modfilecache** to cache files
- separate virtual hosts for anonymous and authenticated access
- warm cache for byte-range access
- **100x faster than *iget***

Project-Specific Storage

Use Case

A project wants to store its data in the Data Store.

- 100 TB of data
- replicas stored locally at two institutions

CyVerse Solution

- project provides institutional storage servers
- CyVerse configures storage servers
 - catalog consumers hosting storage resources
 - project uses replication resource
 - policy to ensure data localities
 - separate iRODS service account

Data Store of Tomorrow

Steps toward utopia

- Increase interoperability
- Reduce accidental complexity
 - See [Your app makes me fat](#)
- Shorten scientific analysis feedback loop



Upcoming Features

- Thematic Real-time Environmental Distributed Data Services (THREDDS) (see Appendix)
- Bring your own (BYO) infrastructure
 - **BYO storage**
 - BYO compute (*later*)
- **Continuous analysis**



User-Provided, S3-Compliant Storage

Use Case

User wants to analyze their cloud data using CyVerse cyberinfrastructure.

- data hosted in an S3-compliant storage system, e.g. Google Cloud Storage
- moving them to Data Store is not feasible

CyVerse Solution

Use **iRODS S3 Resource Plugin** and **Filesystem Scanner**.


- *cacheless, detached* S3 resource for scalability
- Filesystem Scanner registers data in place
- Filesystem Scanner runs on cloud platform to avoid egress costs
- project owns cloud access credentials and responsible for accrued costs

More details in Appendix




CyVerse Continuous Analysis

CI/CD PIPELINE



Learning curves of some Container Orchestration Engines



Why “Continuous Analysis”?

- “Reproducibility of computational workflows is automated using continuous analysis”, CS Greene et al, Nature Biotechnology, June 2016 (<http://dx.doi.org/10.1038/nbt.3780>)
 - Used github and drone to demonstrate “continuous analysis”, ci/cd for science
 - Code and data changes -> re-execute analysis and version everything
 - Authors admit limitations in dealing with data sets, though not impossible
- Scientists and researchers want event-driven analysis (data growth, sensors data, etc)
- Containers are becoming the de facto standard as units of reproducible compute
- Kubernetes is becoming the de facto standard for orchestrating containers
- Container orchestration and CI/CD technologies are difficult to use, esp for a scientist and mortals who don’t know yaml (or json)

Why Continuous Analysis (cont.)

- Lessons learned
 - Jetstream/Atmosphere (multi-cloud, ad hoc interactive environments, allocations)
 - Containerized workflows
 - Data management
- Scientists need infrastructure to create, manage, and share in this emerging Kubernetes-native analyses in a managed fashion
- Complements the CyVerse's ecosystem, including Discovery Environment, Bisque, etc

Example User Stories

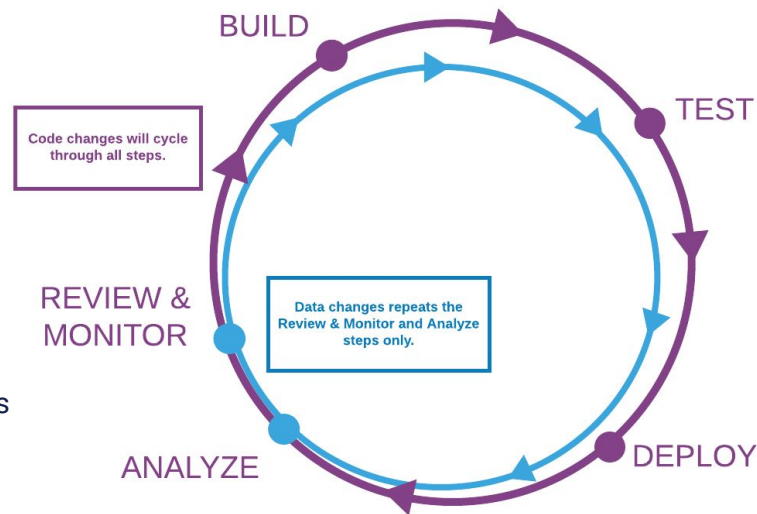
- I want my analyses to launch every time my workflow changes, my data changes, new ML training data is available, or every hour
- I want my analyses to always be “available” and only be “charged” for the resources I actually use
- I want to launch or transfer my analyses onto Jetstream/AWS/GCP/IoT/my own project’s servers
- I want to use Argo, Airflow, Snakemake, or Makeflow workflows with Kubernetes and scale as I define it

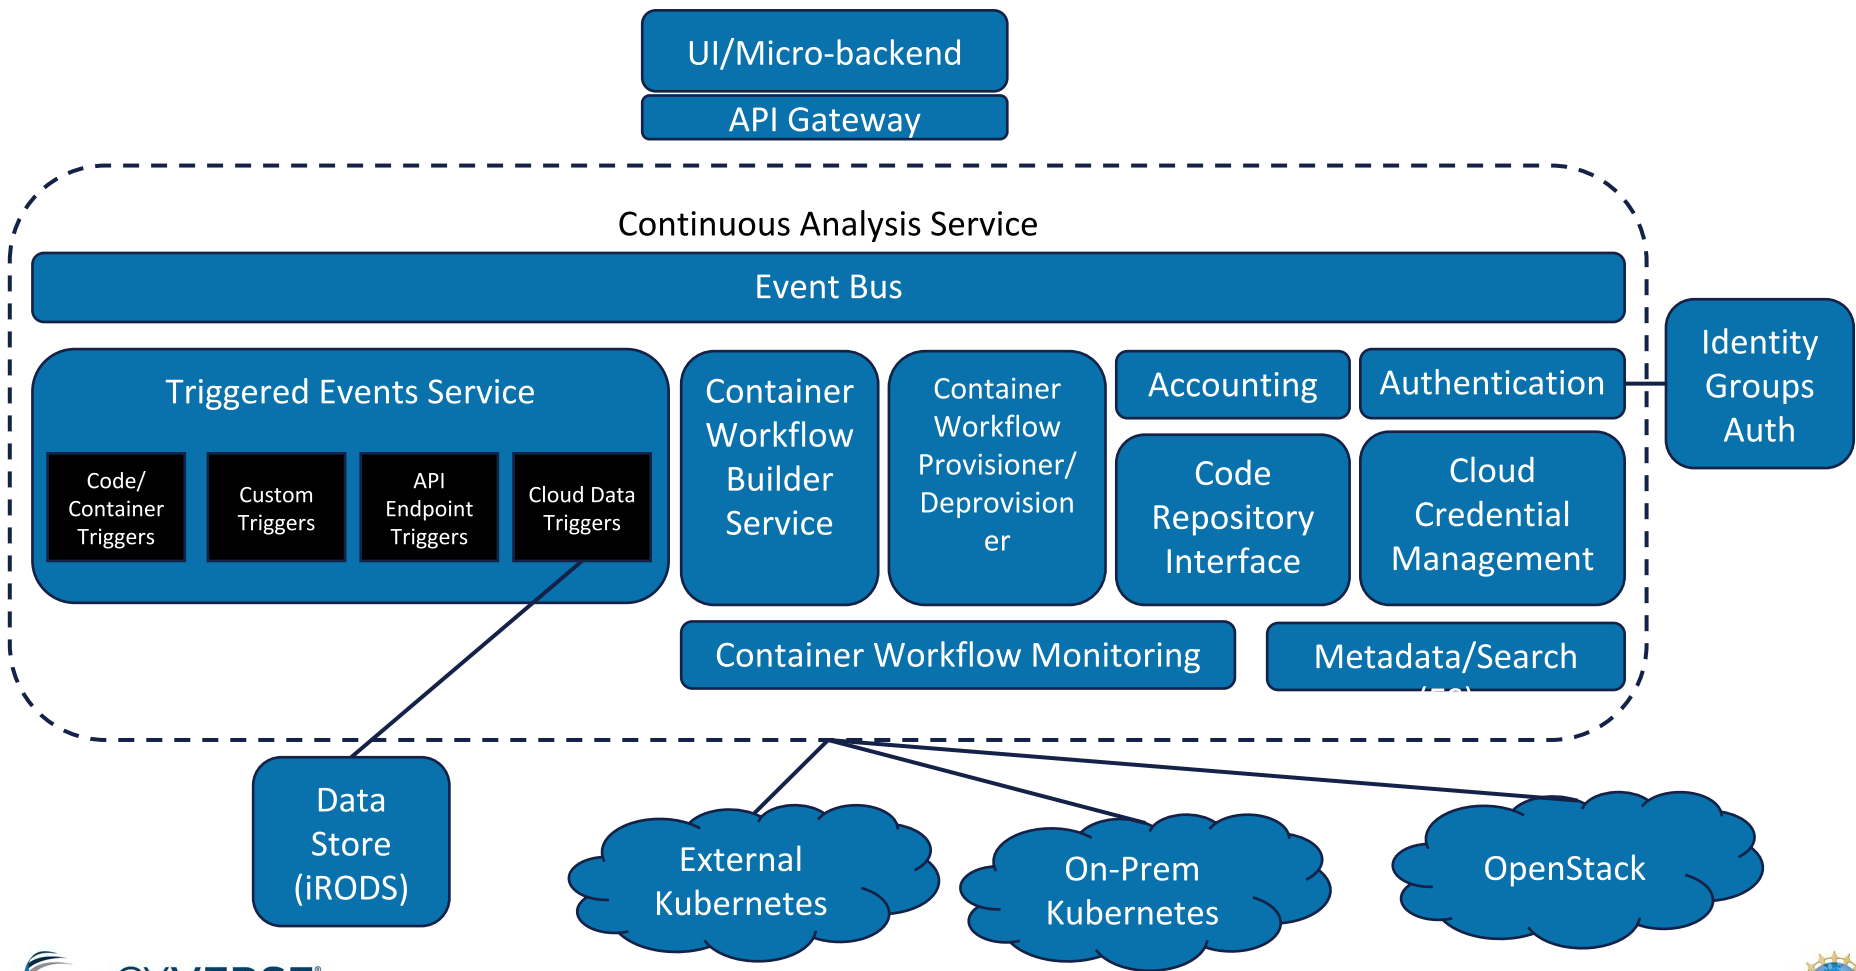
What is Continuous Analysis

Event-driven backend-as-a-service (BaaS) platform that will allow users to create, manage, deploy containerized analyses to any (kubernetes) cloud.

High level Capabilities:

- Multi-cloud (and iRODS integrated)
- Auto-scaling and Scale to zero
- Event-driven aka Continuous analysis (CI/CD for science)
 - Data events, workflow events, periodic, external events
- Kubernetes/Cloud Native
 - Custom Resource Definition (CRD)
 - Supports k8s CRD workflows: standard k8s, Argo workflows
- Git for workflow persistence
- Support for federated identity (via keycloak)
- CyVerse-features: api, sharing/permissions, interop, etc





Current Status

- Currently, in development
 - REST API is the initial focus (not so easy)
 - Command line interface (somewhat easier)
 - Easy to use UI
- Limited release in Q4 2020



Questions?

The slide features decorative curved lines in blue and grey in the top-left and bottom-right corners. The word "Appendix" is centered in a dark blue, sans-serif font.

Appendix

Service Integration

Use Case

A *Powered by CyVerse* service needs to access its users' data.

- not controlled by CyVerse, no admin access to data
- read-write access to *its* user data

CyVerse Solution

- service assigned rodsuser type iRODS account
- user opts into service through *User Portal*
- shared collection for user and service
 - owned by user in home collection
 - policy gives service write on contents
 - user has write permission, discourage delete, breaking service access

Example Application Integration

Use Case

[Sparc'd](#) is a desktop application supporting wildlife conservation created by Susan Malusa.

- manages sets of camera trap images
 - sizeable sets of small files
 - each set is tagged with metadata
 - supports sharing
 - images cannot be public, protect endangered species from poaching
- intended users are citizen scientists
 - volunteers, low frustration tolerance
 - require efficient uploads

CyVerse Solution

- project collection managed by Sparc'd creator who gets own permission on contents, enforced by policy
- “tar pipe” style upload
 - Sparc'd packs images in one or more tar files
 - asynchronous rule unpacks, registers images
- metadata attached in bulk
 - uploaded as CSV in each tar file
 - applied by image registration rule

THREDDS Support for NetCDF Data Sets

Use Case

A project uses NetCDF files to store its public data sets.

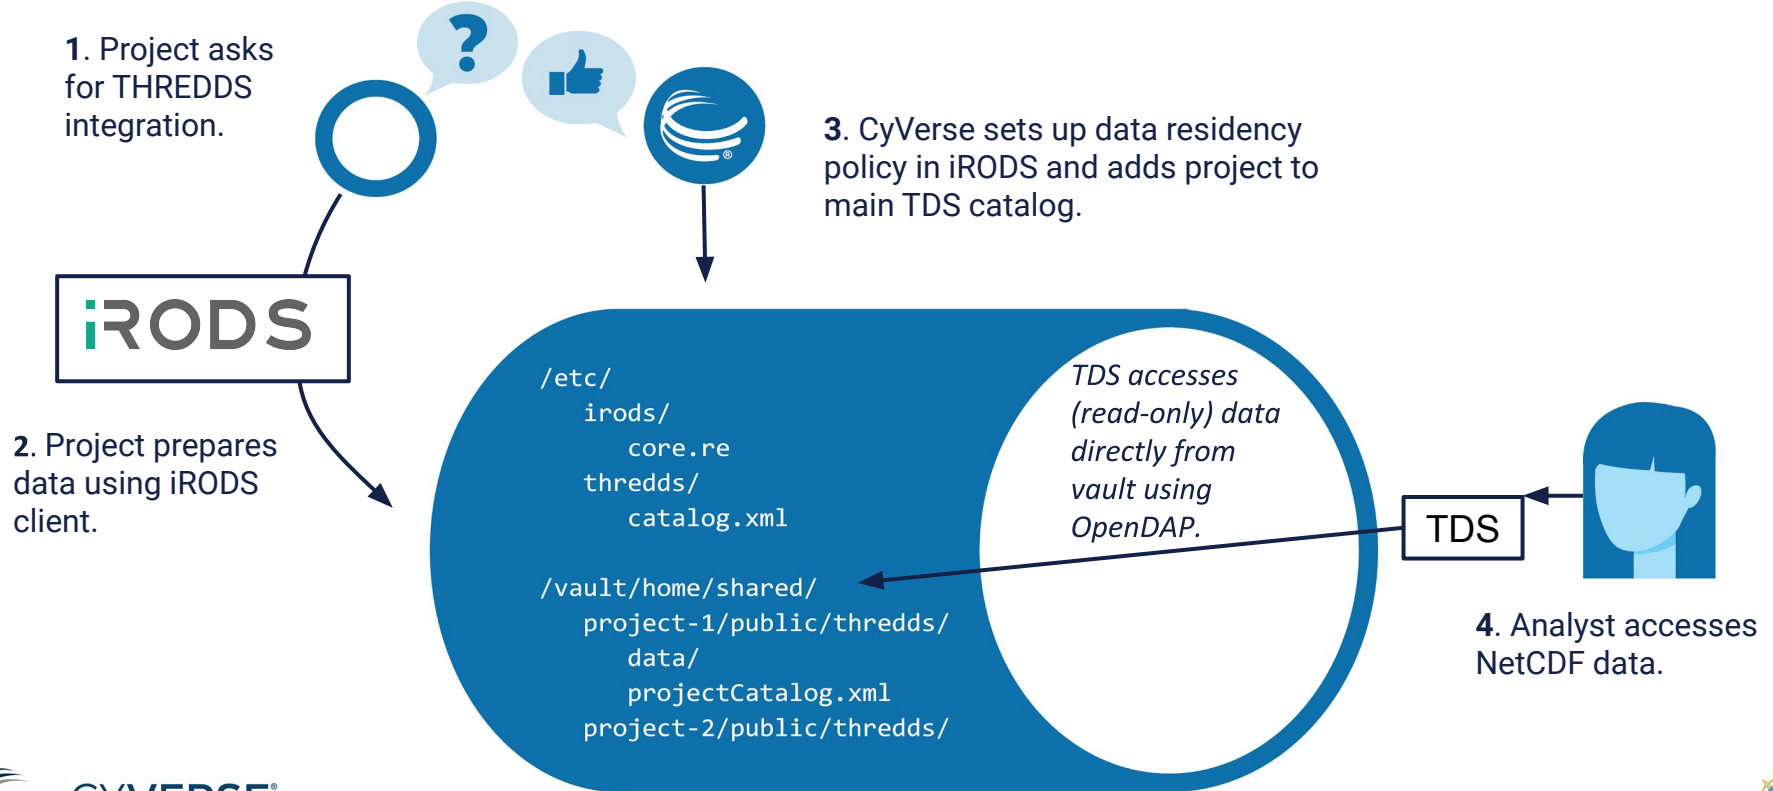
- files are multi-gigabyte sized
- only portions of some files needed at a time

CyVerse Solution

THREDDS Data Server (TDS) provides a collection of web services for accessing various types of datasets including NetCDF.

- iRODS resource server and TDS share host
- TDS has direct, read-only access to iRODS vault
- THREDDS data description files in vault
- project manages served data through iRODS
- analyst accesses data through TDS

THREDDS Integration Process



User-Specific S3 Resource Creation Process

