



ALPHA

CyVerse Tech Preview Webinar  
Edwin Skidmore, CyVerse Director of Infrastructure  
February 25, 2022

# Tech Preview

- Background
- Challenges of cloud
- Solution
- Demo
- Plans

# Then: Research in the Cloud 2008

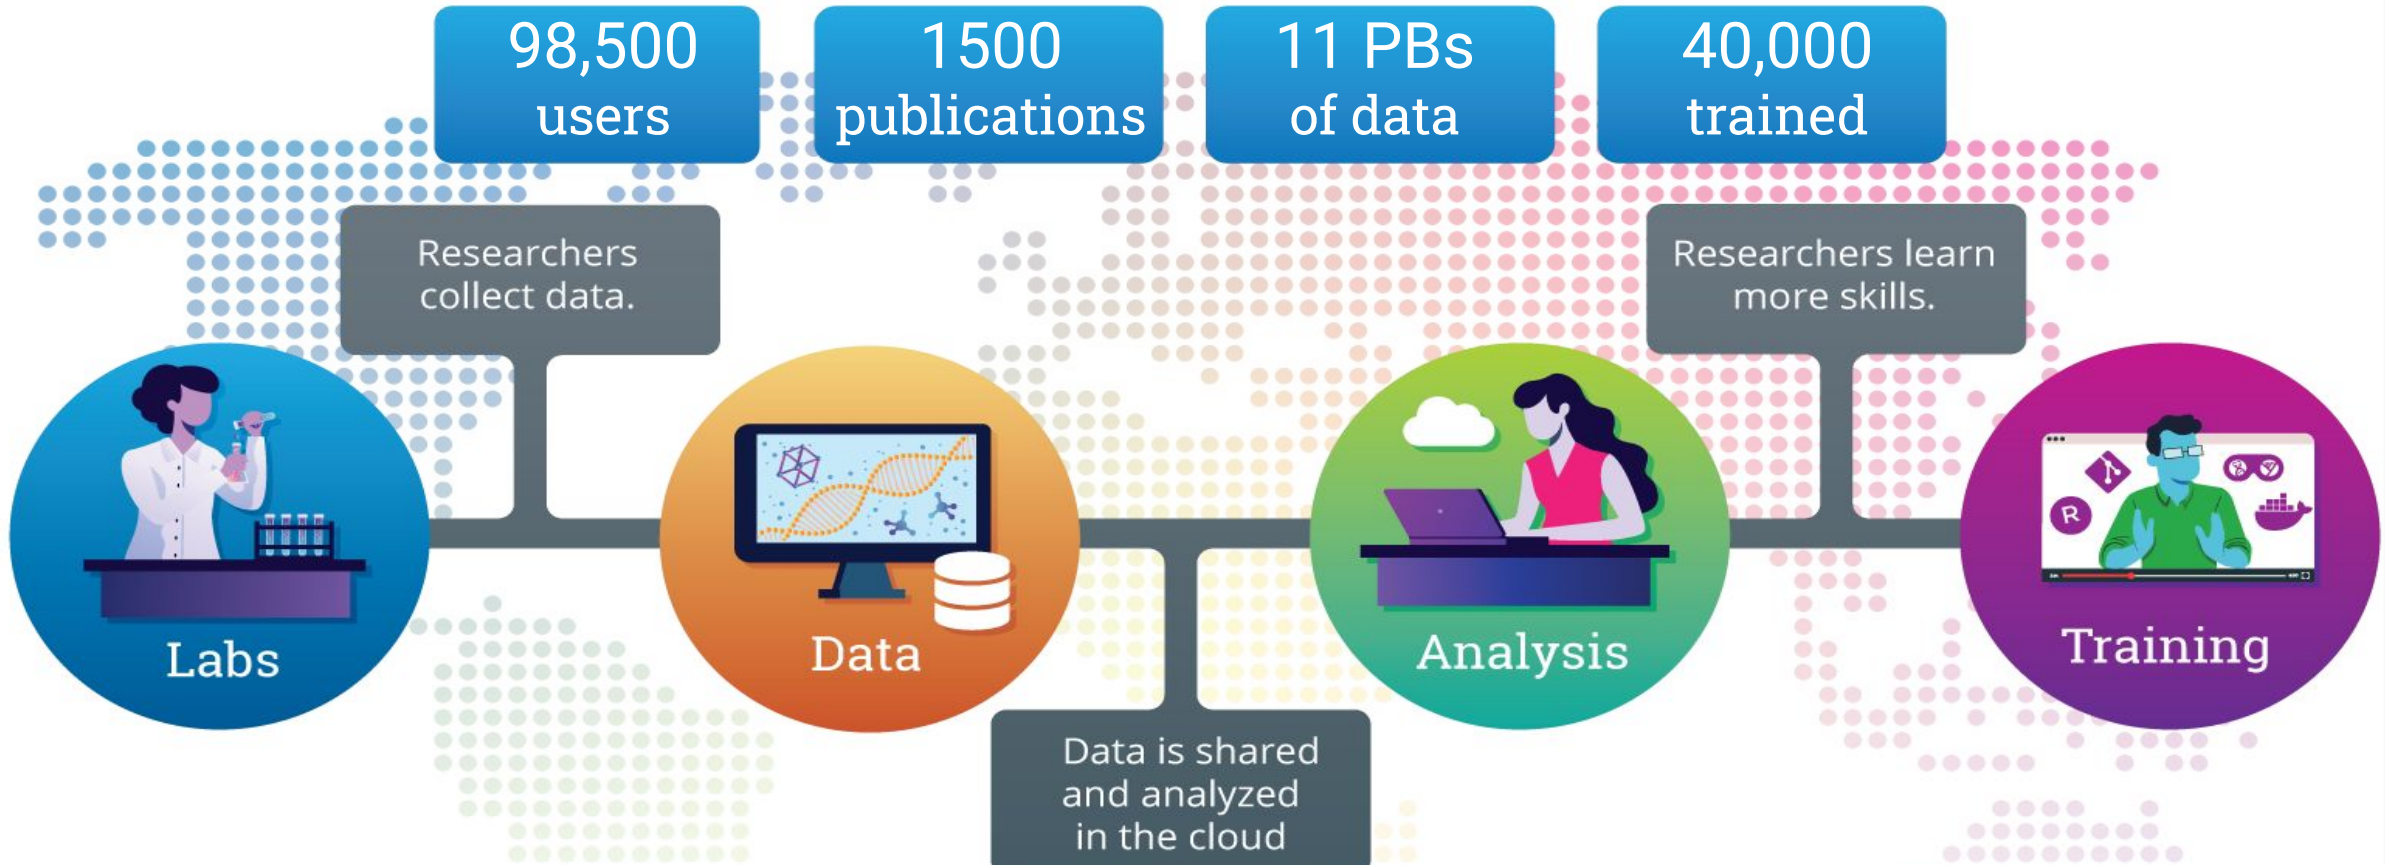
- Uncommon practice and access
- Difficult
- Expensive
- What scientific frameworks?
- CyVerse
  - Data Store (scalable, federated data storage)
  - Discovery Environment (scalable, distributed analytical and data management platform)
  - Atmosphere (scalable, multi/hybrid cloud)



EUCALYPTUS

# Now: Research in the Cloud 2022

- Common practice and better access
- Slightly easier
- Slightly less expensive to use
- Multi-cloud frameworks
  - Software stacks (e.g. web component with database backends, like jupyter or mlflow)
- Science in the cloud is more sophisticated



CYVERSE<sup>®</sup> Professional

- Federation with local and commercial cloud and high-performance computing
- Integration with local user identity management systems
- Security compliance

CYVERSE<sup>®</sup> Health

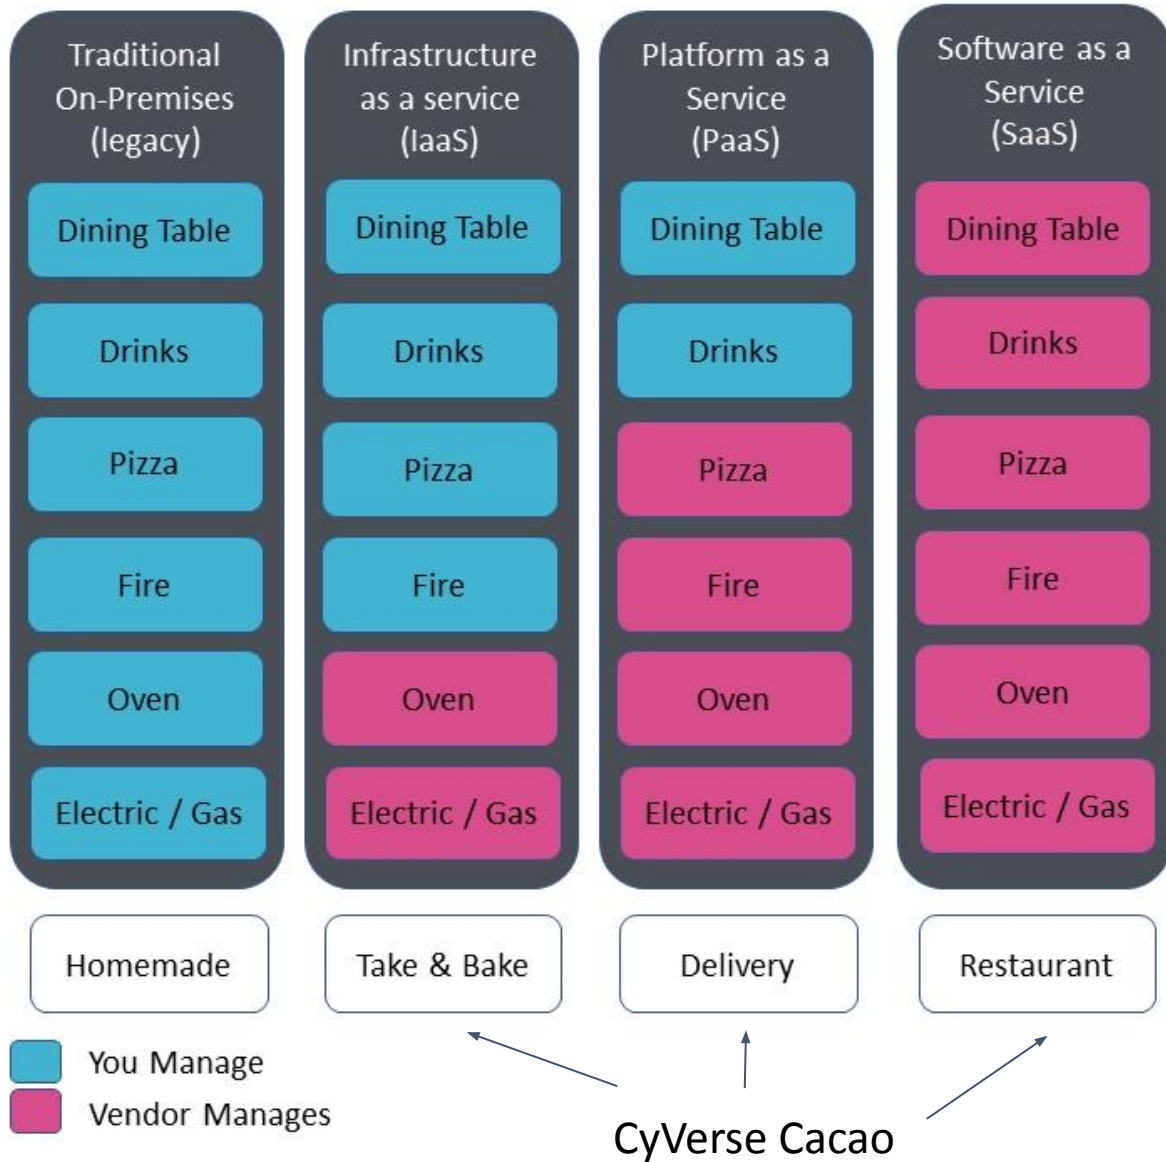
- Secure Perimeter (VPN)
- HIPAA Compliant Platforms inside of Perimeter
- XNAT
- Training

CYVERSE<sup>®</sup> Defense

- ITAR Compliance
- Receive data from multiple sensors
- Policy-based data visibility and sharing
- Support for multiple teams and data partitioning

# Challenges and Open Questions

- “Cloud” is both overloaded and very specific (vendor and type)
  - AWS, GCP, Openstack (Jetstream), Kubernetes
- Oftentimes science in the cloud is about being opportunistic
  - Cost containment
- Resources are scattered across multiple clouds
  - storage in one place, compute in another (how to aggregate?)
  - data proximal compute
- Multi-cloud? Hybrid cloud? Edge cloud and IoT?
- Event-driven?
- Multi-tenancy?
- Sharing?




# Pizza-as-a-Service

From Albert Barron. "Pizza As A Service", July 30, 2014.  
<https://www.linkedin.com/pulse/20140730172610-9679881-pizza-as-a-service/>

“order a red velvet 3-layer cake”

“order party supplies”




“reserve band”

“find caterer and order food”


“reserve venue”






"I want this theme"

"I want a red velvet 3-layer cake"



"order party supplies"



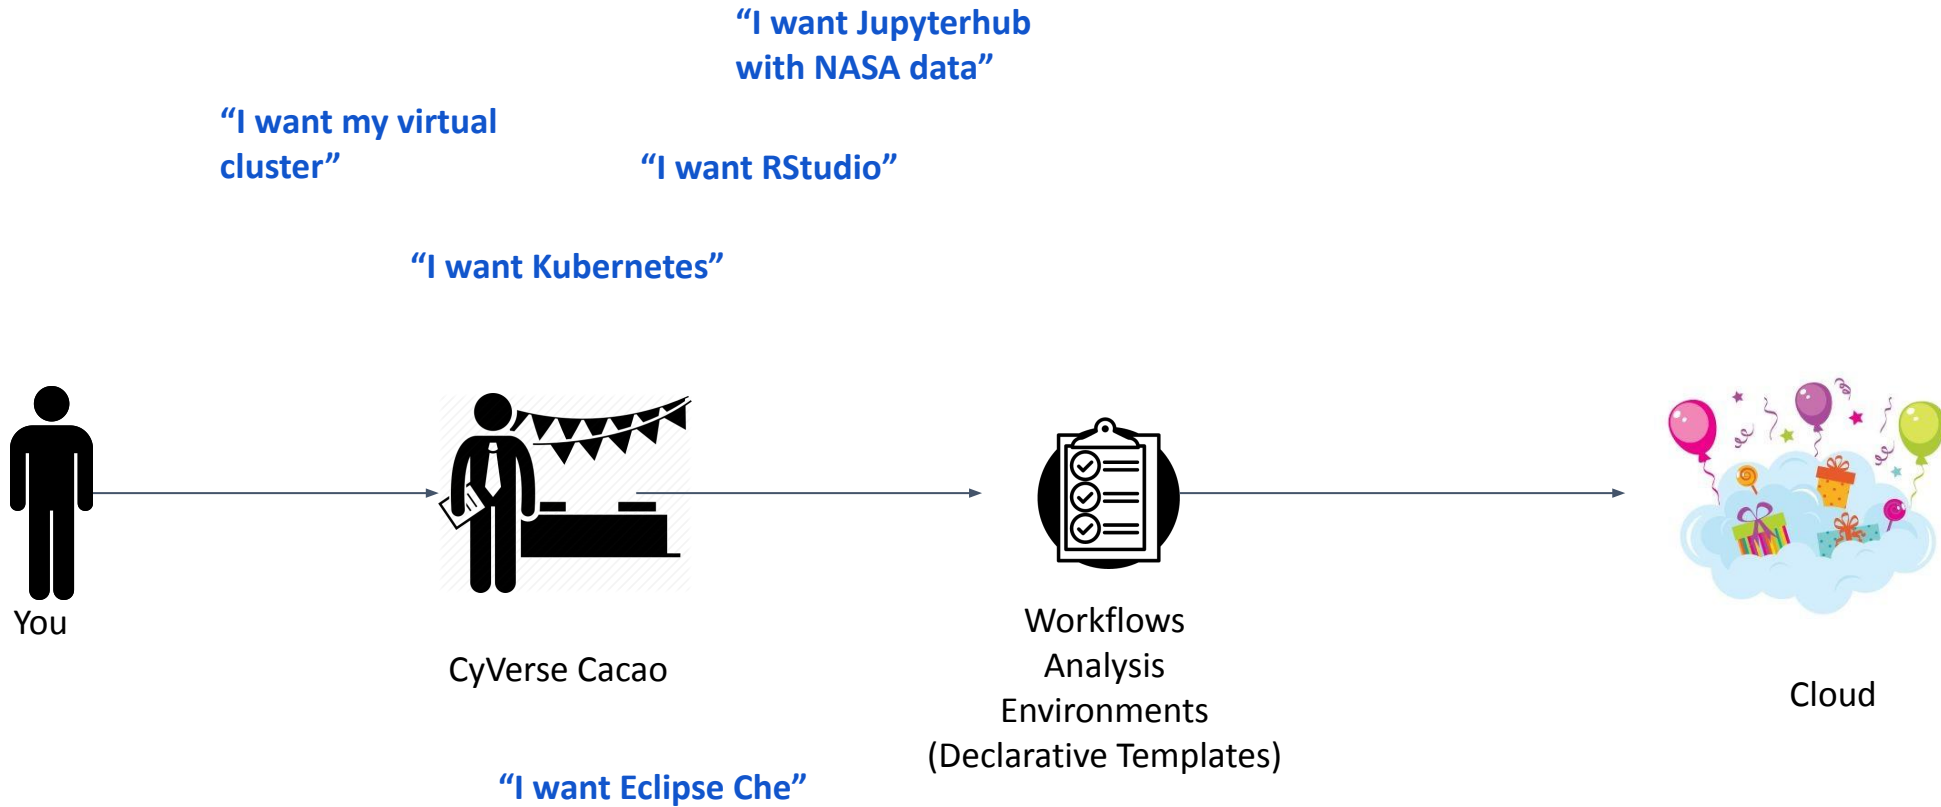
"I want this type of music"



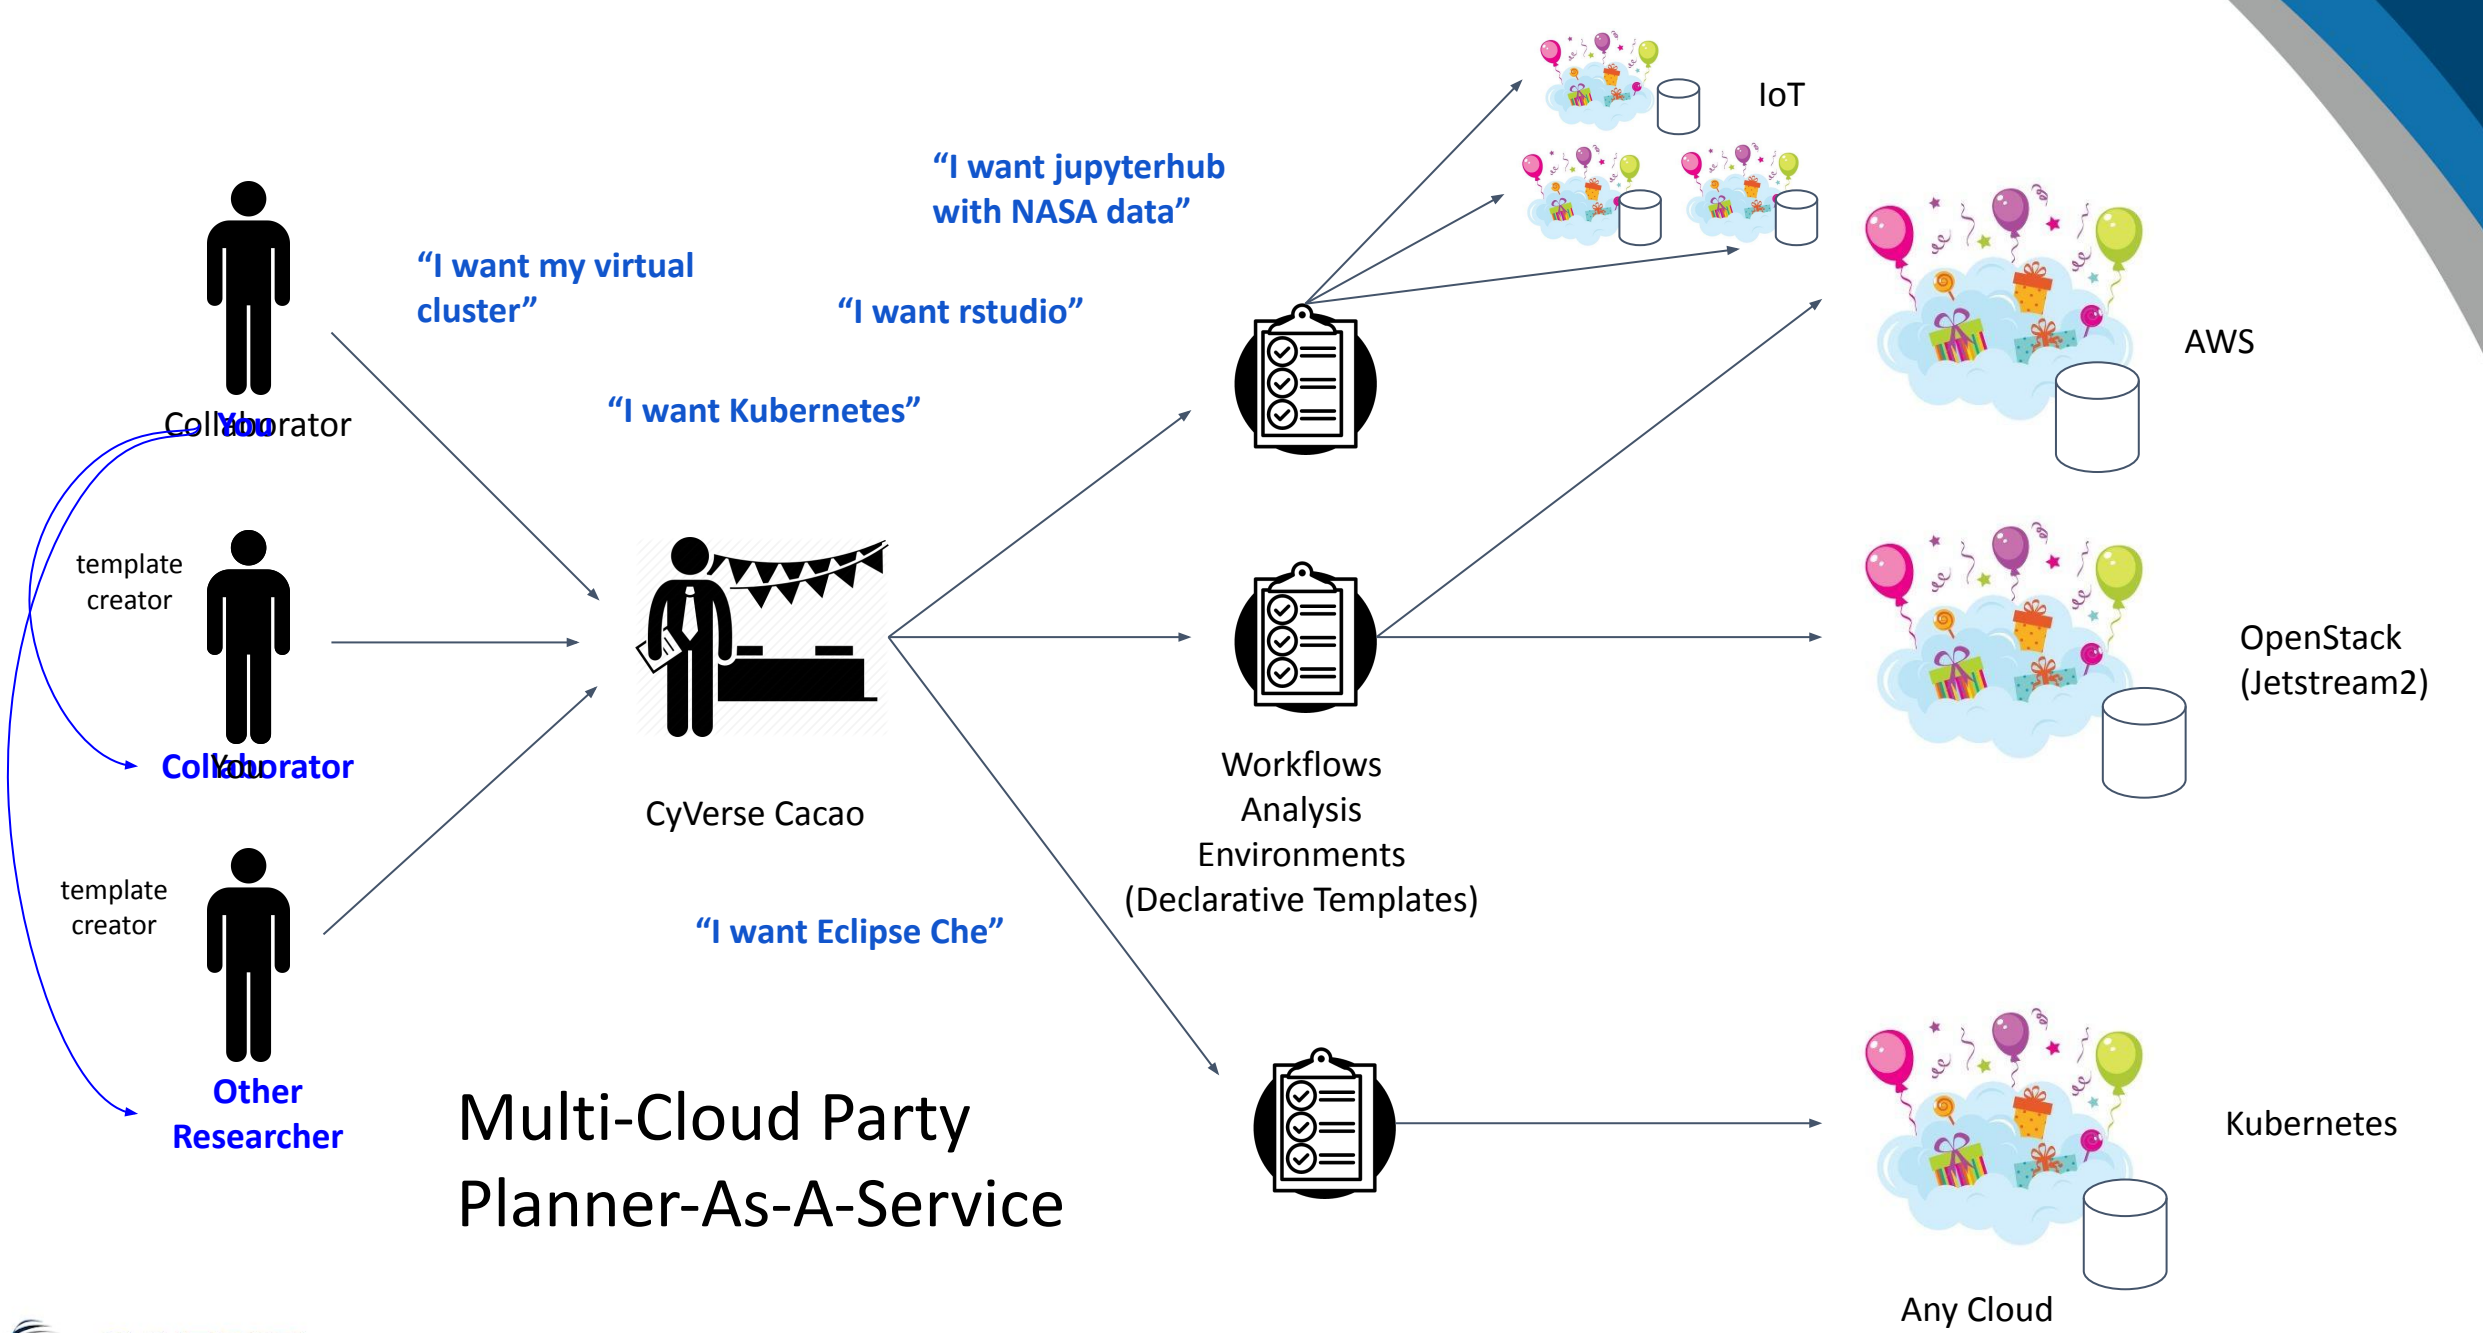
"I want this type of food"

"I want this venue"





## Multi-Cloud Party Planner-as-a-Service



# What is Cacao, really?

An **event-driven multi-cloud** service that enables researchers and educators to effortlessly manage, scale and share their tools and workflows to any research-capable cloud **using declarative templates**.



kubernetes



mongoDB®



HashiCorp  
Terraform



elastic



HashiCorp  
Vault



PostgreSQL



ANSIBLE



argo



fluentd



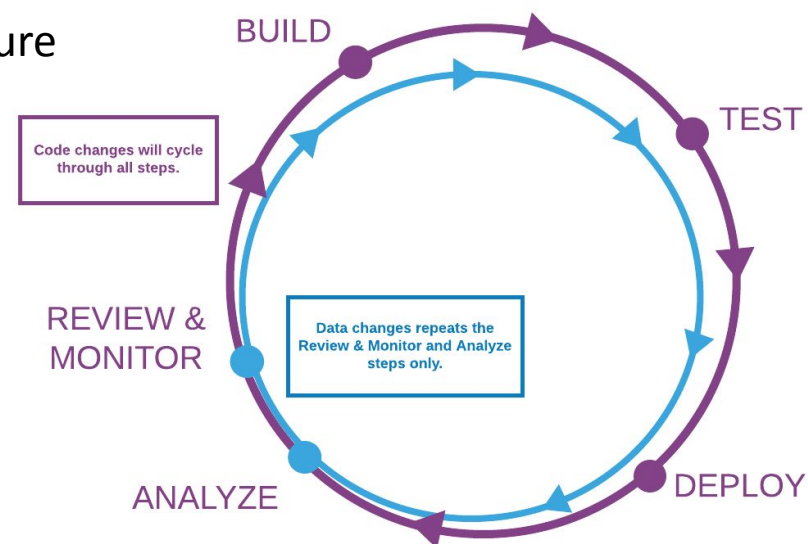
kibana



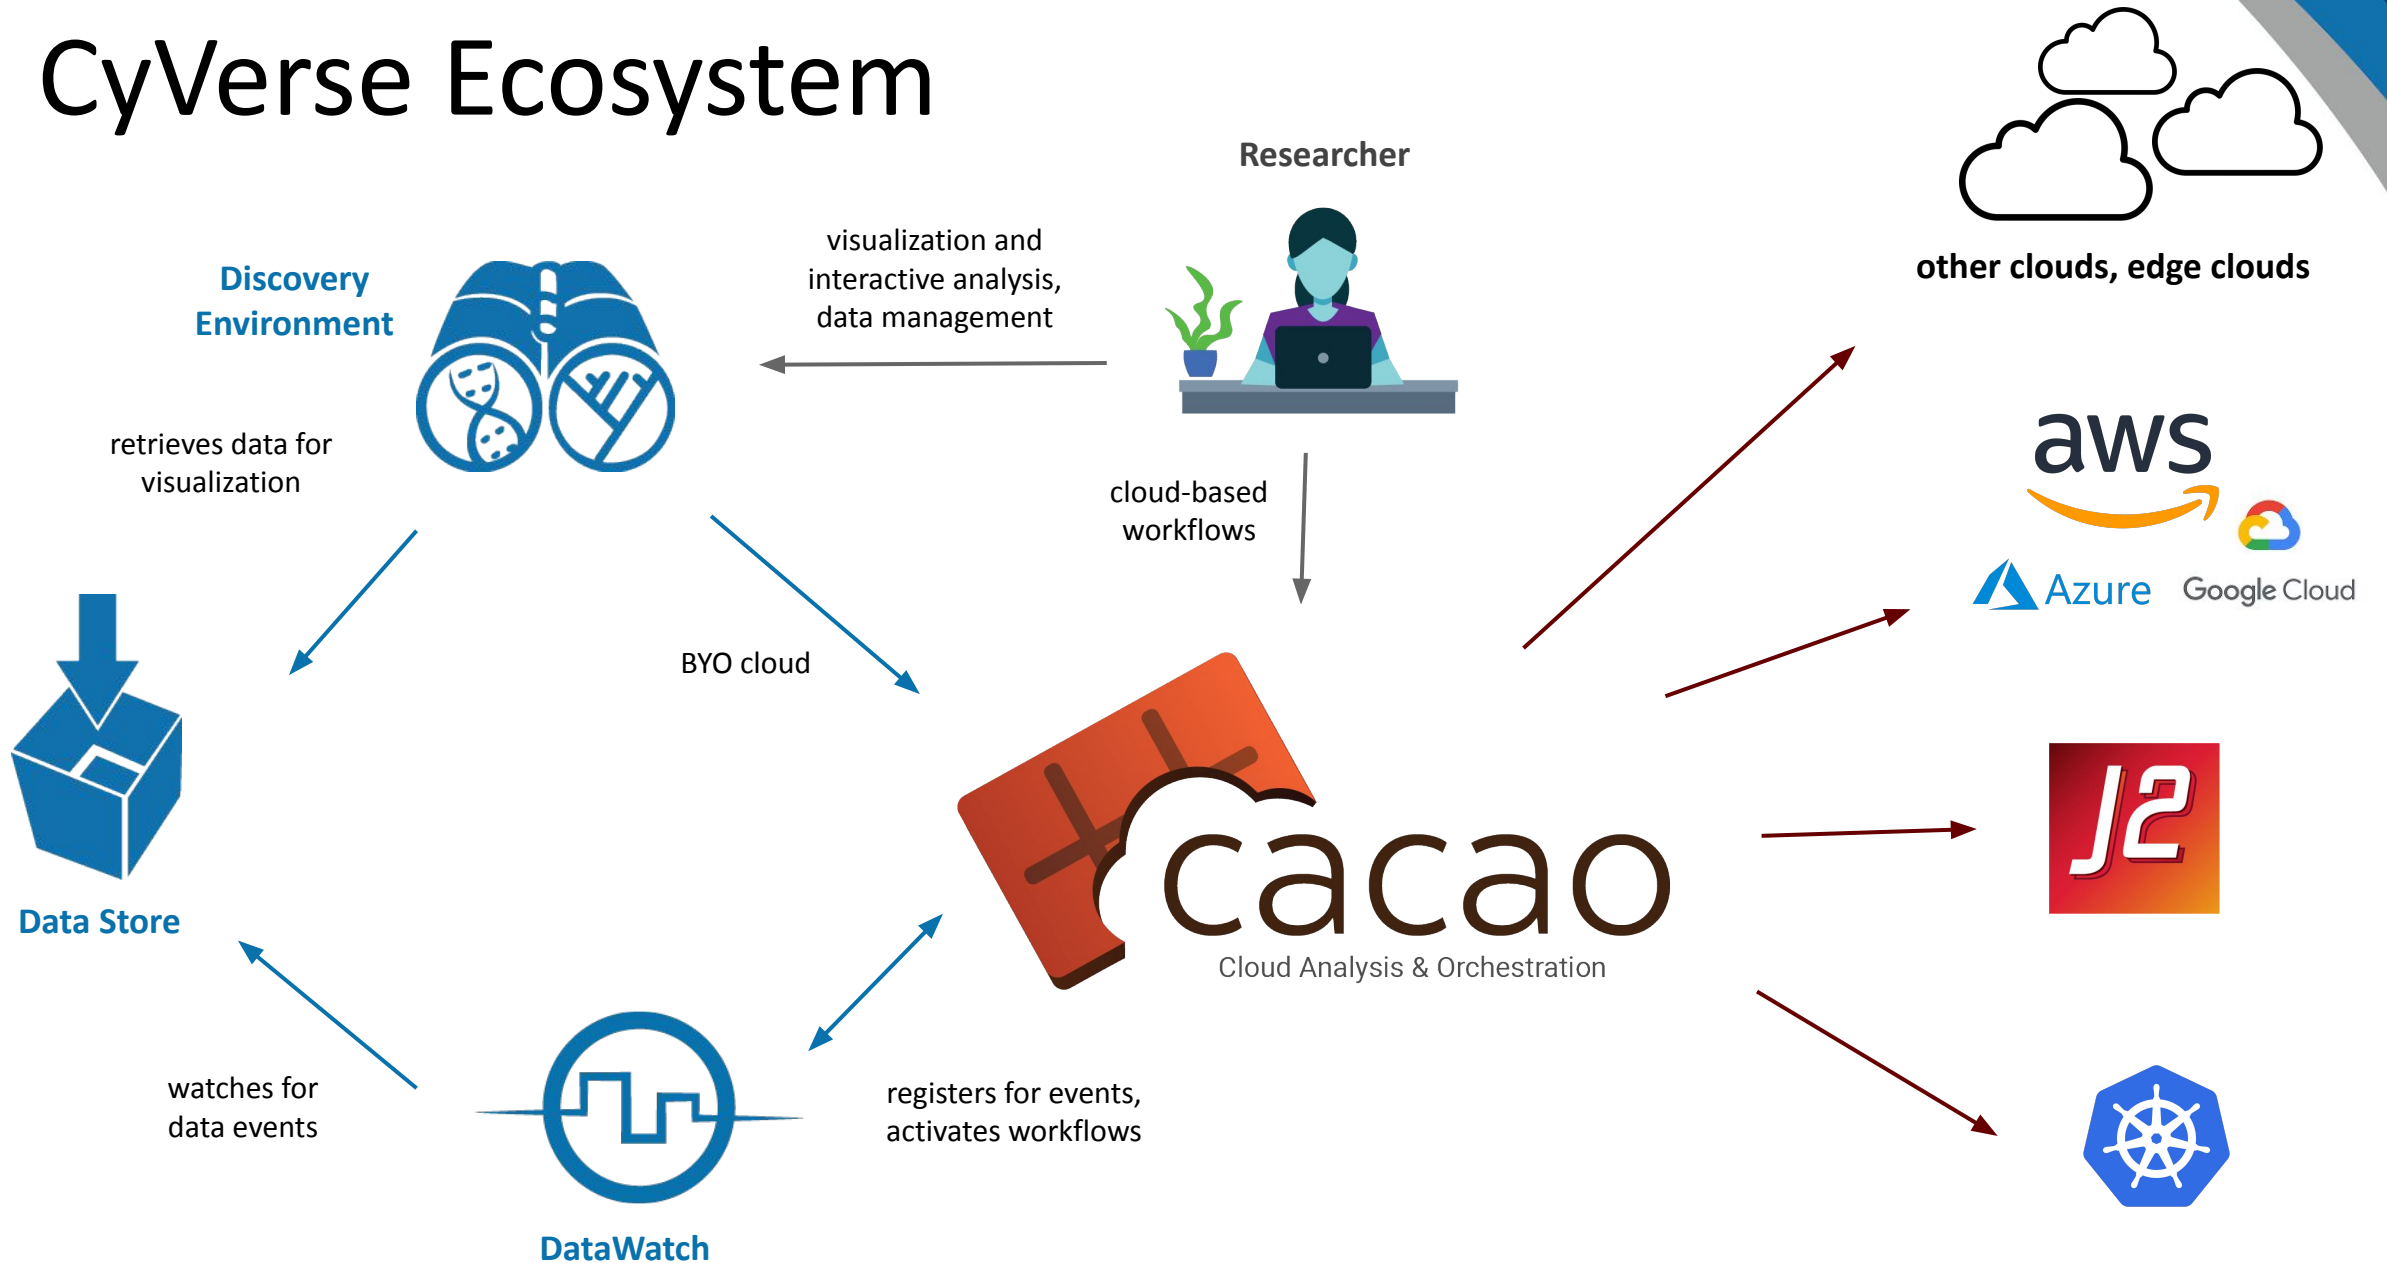
redis

# Characteristics and Capabilities

- Material UI, focused on simplicity, built on Next.js
- Event-driven, cloud native architecture built on Go, Kubernetes, Nats
  - highly extensible: templates & DSLs, future-resistant architecture
  - highly scalable
- Meta/Hybrid-cloud (via integration with declarative templates)
  - Terraform
  - k8s, Argo, openfaas
- Scale to zero and auto-scaling based on metrics\*\*
- Continuous analysis (CI/CD for science)
  - data events, code change events, periodic, external sources
- Git for template and workflow persistence
- OAuth2 and federated identity: Globus and Keycloak
- REST API & CLI



# CyVerse Ecosystem



# Templates in Cacao


- git is first class persistence
  - supports branches/tags
  - subpaths
  - repo changes reloaded



**git**



**GitLab**



**GitHub**

# Templates

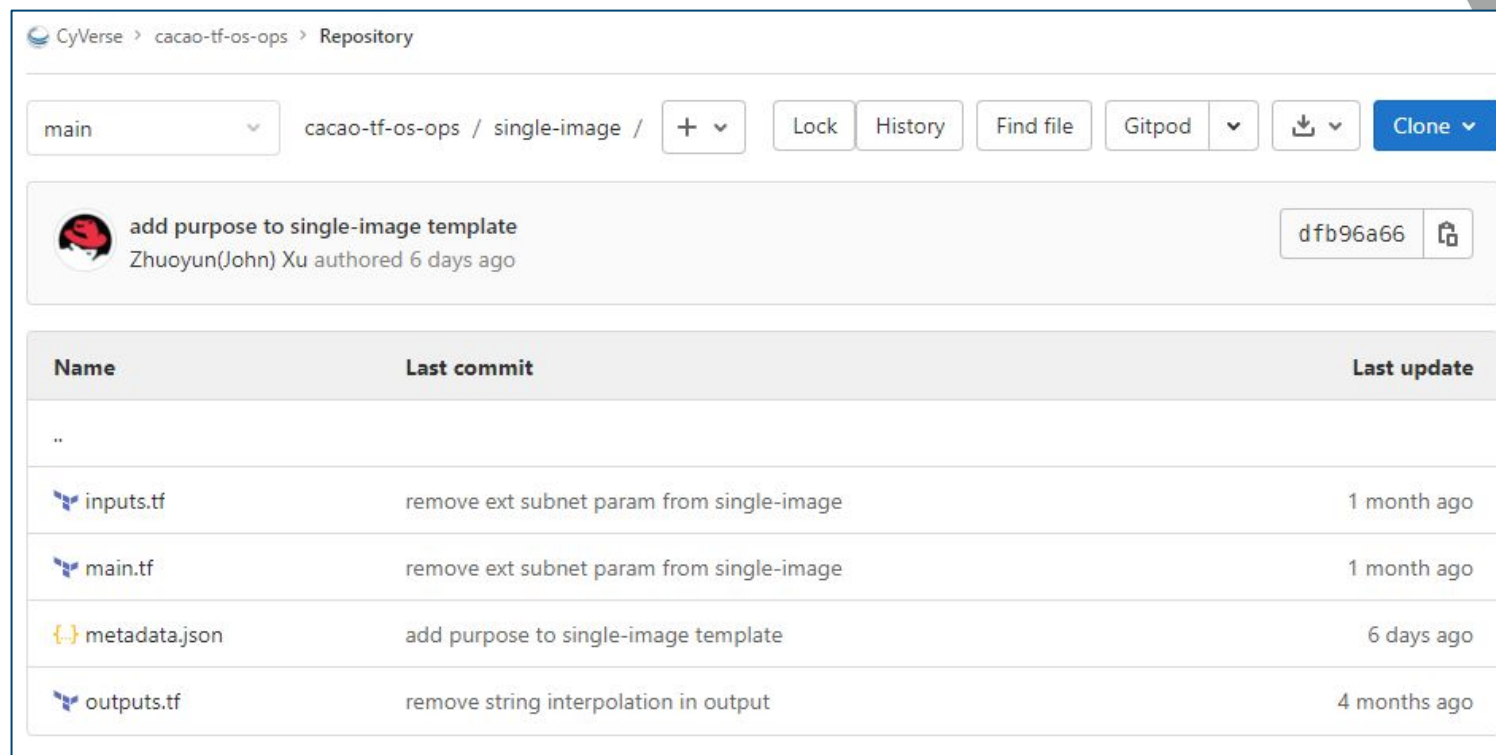
- git is first class persistence
  - supports branches/tags
  - subpaths
  - repo changes reloaded
- metadata-driven

```
{} metadata.json ×
cacao-tf-os-ops > single-image > {} metadata.json > ...
1  {
2  "name": "single image openstack instances",
3  "author": "Edwin Skidmore",
4  "author_email": "edwin@cyverse.org",
5  "description": "launches a single image on a single cloud",
6  "template_type": "openstack_terraform",
7  "purpose": "compute",
8  "cacao_post_tasks": [
9    {
10     "type": "ansible",
11     "location_type": "cacao_atmosphere_legacy"
12   }
13 ],
14 "parameters": {
15   "username": {
16     "type": "cacao_username",
17     "default": "username",
18     "ui_label": "CACAO username",
19     "description": "CACAO username"
20   },
21   "instance_name": {
22     "type": "string",
23     "ui_label": "Instance name",
24     "description": "Instance name"
25   },
26   "instance_count": {
```



# Templates

- git is first class persistence
  - supports branches/tags
  - subpaths
  - repo changes reloaded
- metadata-driven
- creation & customization
  - CyVerse curated
  - community-provided
  - easy to create from scratch or fork and modify




The screenshot shows a web interface for a CyVerse repository. The breadcrumb path is 'CyVerse > cacao-tf-os-ops > Repository'. The current branch is 'main'. The repository path is 'cacao-tf-os-ops / single-image /'. There are buttons for '+', 'Lock', 'History', 'Find file', 'Gitpod', a download icon, and 'Clone'. A commit message 'add purpose to single-image template' by 'Zhuoyun(John) Xu' is shown, authored 6 days ago, with a commit hash 'dfb96a66'. Below this is a table listing files and their last commit details.

| Name          | Last commit                               | Last update  |
|---------------|---|--------------|
| ..            |   |              |
| inputs.tf     | remove ext subnet param from single-image | 1 month ago  |
| main.tf       | remove ext subnet param from single-image | 1 month ago  |
| metadata.json | add purpose to single-image template      | 6 days ago   |
| outputs.tf    | remove string interpolation in output     | 4 months ago |

# Templates


- git is first class persistence
  - supports branches/tags
  - subpaths
  - repo changes reloaded
- metadata-driven
- creation & customization
  - initially, provided by us
  - community-provided
  - from scratch or fork
- cacao validation, testing, cloud filtering



# Templates

- git is first class persistence
  - supports branches/tags
  - subpaths
  - repo changes reloaded
- metadata-driven
- creation & customization
  - initially, provided by us
  - community-provided
  - from scratch or fork
- cacao validation, testing, provider filtering

- sharing
  - share via github, gitlab, etc (fork)
  - sharing and search within cacao



# Demo

- Take a walkthrough of an early version of Cacao
  - Jetstream 2
  - Basic template that will launch a vm using a terraform+ansible template:

<https://gitlab.com/cyverse/cacao-tf-os-ops/-/tree/main/single-image>

**Warning: Demo showing a boring automation of a template will start soon.\***

\* Unabashedly derived from Al Sweigart, author of “Automate the Boring Stuff” books

# Zero-to-Jupyterhub with OAuth

test-workspace-tacc-20211206

## × New Deployment

✓ Select Workspace — ✓ Select Template — ✓ Select Credential — **4** Set Parameters — 5 Review & Deploy

### Deployment Details

Deployment Name\*

No. of Instances

JS-API-Featured-Ubuntu20-Latest

Size

### JupyterHub Auth Type

GitHub

Default

Dummy Auth

CyVerse (KeyCloak)

### Allowed Users

Add Usernames

edwin × nirav ×

mariah ×

### JupyterHub Configuration

OAuth2 Client ID

OAuth2 Client Secret

Image

Tag

[LEARN ABOUT AUTH TYPES AND CONFIGURATION](#)

PREVIOUS

Advanced Options

Screen Mockup  
Not yet in production

# Kubernetes-on-demand

test-workspace-tacc-20211206

## × New Deployment

✓ Select Workspace — ✓ Select Template — ✓ Select Credential — **4** Set Parameters — 5 Review & Deploy

### Deployment Details

Deployment Name\*

No. of Instances

JS-API-Featured-Ubuntu20-Latest

Size


Open Ports

Resource Definition  Disable  Enable

URL  Raw Text

Yaml

PREVIOUS

 Screen Mockup  
Not yet in production

# Cloud Management

The screenshot displays the Jetstream2 Cloud Management dashboard. The interface is organized into several sections:

- Header:** Features the Jetstream2 logo and a navigation menu on the left with options: Home, Deployments, Clouds (highlighted), Templates, Credentials, and Help.
- Clouds Section:** Includes an Openstack icon and three cloud configurations:
  - Jetstream IU:** jetstream\_indiana\_university, TG-TRA160003, Jetstream Cloud for Indiana University.
  - Jetstream-TACC:** jetstream\_TACC, TG-TRA160003, Jetstream Cloud for TACC.
  - Jetstream2:** jetstream2, TG-TRA160003, Jetstream2 Cloud.
- Kubernetes Section:** Includes one configuration:
  - CyVerse k8S Continuous Cloud:** cyverse\_kubernetes\_continuous\_cloud, TG-TRA160003, CyVerse Kubernetes Cloud for Continuous Analysis.
- Public Clouds Section:** Includes two configurations:
  - Google Kubernetes Engine (GKE):** google\_kubernetes\_engine, TG-TRA160003, Google Cloud for Continuous Analysis.
  - AWS Elastic Kubernetes Service (EKS):** AWS\_elastic\_kubernetes\_service, TG-TRA160003, CyVerse Kubernetes Cloud for Continuous Analysis.

An orange notification box in the bottom right corner reads: "Screen Mo Not yet in pro".

# Searchable template catalog

The screenshot displays the Jetstream2 web interface. On the left is a dark sidebar with navigation links: Home, Deployments, Clouds, Templates (highlighted in red), Credentials, and Help. The main content area features the Jetstream2 logo at the top, followed by a 'Templates' section header. Below this is a search bar and the text 'Showing 100 of 184 templates'. The 'Featured Templates' section lists three items:

- Simple launch of one or more VMs**: Created 12/12/2022 00:00 by username. Purpose: compute. Description: launches a k3s kubernetes cluster. Versions: 5. Tags: K3, Compute, Tag, Tag.
- Launch a Jupyter Hub**: Created 12/12/2022 00:00 by username. Purpose: compute. Description: launches a k3s kubernetes cluster. Versions: 5. Tags: K3, Compute, Tag, Tag.
- Launch a Kubernetes-on-demand Cluster**: Created 12/12/2022 00:00 by username. Purpose: compute. Description: launches a k3s kubernetes cluster. Versions: 5. Tags: K3, Compute, Tag, Tag.

The 'All Templates' section lists three items:

- kubernetes (k3s)**: Created 12/12/2022 00:00 by username. Purpose: compute. Description: launches a k3s kubernetes cluster. Versions: 5. Tags: K3, Compute, Tag, Tag.
- Another Template**: Created 12/12/2022 00:00 by username. Purpose: compute. Description: launches a k3s kubernetes cluster. Versions: 5. Tags: K3, Compute, Tag, Tag.
- Another Template**: Created 12/12/2022 00:00 by username. Purpose: compute. Description: launches a k3s kubernetes cluster. Versions: 5. Tags: K3, Compute, Tag, Tag.

At the bottom right, there is an orange notification box with an information icon and the text 'Screen Mo Not yet in pro'.



# Cacao Projects

**Jetstream 2:** NSF (OpenStack) cloud for science and engineering

**HydroGEN:** ML-based simulation platform to simulate hydrological scenarios with the goal of elastic scaling to AWS

**COALESCE** (COntext Aware LEarning for Sustainable CybEr-agriculture systems): innovation by combining digital agricultural tools, like robotics, and physical tools, like sensors and drones; will have distributed computing needs, including ML compute to the edge (farm sites)

**AIIRA** (AI Institute for Resilient Agriculture): creating and deploying AI-driven predictive digital twins for modeling plants to increase resiliency of the nation's agricultural systems; MLOps in the cloud

# Cacao Benefits for Users

- Bring your own cloud
  - bring your own credentials
  - connect on-premise clouds or IoT (your own Kubernetes cluster)
  - share cloud access
- Federated identity (e.g., Keycloak, Globus)
- Modes of provisioning (on-demand, persistent, event-driven)
- Cloud cost-containment
  - Event-driven
  - Budgeting (notifications or enforcement)
- Programmable, either using API or CLI
- Template reuse

# Cacao Benefits for Infrastructure and Template Creators

- Connect to your cloud
- Create templates for your own community or group
  - Store your templates in your own repo
  - Easily fork existing templates and customize
  - Share templates with users
  - Publish templates and have them be searchable

# Future & Roadmap

- “Alpha” -- now heavily in development
- Improve the performance of deployments
  - exposing more of the underlying execution of a template as it happens
- Complete integration with other CyVerse services
- Complete and expose event triggers e.g. for data events
- Integrate other template engines and frameworks, esp ML-related ones
- Searchable template catalog, along with template management
- Add OpenStack-specific operations, like instance snapshotting
- Template-Cloud Resource reconciliation (within reason)

# Your input is needed

- Share with us!
  - your use cases
  - your template languages that fill a great need (future roadmap)
  - your templates or suggested templates (CyVerse-curated)
- Current focus: terraform+openstack (ansible optional)
- Near future
  - terraform+aws
  - kubernetes native
  - argo
- Beta testers
- Developers needed as well!

# Next Steps

## **When will Cacao be available?**

Cacao (Alpha) will be available very soon

## **Where?**

Available through Jetstream2

## **Who can access?**

Those with XSEDE allocation to get access to Jetstream2

# Meet the Team

Matt Bomhoff

Jeremy Frady

John Wregglesworth

Illyoung Choi

Sarah Roberts

John Xu

Sean Davey

Mariah Wall

Edwin Skidmore

## Students

Ryan Schneider

Terry Melo Cruz

# Thank you!