Billing and Secure Collaboration

Managing cloud costs and collaborating securely

https://terra.bio

BROAD INSTITUTE

Session Goals

- What cloud computing vendors work with Terra?
- What costs money in Terra and how do you pay for it?
- What is a Terra billing project and how is it connected with a workspace?
- Where can you see how much you've spent?
- Any recommendations for getting started with running analyses and securely sharing workspaces?

Cloud Vendors

Pay for what you use in “the cloud”

- Terra is free to sign up for and browse
- Computing services cost money
- Cloud providers charge for usage

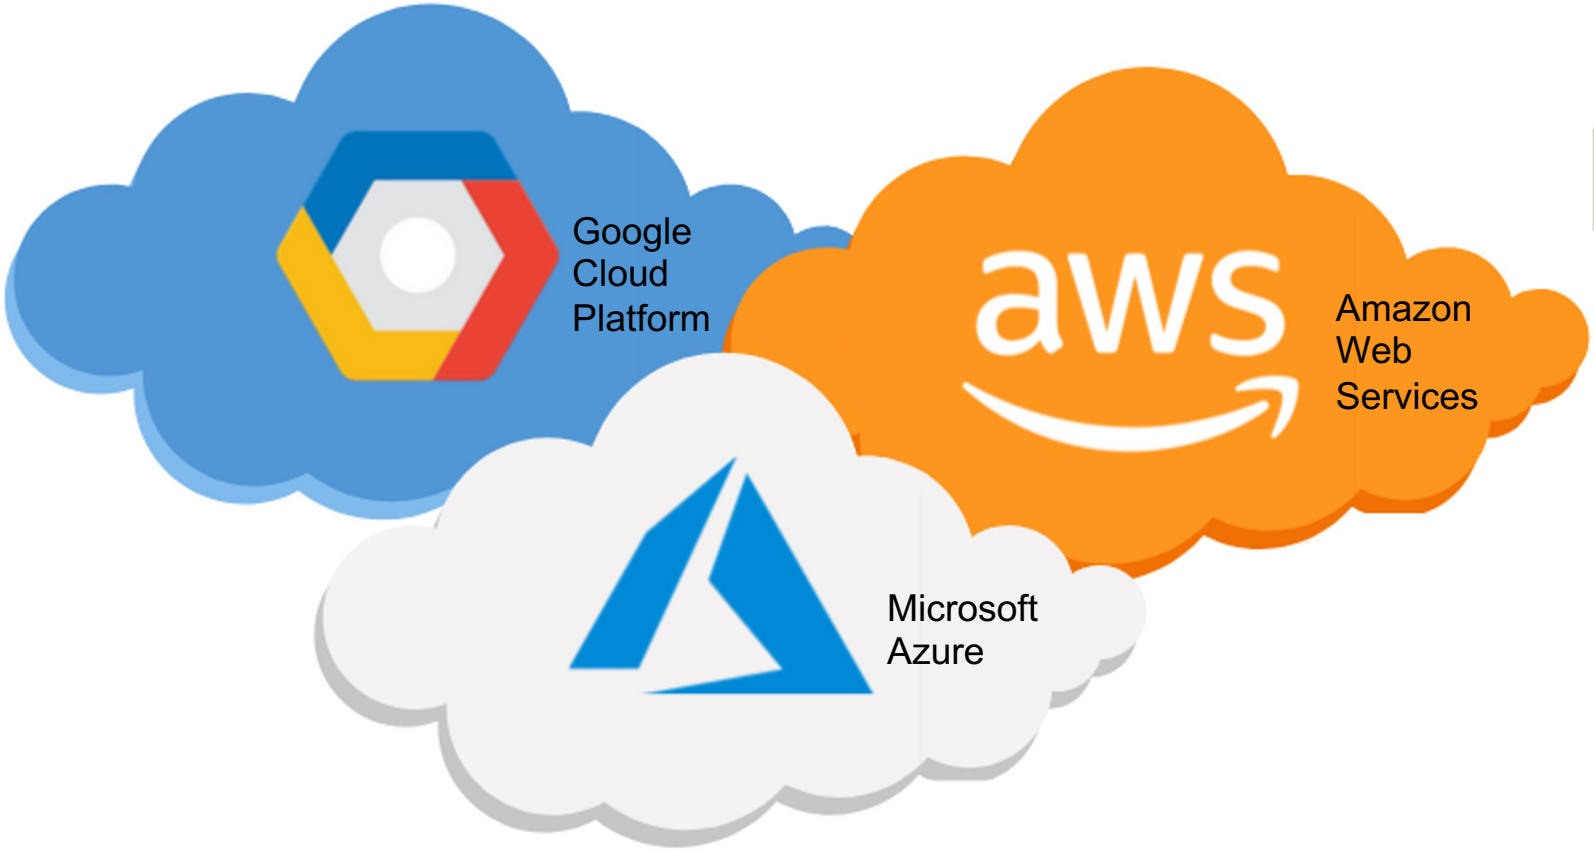

https://terra.bio

BROAD INSTITUTE

What are the costs of working on Terra?

1. **Compute costs**
   - Instance type (high performance costs more; preemptibles cost less)
   - CPU hours
   - Memory & Disk
2. **Data storage**
   - Anything in the Workspace bucket
   - Workflow submission outputs (whether the workflow succeeds or fails)
3. **Data egress** (i.e. moving data) costs
   - Downloading data to a non GCP location
   - Moving hosted data to a GCS bucket in a different region

Costs are organized by workspace

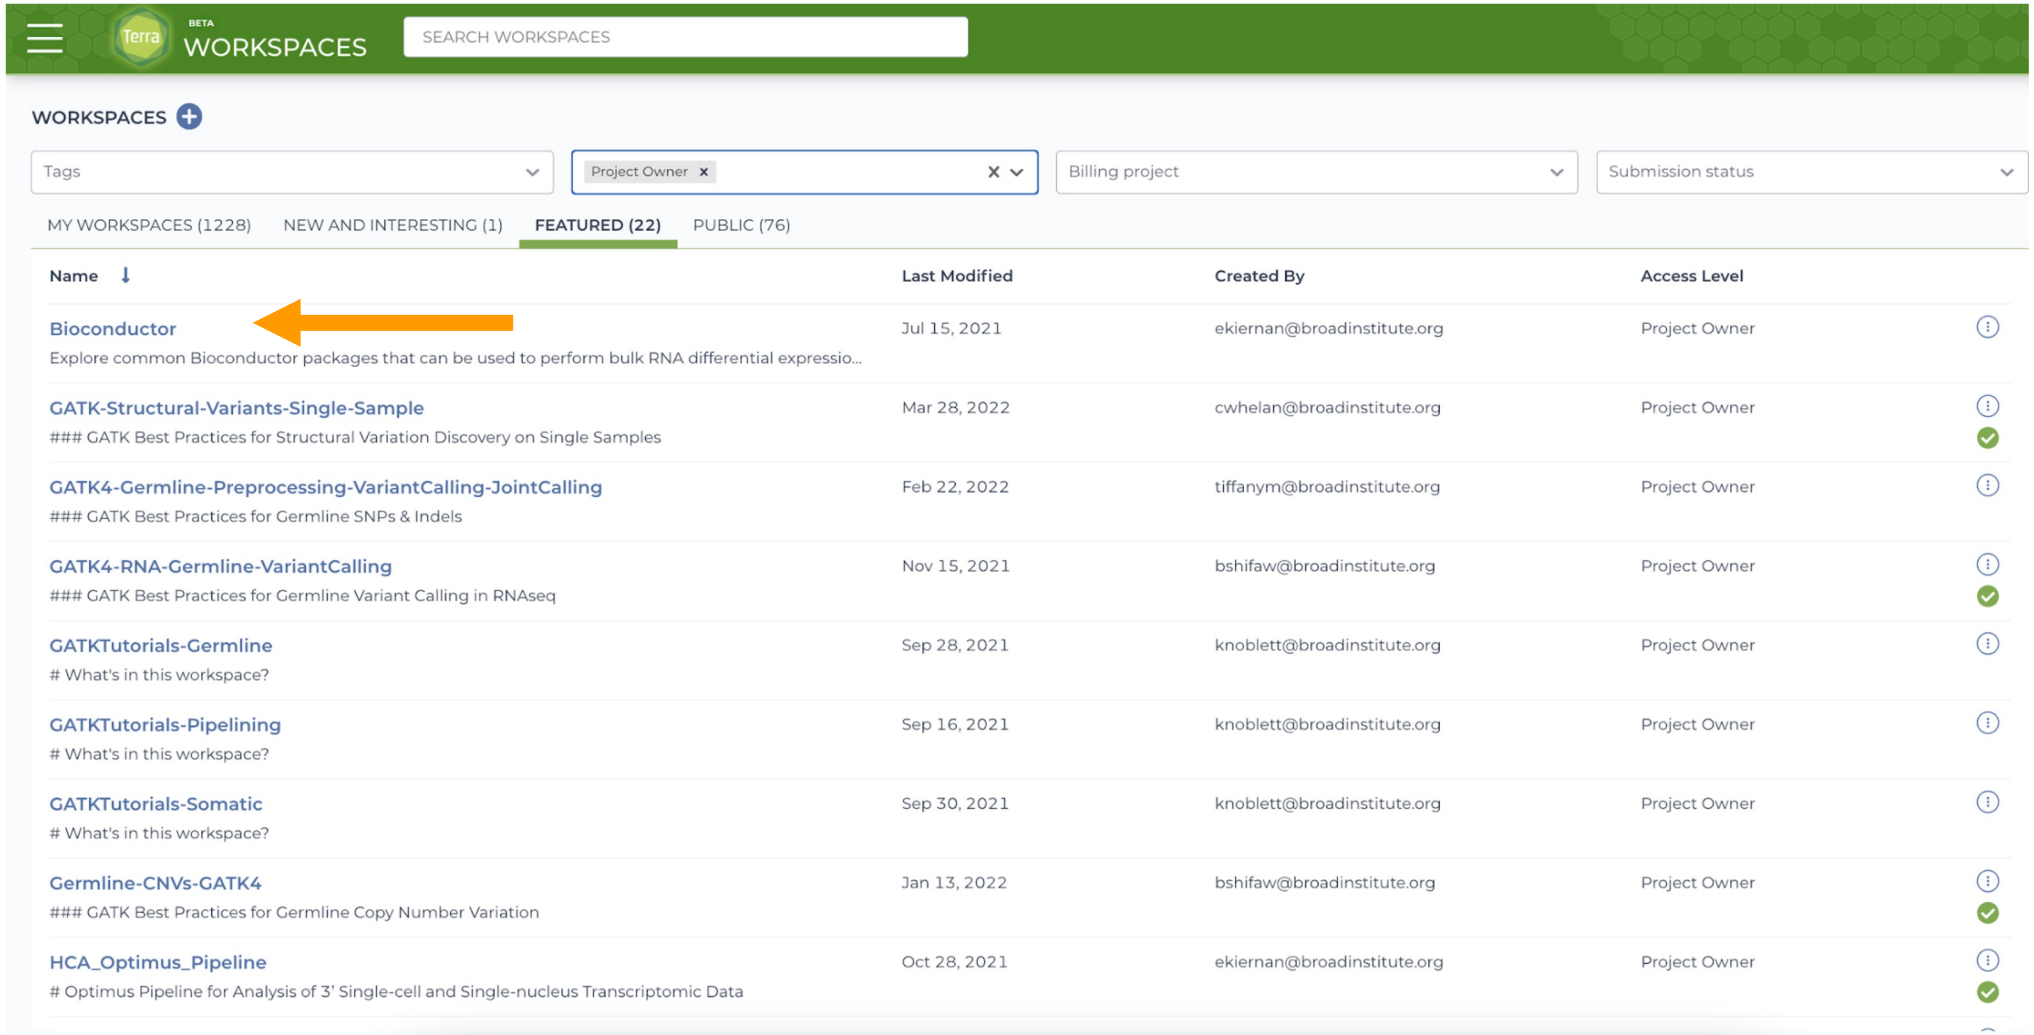

https://terra.bio

BROAD INSTITUTE

Terra Billing Project

- Terra billing project – a feature that links a cloud billing account with a workspace

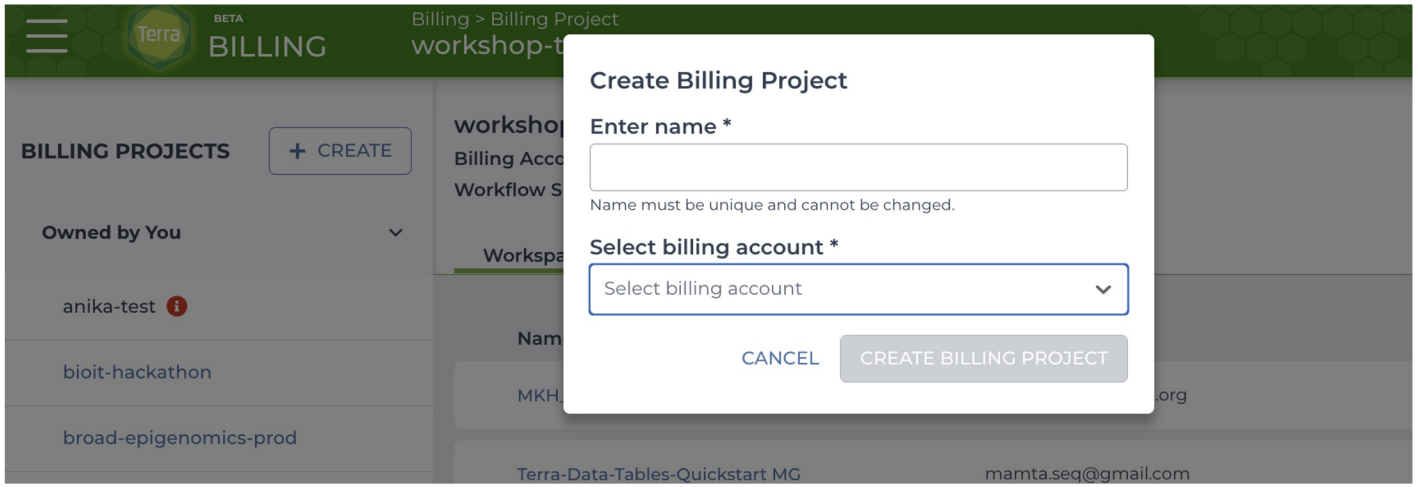

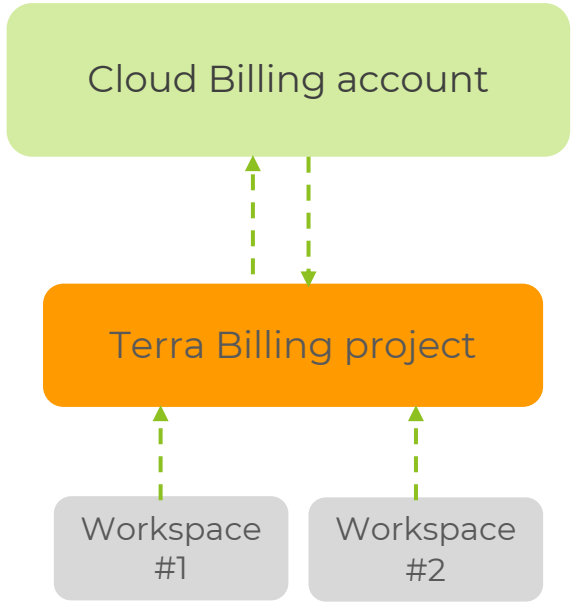

https://terra.bio

BROAD INSTITUTE

Compute Costs

Estimating workflow spend

- Find cost estimates in featured workspace **dashboard descriptions**

CRAM_To_BAM

Sample Name	Sample Total Size	Time	Cost $
NA12878_small	6.84 MB	00:07:00	0.02

- Use **built-in cost reporting** in the Job History page (see how to set it up here)

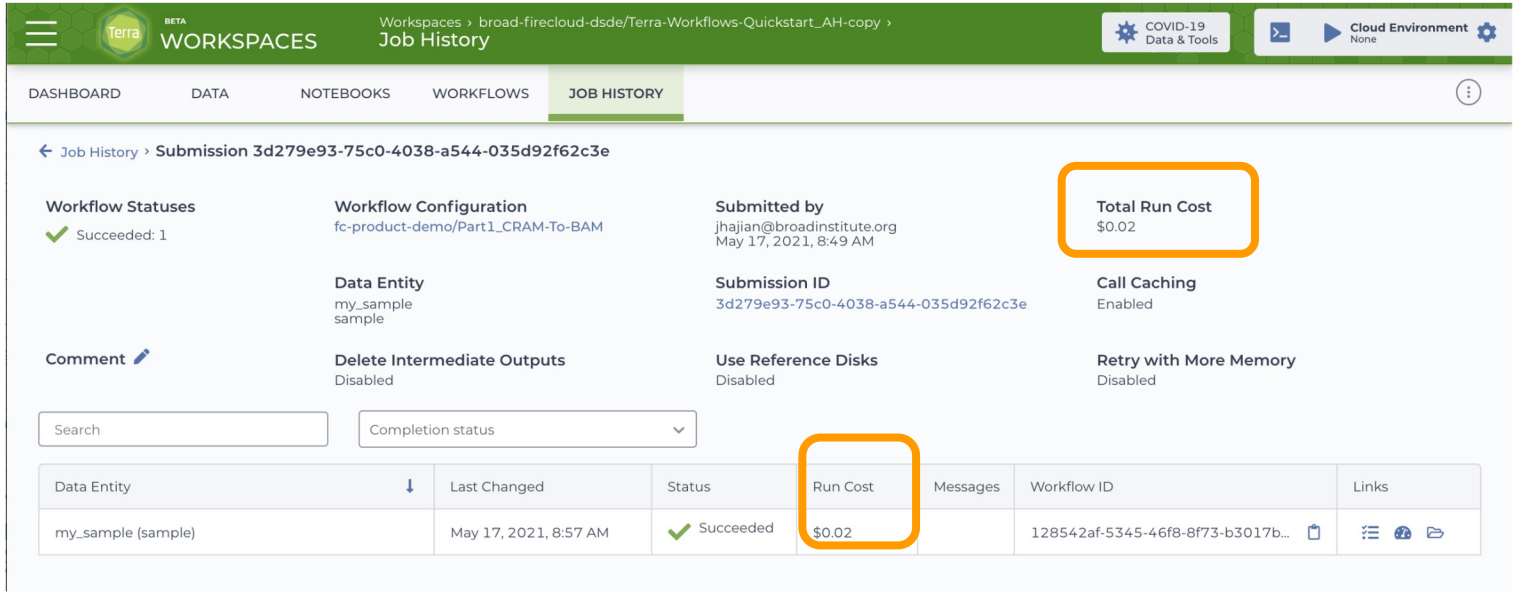

1) ***Run** on a **small sample***
2) ***Multiply** to scale*
3) ***Add buffer** for failed jobs*

Estimating interactive analysis costs

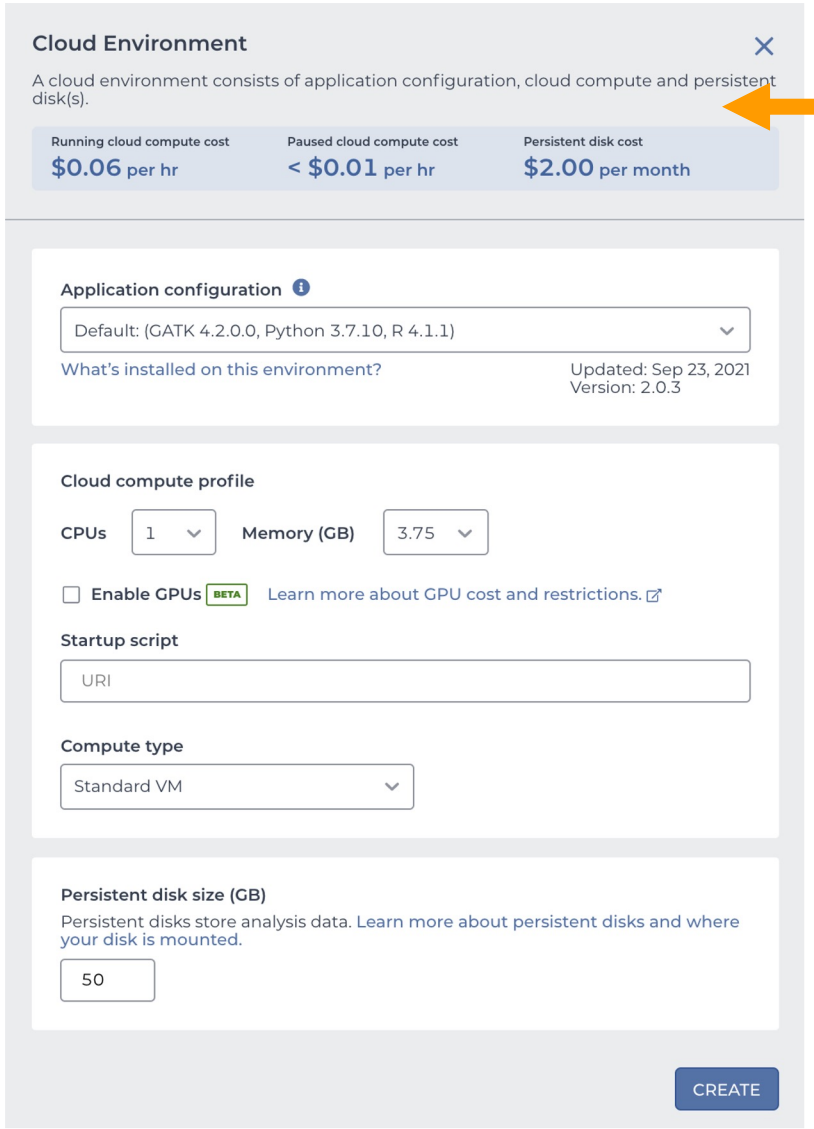

Adjust settings *to optimize cost*
(Compute Type, CPU, and disk)

Estimate costs *using real-time cost/hour in Cloud Environment widget*

- Updates based on the machine configuration you choose
- Total cost (estimate) =
(cost per hour) x (hours the instance will be active) + cost of the Persistent Disk
- Autopause function saves money! Delete when you are done!

https://terra.bio

See *Understanding and adjusting your Cloud Environment* in Terra Support

BROAD INSTITUTE

Storage Costs

Estimating cloud storage costs

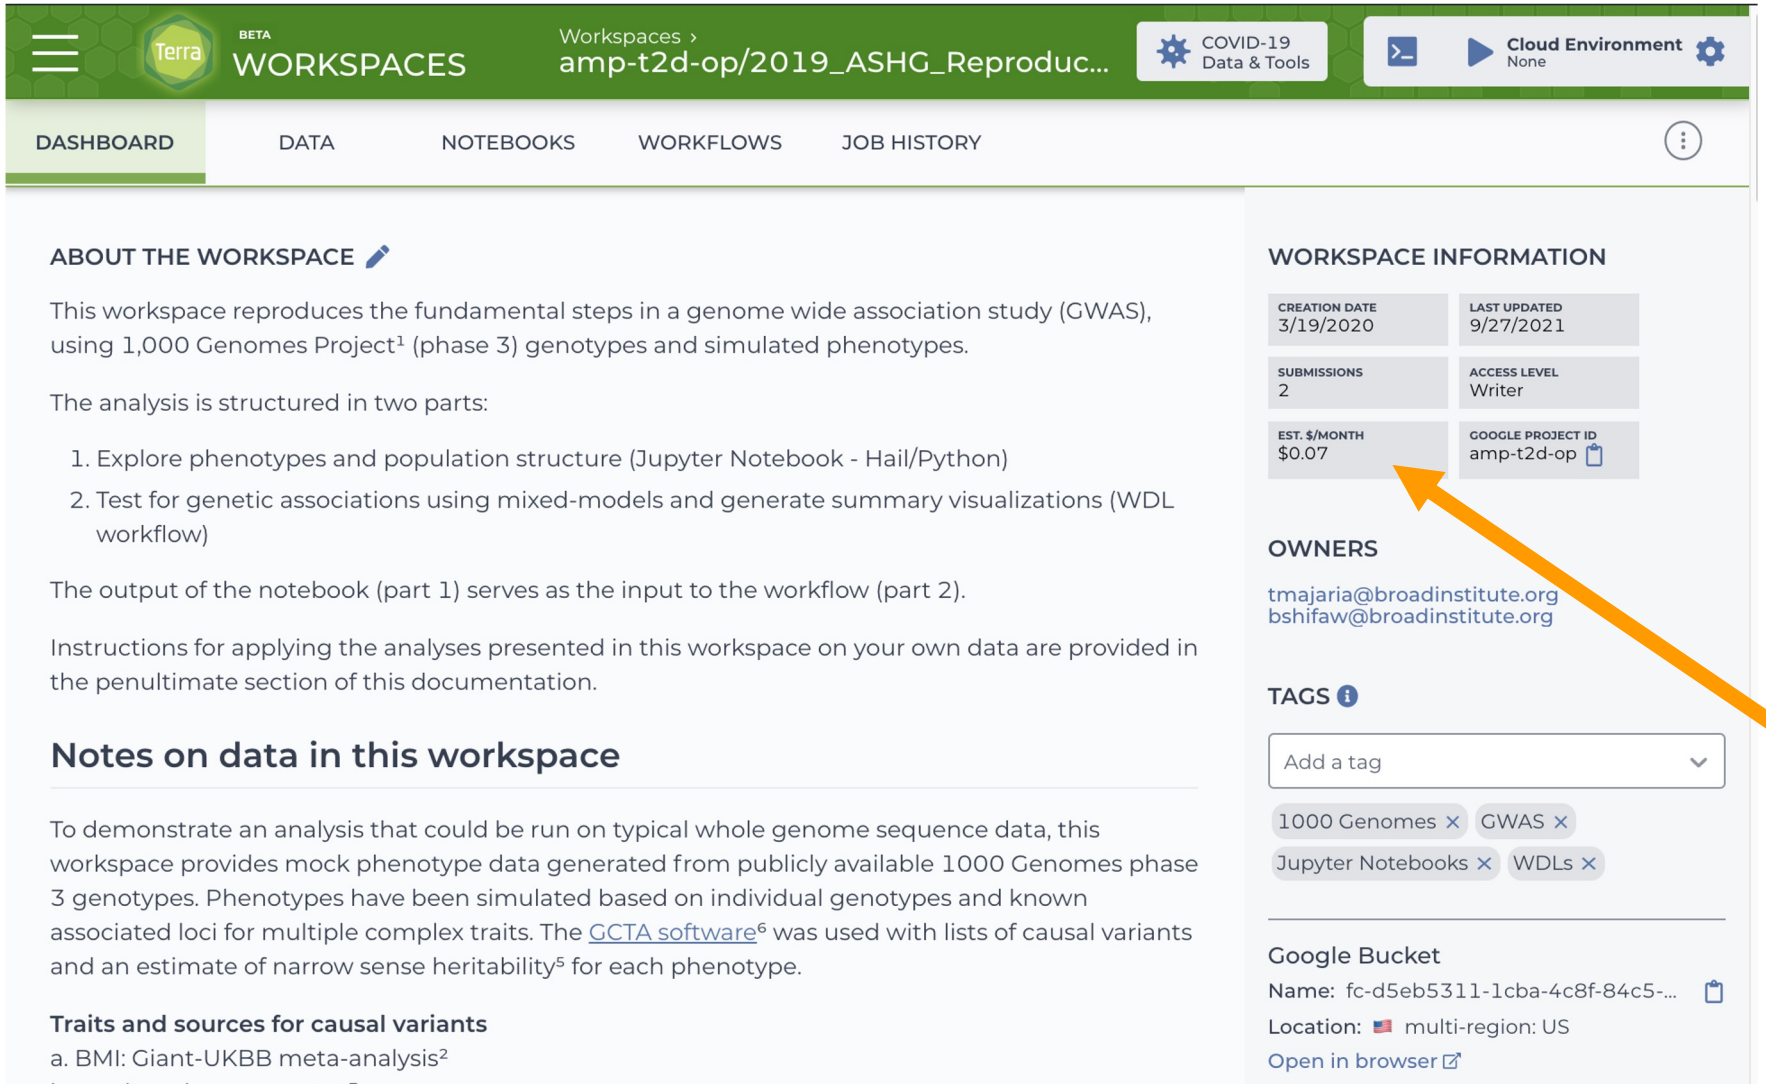

***Estimated** cloud storage costs for your workspace*

Egress Costs

Estimating egress costs

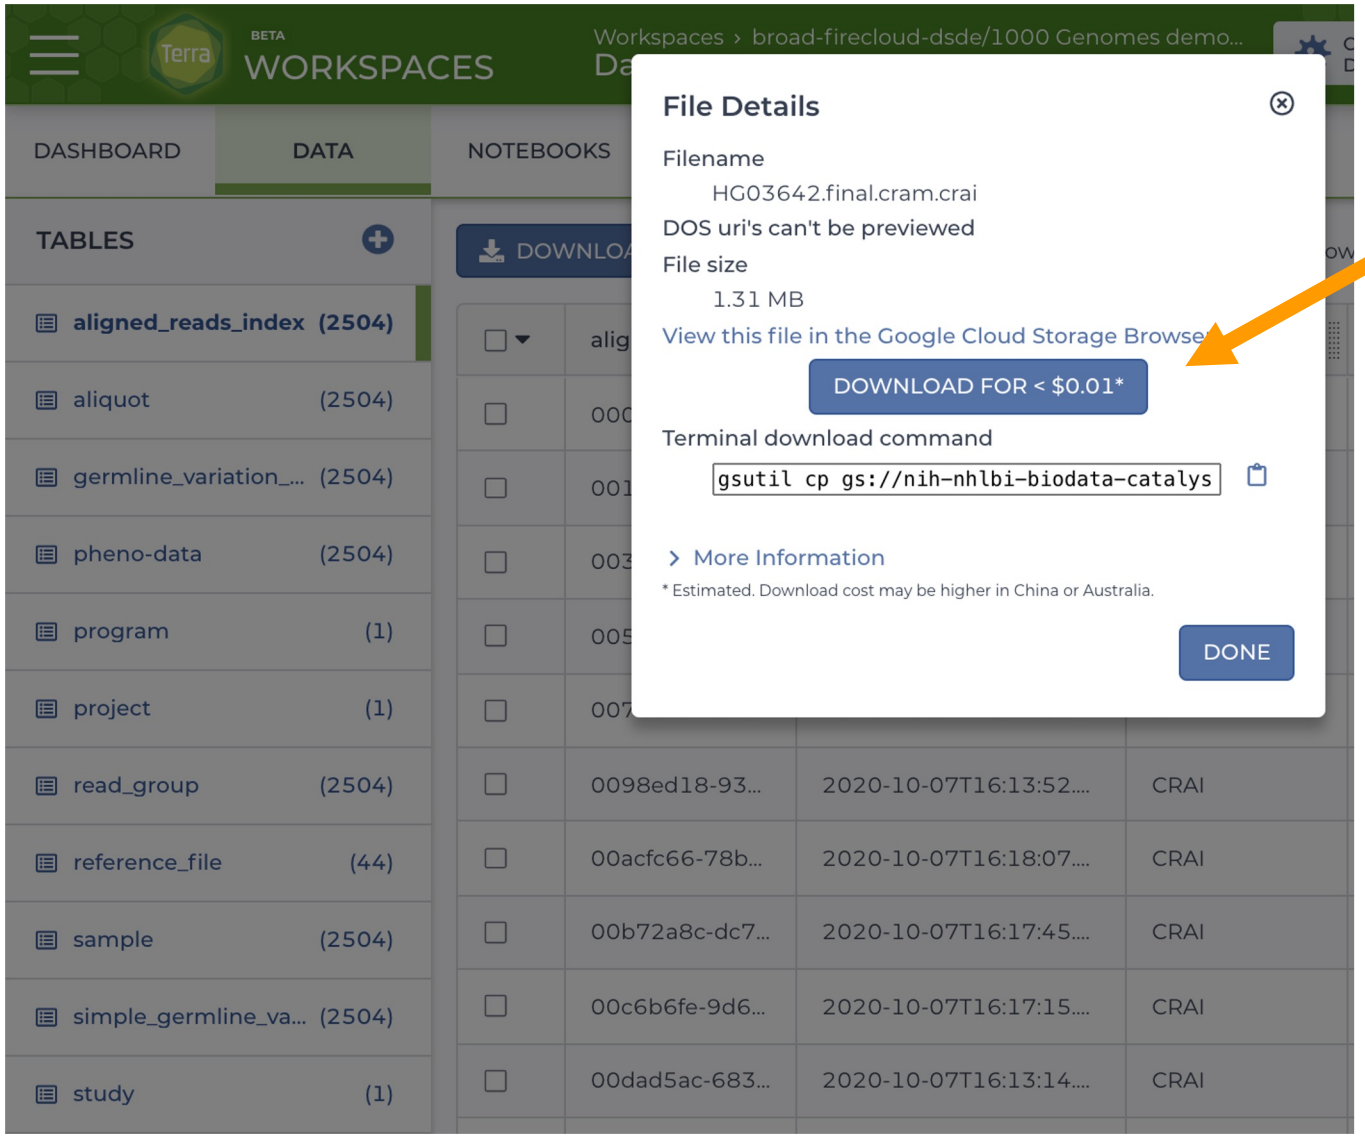

Find price to egress

Will you download derived data to save locally or elsewhere?

For more information, see *Understanding and controlling cloud costs* in Terra Support

Roles & permissions

Terra Billing project roles

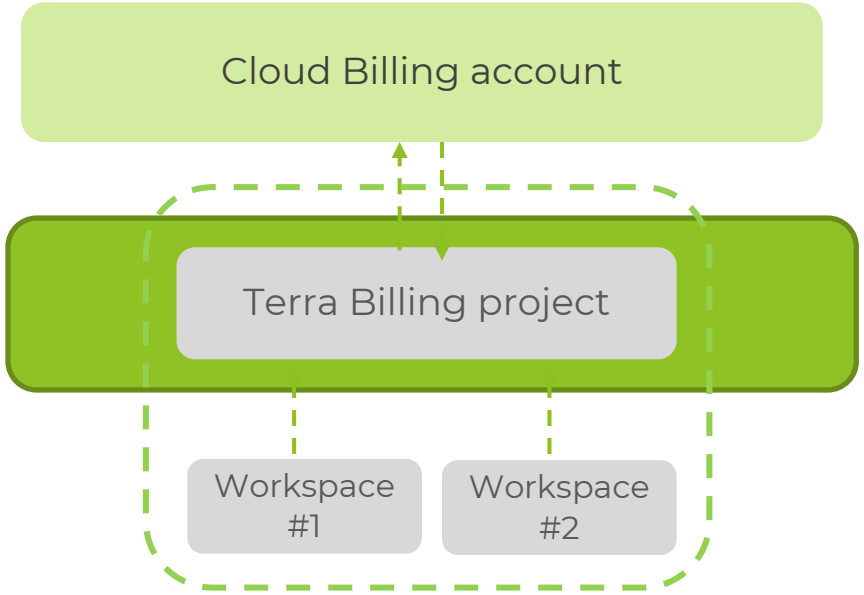

Billing roles determine who can create/manage resources

- Owners & Users can create Terra workspaces
- Owners can grant Terra Billing project permissions to other collaborators

Terra Billing project roles

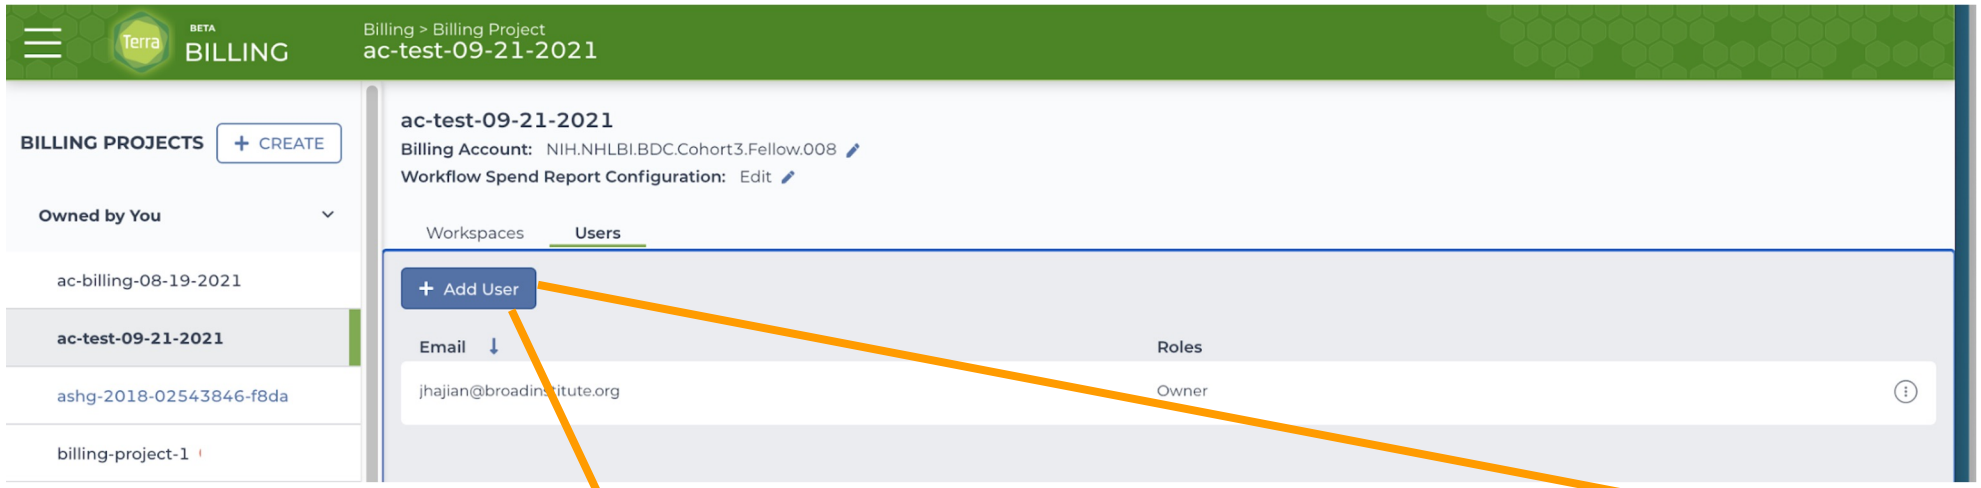

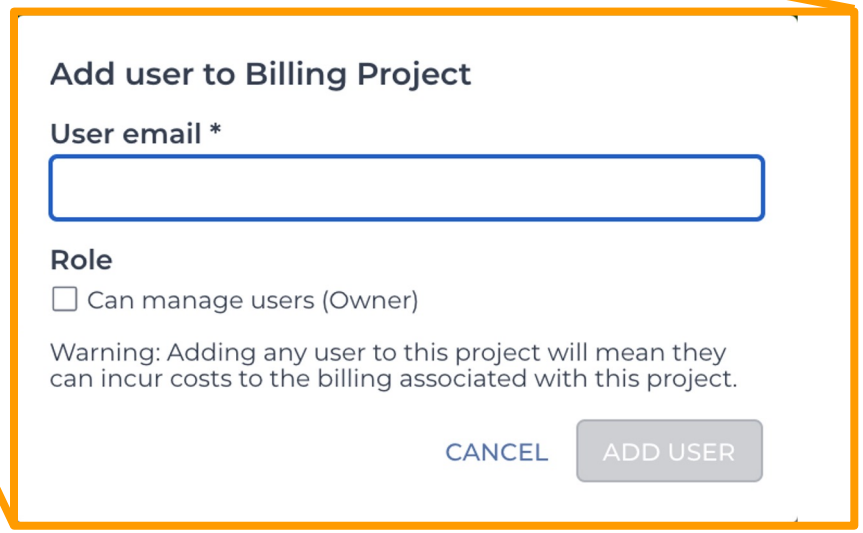

Roles include

- Owner (i.e. Project Owner)
- User

https://terra.bio

BROAD INSTITUTE

Workspace roles

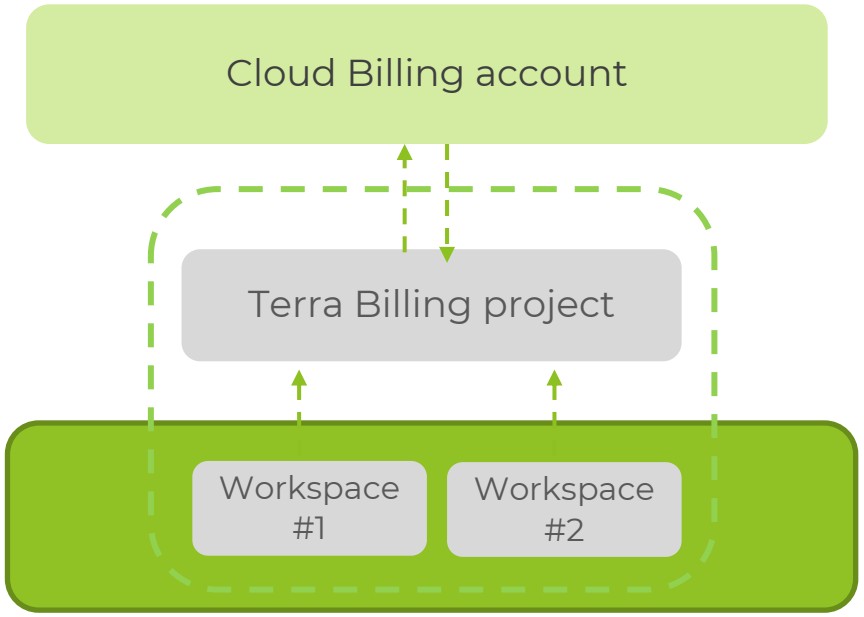

Workspace roles determine who can perform actions with a cost (spend money)

- **Roles** include
  - Owner
  - Writer
  - Reader
  - Can share
  - Can compute

Costs are charged to the workspace Billing project – not the user!

Streamline shared access with Groups

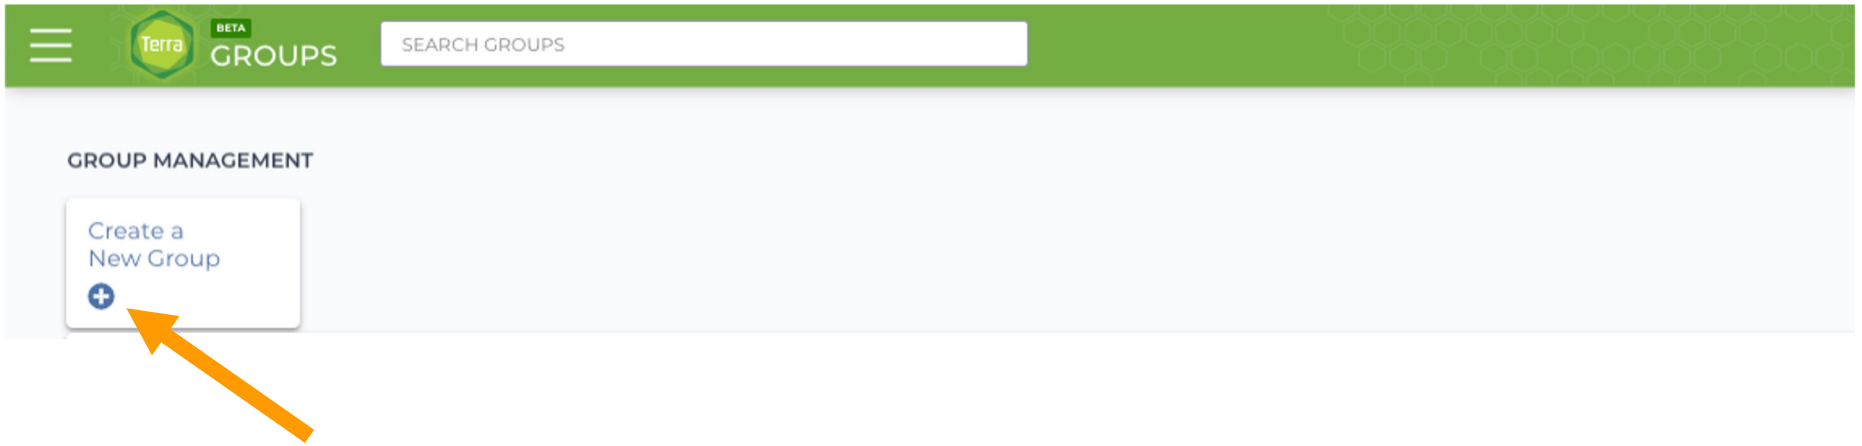

- Add groups at **Main Menu navigation** > **groups**
- Assign **roles** en masse for a given **resource**
- Can be updated as needed
  - Member - Has access to any resources granted group permission
  - Admin - May add or remove members or other admins to the group - also members of the group
- Permissions across all resources updates when group is updated

To learn more about using permissions to manage shared resources, see https://support.terra.bio/hc/en-us/articles/360024617851

Authorization Domains can contain one or more groups

- Authorization Domain-

A group or multiple groups that are allowed to access a workspace.

- Example ->

If you are in both ASHG18-dev and the Broad_Product_Team groups, you can access this workspace if you are given Owner, Writer, or Reader role.

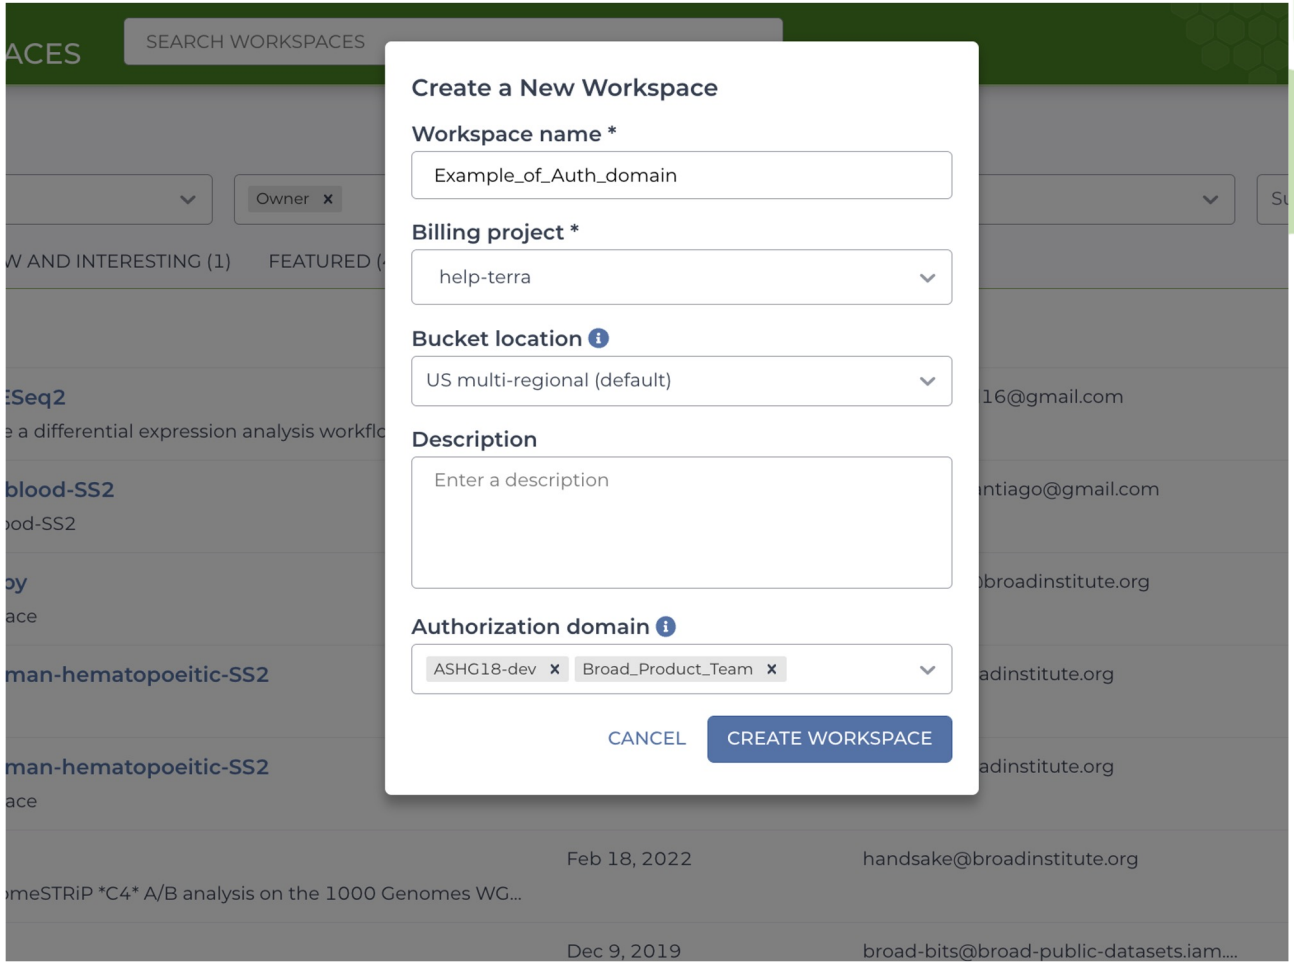

https://terra.bio

BROAD INSTITUTE

Setting up billing

Options for working in Terra

How to set up a Cloud Billing account

How to pay for costs when working in Terra

- A collaborator can grant you “can compute” permission in a shared workspace
  - **Your** costs are covered by **their** Billing project

To browse showcase workspaces (no cost!), see the Showcase workspace library

How to pay for costs when working in Terra

- A collaborator can grant you “can compute” permission in a shared workspace
  - **Your** costs are covered by **their** Billing project
- The owner of an existing Terra Billing project can authorize you to use their Billing project
  - You can create/clone a workspace to work in
  - **Your** workspace costs are covered by **their** Billing project

To browse showcase workspaces (no cost!), see the Showcase workspace library

How to pay for costs when working in Terra

- A collaborator can grant you “can compute” permission in a shared workspace
  - **Your** costs are covered by **their** Billing project).
- The owner of an existing Terra Billing project can authorize you to use their Billing project
  - You can create/clone a workspace to work in
  - **Your** workspace costs are covered by **their** Billing project
- Use your own Cloud Billing
  - Set up a Cloud Billing account
  - Link to Terra
  - Create a new Terra Billing Project
  - Create or clone a workspace

To browse showcase workspaces (no cost!), see the Showcase workspace library

How to set up a Google Cloud Billing account

- Option 1: Create your own Cloud Billing account on GCP console
  - $300 Google Getting Started credits
  - Institution or personal credit card

- Option 2: Use a third-party reseller (such as STRIDES for NIH or NIH funded researchers or Onix Networking)

To learn more, see: How to set up billing in Terra

For step-by-step instructions to set up a Google Cloud Billing account and link it to Terra, see: Set up billing with $300 Google credits to explore Terra

Session Goals

Revisit Session Goals

- What cloud vendors work with Terra?
- What costs money in Terra and how do you pay for it?
- What is a Terra billing project and how is it connected with a workspace?
- Where can you see how much you've spent in Terra?
- Any recommendations for getting started with running workflows, cloud environments, and securely sharing workspaces?

Revisit Session Goals

- What cloud vendors work with Terra?

Revisit Session Goals

- What cloud vendors work with Terra?

Google Cloud Platform and Azure (soon)

Revisit Session Goals

- What cloud vendors work with Terra?

Google Cloud Platform and Azure (soon)

- What costs money in Terra and how do you pay for it?

Revisit Session Goals

- What cloud vendors work with Terra?

Google Cloud Platform and Azure (soon)

- What costs money in Terra and how do you pay for it?

Compute, data storage, and egress cost money.

Pay with Terra Billing project linked to Cloud billing account.

Revisit Session Goals

- What is a Terra billing project and how is it connected with a workspace?

Revisit Session Goals

- What is a Terra billing project and how is it connected with a workspace?

A feature in Terra that connects a Cloud billing account and a workspace so that costs accrued in the workspace get paid for.

Revisit Session Goals

- What is a Terra billing project and how is it connected with a workspace?

A feature in Terra that connects a Cloud billing account and a workspace so that costs accrued in the workspace get paid for.

- Where can you see how much you've spent in Terra?

Revisit Session Goals

- What is a Terra billing project and how is it connected with a workspace?

A feature in Terra that connects a Cloud billing account and a workspace so that costs accrued in the workspace get paid for.

- Where can you see how much you've spent in Terra?

Combination of the Job History tab & Workspace dashboard. The source of truth is in the Google Cloud console.

Revisit Session Goals

- Any recommendations for getting started with running workflows, cloud environments, and securely sharing workspaces?

Revisit Session Goals

- Any recommendations for getting started with running workflows, cloud environments, and securely sharing workspaces?

Start small. Run workflows you are interested in on one sample first.

See how much it cost. Then scale up.

When you are done using a cloud environment, delete it.

Use groups to share and only share with folks you are comfortable with having the ability to spend $.

Questions?

https://terra.bio

BROAD INSTITUTE