



Callysto

Callysto is managed by Cybera and The Pacific Institute for the Mathematical Sciences (PIMS), and funded by Innovation, Science and Economic Development Canada through the CanCode program.

cybera

cybera.ca | @cybera

 **pims**

pims.math.ca | @pimsmath


With funding from | Avec un financement du

Canada

The Callysto Project

- Targeting K-12 audiences.
- Making open educational infrastructure and learning resources for computational thinking available and accessible.
 - ◆ Such as JupyterHub
- Hosting teacher training workshops across Saskatchewan, Alberta & British Columbia.
- Working with teachers to use Callysto in the classroom.

cybera



Pacific Institute *for the*
Mathematical Sciences

Lots of people involved!

- Workshop leaders
- Notebook Developers
- Project Managers
- System Admins



Lots of people involved!

- Workshop leaders
- Notebook Developers
- Project Managers

→ System Admins






~~25 QUICK & EASY DINNER IDEAS~~
IN 20 MINUTES OR LESS!



**Callysto
Infrastructure**







Easy! Right?

Infrastructure Summary



Existing Work

Syzygy

- PIMS JupyterHub project for post-secondaries.
- Its success helped inspire the Callysto project.
- Re-used and transferred a lot of infrastructure knowledge gained from this project.



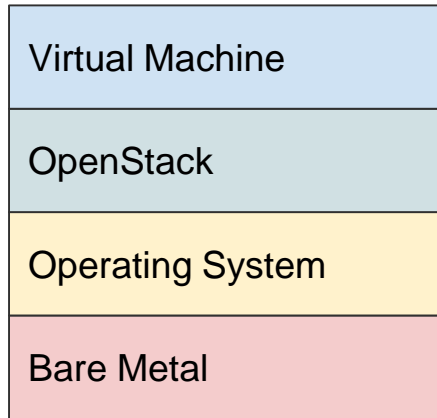
Existing Work

Berkeley

- Data 8: <https://github.com/data-8>
 - Unfortunately, most of this requires Kubernetes (and Google Cloud).
- Jupyter: <https://github.com/jupyter/>
 - <https://github.com/jupyter/docker-stacks>



Cybera already has a cloud





And we didn't feel like adding another stack to our list of responsibilities.




```
[ pttu2u@clavius ~/work/callysto-infra ] $ make help
ansible/exec           Executes $MODULE on $GROUP with $ARGS in $ENV
ansible/get-ipv4       Returns the ipv4 address of $HOST in $ENV
ansible/get-ssh-user   Returns the ssh user of $HOST in $ENV
ansible/hosts/group   Lists the hosts in $GROUP in $ENV
ansible/hosts         Lists all hosts in $ENV
ansible/hosts/playbook Lists the hosts in $PLAYBOOK in $ENV
ansible/list-playbooks Lists plays
ansible/list-tasks     Lists the tasks in a $PLAYBOOK in $ENV
ansible/ping          Pings the hosts in $GROUP in $ENV
ansible/playbook/check Runs $PLAYBOOK on $GROUP in check-mode in $ENV
ansible/playbook       Runs $PLAYBOOK on $GROUP for reals in $ENV
ansible/setup         Runs any first steps when this repository is first cloned
backup                Performs a backup of sensitive data
packer/build/hub       Builds the JupyterHub OpenStack image
quota/get             Gets a quota for $USER in $ENV
quota/set             Sets a quota to $REFQUOTA for $USER in $ENV
ssh/shell             SSH to a given $HOST in $ENV
terraform/apply       Runs "terraform apply" for $ENV
terraform/auto-apply  Runs "terraform apply -auto-approve" for $ENV
terraform/auto-destroy Runs "terraform destroy -auto-approve" for $ENV
terraform/destroy     Runs "terraform destroy" for $ENV
terraform/hub/new/dev  Creates a new dev Terraform $ENV
terraform/hub/new/prod Creates a new production Terraform $ENV
terraform/hub/rebuild Rebuilds an instance and volumes in $ENV
terraform/init        Runs "terraform init" for $ENV
terraform/list-targets Lists available targets in $ENV
terraform/plan        Runs "terraform plan" for $ENV
terraform/setup       Runs any first steps when this repository is first cloned
terraform/show        Runs "terraform show" for $ENV
terraform/taint       Runs "terraform taint" on a $TARGET in $ENV
user/findhash         Finds a hash for $USER in $ENV
```









prod




ruling-donkey




mature-redfish



boss-flounder




flowing-cowbird



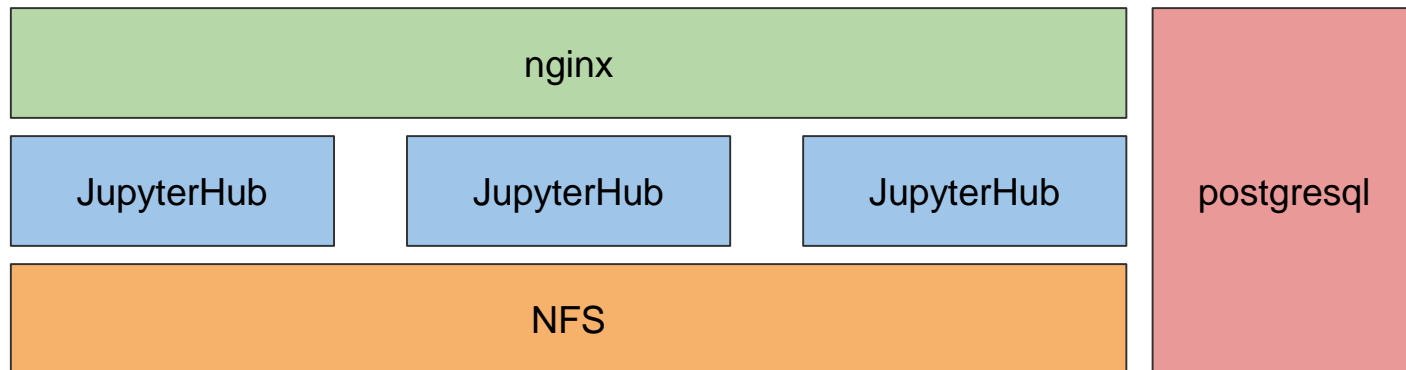
obliging-wildcat

How Do We Scale?

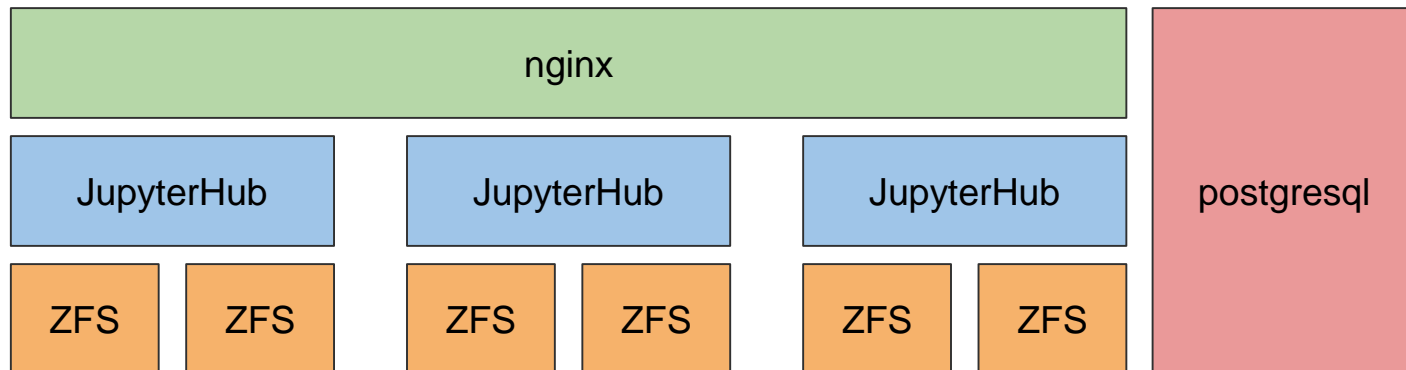


prod


Berkeley / Data 8 Solution



Callysto's Plan A



Callysto's Plan B



Object Storage? How?

- <https://github.com/danielfrg/s3contents>
 - Supports S3 and GCS
- This is a *notebook* driver - not a JupyterHub driver.
- But JupyterHub can pass options to each notebook.




OpenStack Swift

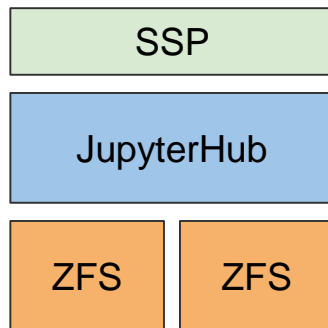
- <https://github.com/callysto/s3contents/tree/swiftfs>
- Pretty easy to do:

```
def mv(self, old_path, new_path):  
    old_path = self.path(old_path)  
    new_path = self.path(new_path)  
  
    self.cp(old_path, new_path)  
    self.rm(old_path)
```

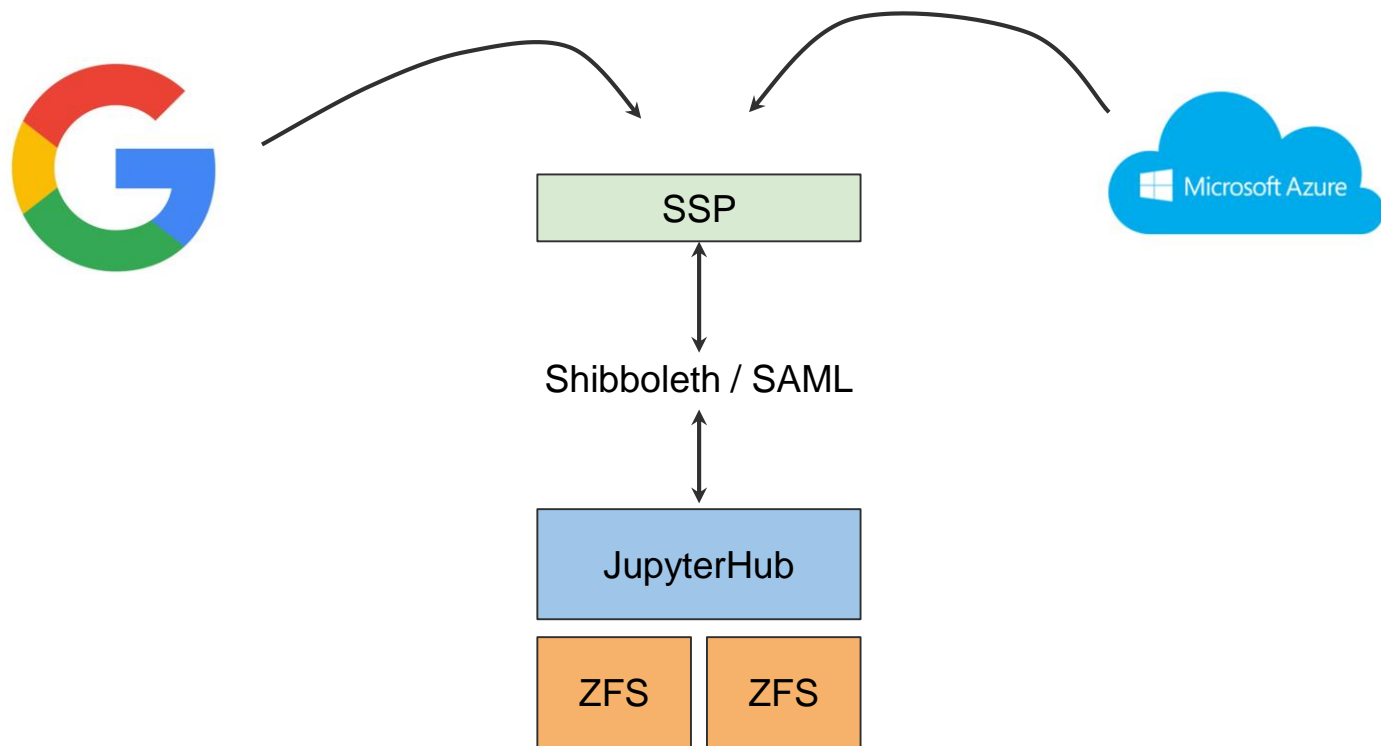
Authentication



Authentication



Authentication



Authentication



Welcome to Callysto!

English | [Français](#)

How would you like to sign in? Please select an authentication source:



Login with Institution account



Login with Google account



Login with Microsoft account

Register [here](#) to be notified of any upcoming service changes, and maintenance and outage windows.

Having trouble logging in? Email Callysto's technical support: support@callysto.ca.



Storing Sensitive Data



Storing Sensitive Data

We don't.

Storing Sensitive Data

- We leverage the TargetedID field in SimpleSAMLphp
- This causes everyone's username/email to appear as a hash:

```
[root@hub pttty2u]# ls /tank/home
0006c2990300b12d208dd667ffd26ce444a7df53 240452b7495ef31e060f8a3810ca06d845bc24dc
00de3b34339953375b237088d5f37ec0773c9355 24692f1b6126dede9322e4ec1cd07997dc806cb6
021d82b0975363a0d969141578415fd86f311639 24c0ddb3929172748fcd5b1f60b9013808780c6
023d546507aaba15b0be4b3b68d143666ff71bda 24c422d10e83d3fb659daa1a7bbbe99934c68d4f
02b77bcb8ae84ef7fe4300593a32a54c97471d0dd 252ff4724be0823de7d3f4f07461d12e965cf10f
02c67985df85f7d72f1652e4c618ae3fa429952d 253fb93069c2f3f7d1529f2f9fcd3408440d3976
0314e6ddfe5804bb4aa26b5e52f83d40943c6e1c 268cdde28ee4a8fbaaceea44f3f126e1ec1c9002
03a6432bd9e5425e8bb38d062d0a072ba99b0f02 269c63602f1799d6c31df6fc7e0416db4c9f544c
045ec0df0a0b123e6a59a5b941c916fecdc4bd41 27609aa22b6196f01eae4749984190eb91cf027d
047244ce2cc4adc31bc3686e7357da0e06902f71 27bd5f0cb32af8aea48a8818f93cf9c676892e36
0473b2bda3d595c236067696d5b8e07297bb5d64 2849472f2538013f6013ef6636446c11f0cb6531
049c3ca372d6bc2a5da1a942cf9a13e90c89786a 29072fb3830f28dc7024c0dbd4e57df59e252b62
04b4bf6ef85328e3a30a79d94df24233c97c665a 29565b4ddc5cb660888f7c421930c053cd3d96a7
04c9b307d84e38862600d7d0666833d3d1b5e4f5 2967d6000e2bb1455545e3b4d40eb863011b3305
```



Storing Sensitive Data

Advantages

- We are not storing any student Personally Identifiable Information.
- Do not have to get consent of every student who logs in.



Storing Sensitive Data

Disadvantages

- Troubleshooting is hard.
 - We have a script which can generate the hash of an email.
- Cannot record rosters (students belong to a class taught by an instructor).
 - We might develop a feature where students submit their work to an instructor. Instructors are adults, so we can more easily store their information.



Future Plans

- Implement our scaling configuration.
- Develop some kind of grading / work administration system.
- Open edX MOOC to provide asynchronous training to teachers (and students) who want to start using or modifying our Callysto / Jupyter learning resources and/or creating their own Jupyter content for their classrooms.



github.com/callysto/infrastructure



Callysto

Thank you!
joe.topjian@cybera.ca

github.com/callysto/infrastructure