

Ham the query monkey

Let me introduce you to ham: <https://en.wikipedia.org/wiki/Ham>

Mmm, ham. but that's a different kind of ham. Our query monkey is also named Ham: [https://en.wikipedia.org/wiki/Ham_\(chimpanzee\)](https://en.wikipedia.org/wiki/Ham_(chimpanzee))

We need a hardworking space monkey to help us! Ham will automatically run a query, and wait for it to complete. If the query fails, it sends a message in a slack channel (now currently #dm-pdac-nagios). This allows the QServ team to know that the query failed. After that, Ham continues on. If Ham gets three queries to fail in a row, Ham will be quiet but will continue to send queries.

You can tell which queries are Ham's because they should all start with SELECT 'monkey' in them.

To get Ham started on running queries, you need access to the NCSA k8s with kubectl.

First, check out this repo, which contains the k8s configuration:

git clone <https://github.com/lsst-sqre/lsst-tap-demo>

Now you can look for the query monkey deployment, which is in the kube/stable folder. Run the monitor.sh script, and it will start the monkey.

<https://github.com/lsst-sqre/lsst-tap-service/blob/master/kube/stable/monitor.sh>

To stop the monkey, you delete the deployment (which will delete the pod running the monkey) by running this command:

```
kubectl delete deployment querymonkey-deployment -n dax-stable
```

To add queries to the query monkey, add your jinja compliant test sql to this folder and make a pull request:

https://github.com/lsst-sqre/lsst-tap-service/tree/master/integration_test/querymonkey/test_queries