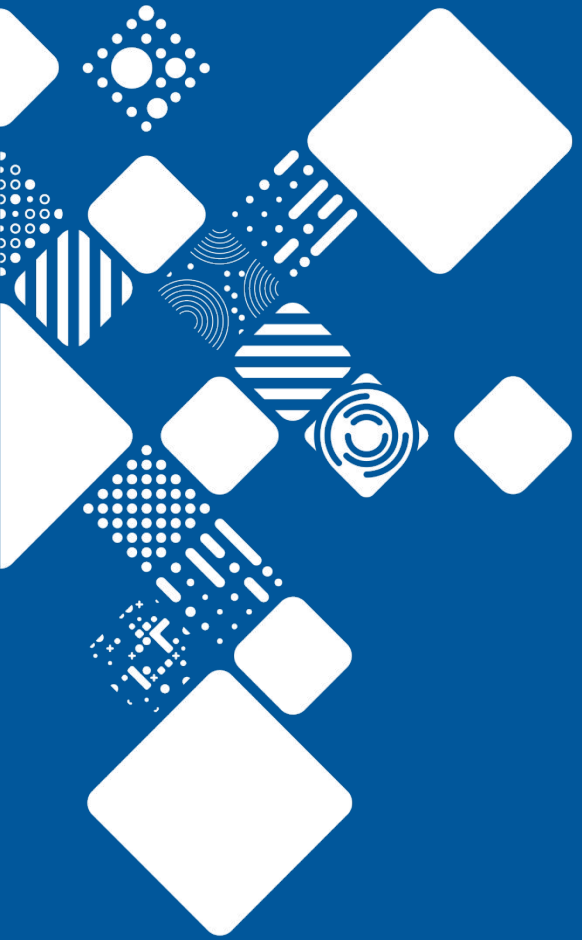

High(ish) Availability For iRODS

John Constable
Informatics Support Group

wellcome sanger institute

Design User Stories

As a Systems Administrator, I want to be able to take an IES down for maintenance without affecting sequencing pipelines

As a Zone Administrator, I don’t want the failure of one server to stop the pipelines

wellcome sanger institute

Current Design

Federated Zones

Federation Master

It’s all on 4.1.10

wellcome sanger institute

Reference Design

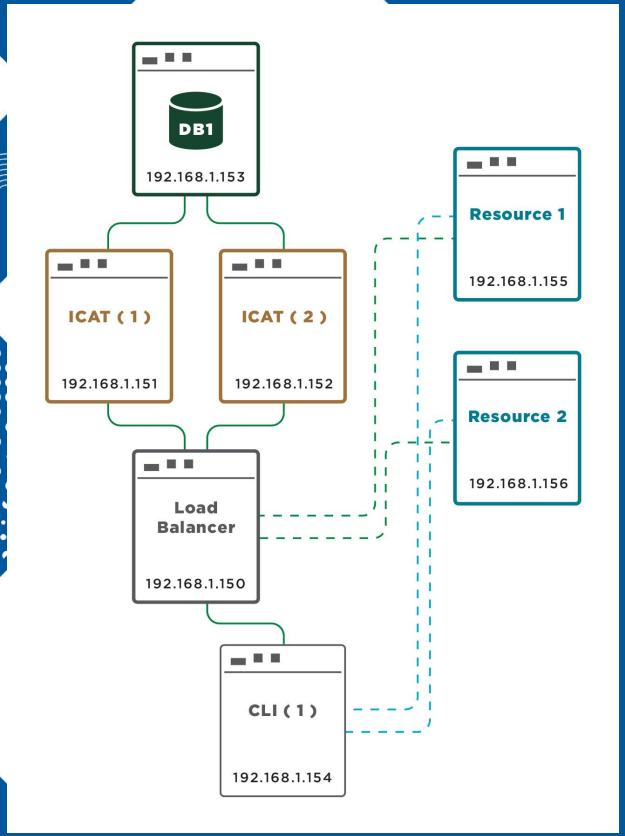

Taken from https://irods.org/2015/07/configuring-irods-for-high-availability/

However, this leaves the load balancer as a single point of failure, if you're not using LBaaS such as on Amazon et al.

wellcome sanger institute

Design

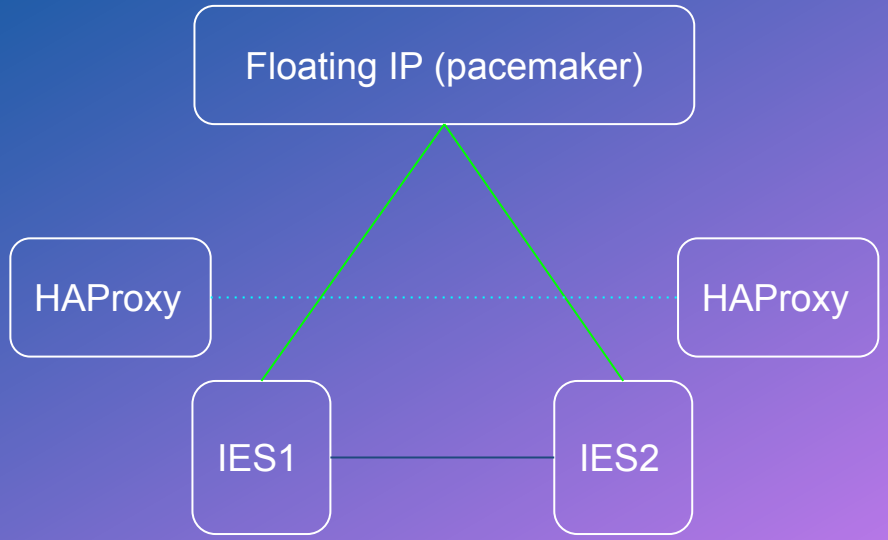

wellcome sanger institute

We ran into a few snags

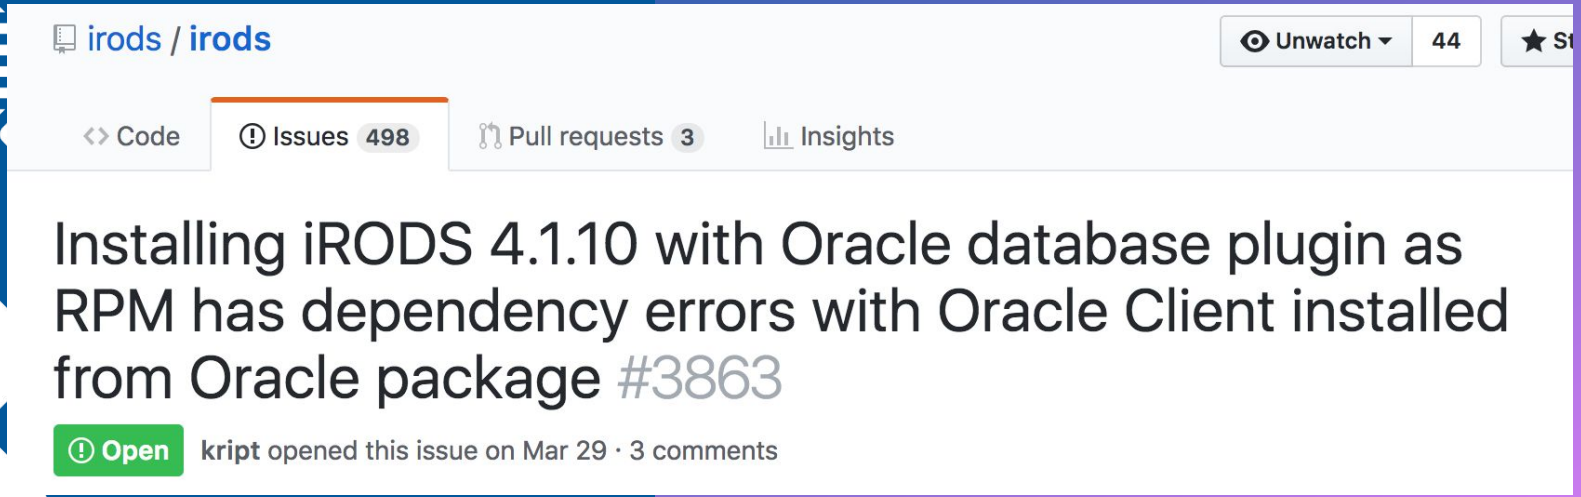

wellcome sanger institute

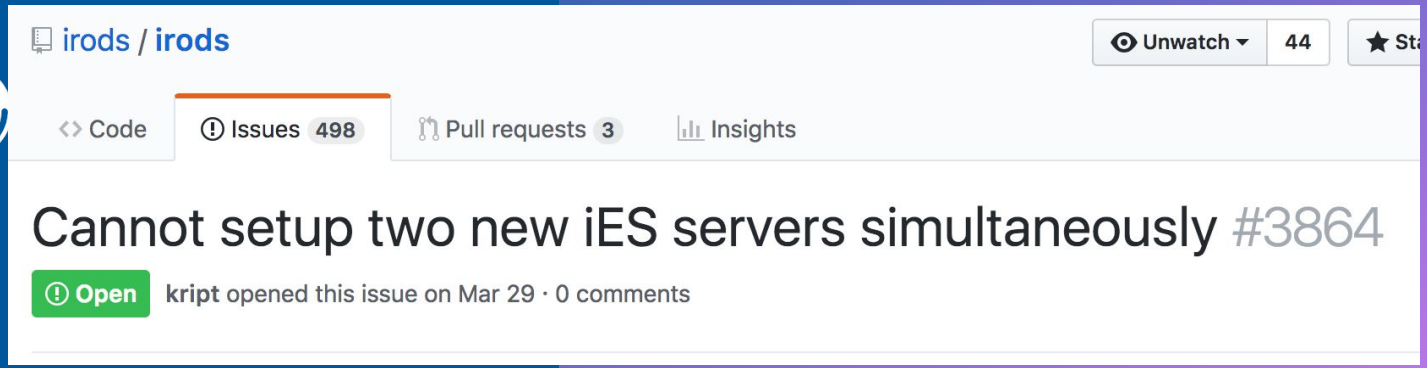

wellcome sanger institute

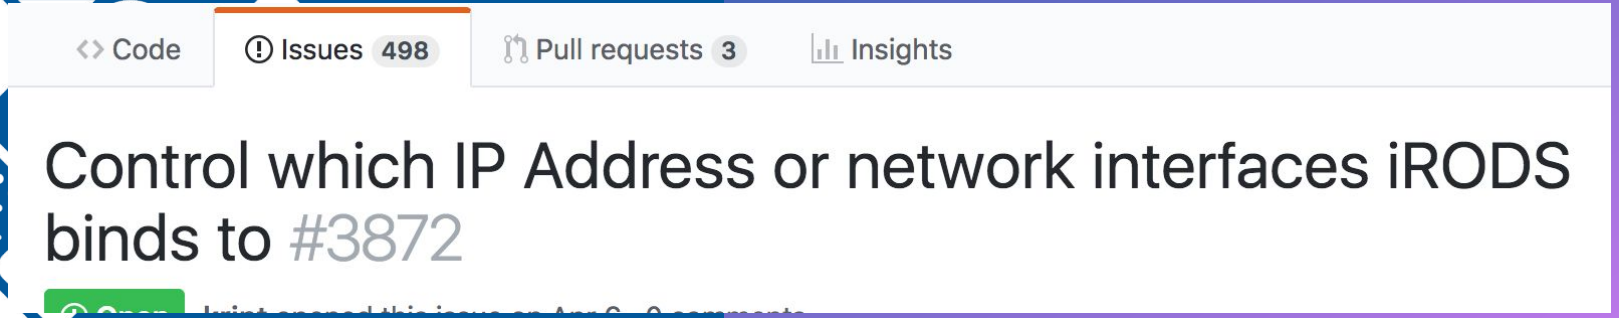
Code
Issues 498
Pull requests 3
Insights
Control which IP Address or network interfaces iRODS binds to #3872

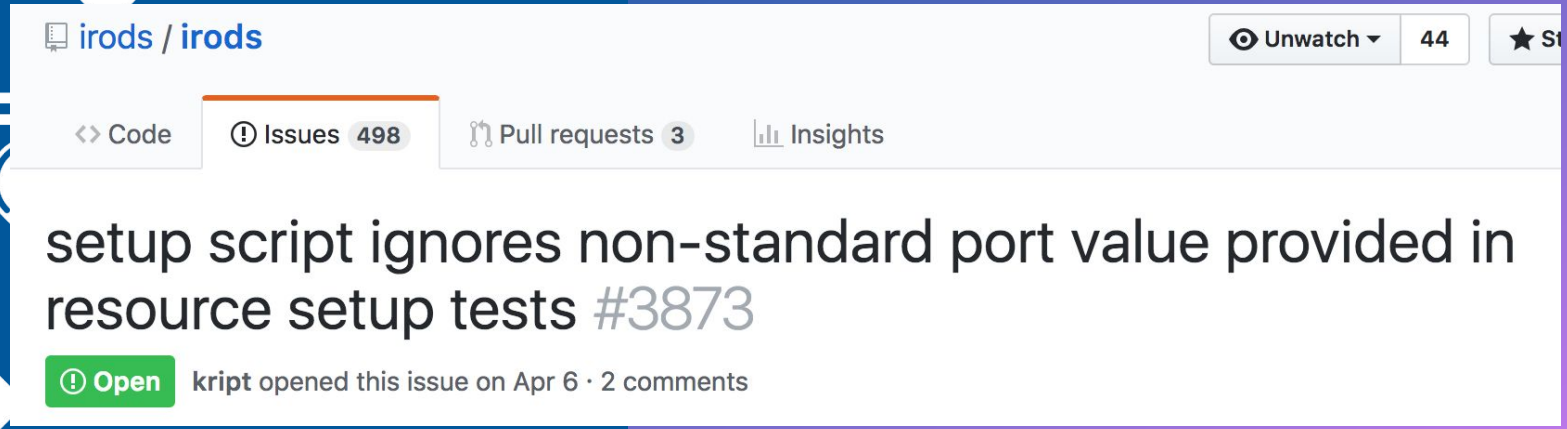

wellcome sanger institute

Not really a snag,
but boy does it make the logs chatty

wellcome sanger institute

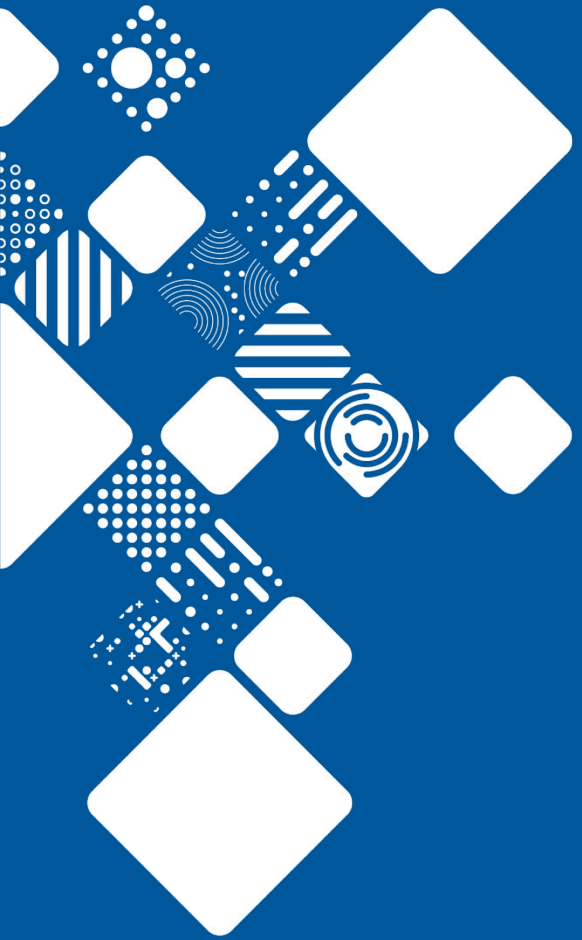

Whatever port you install the IES on,
It expects the rest of the Zone to also be on.

This makes installing HAproxy in front of it somewhat… tricky.

wellcome sanger institute

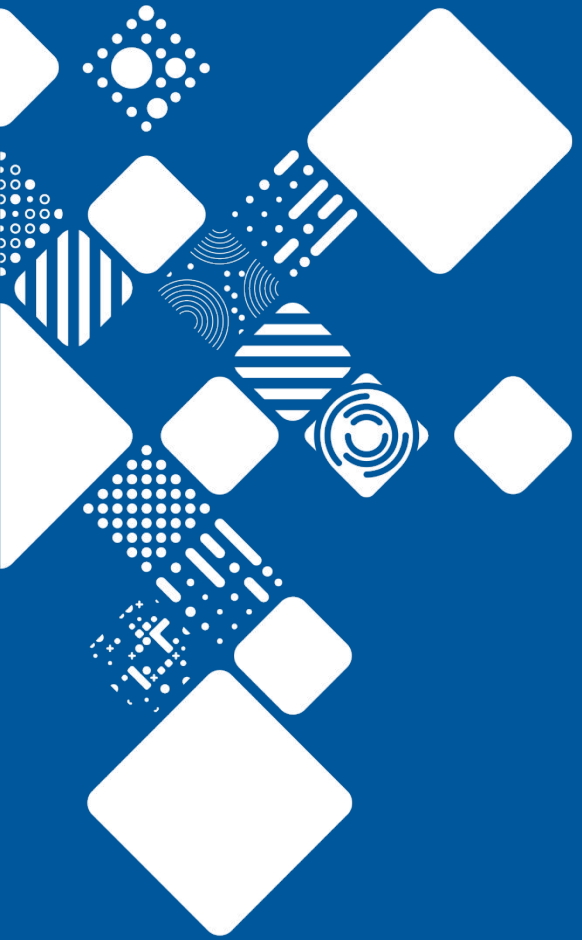

IPtables to the rescue!

wellcome sanger institute

Current Design

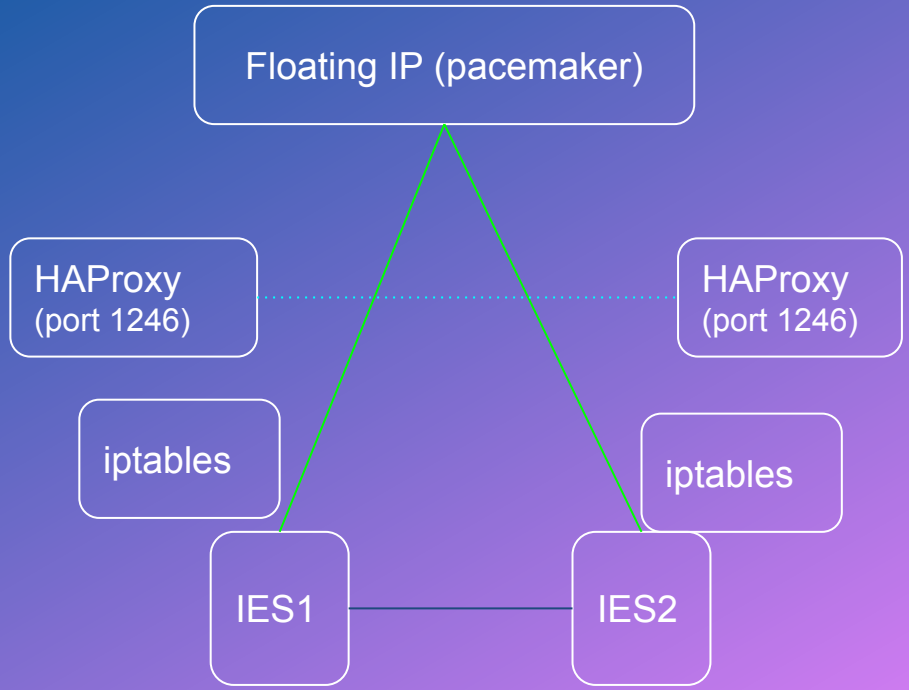

wellcome sanger institute

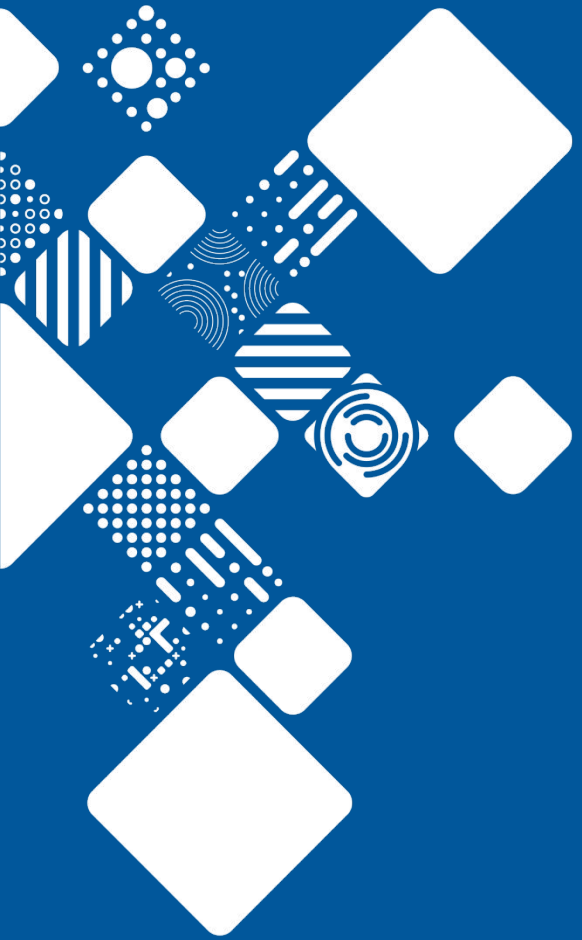

```
- name: "Red Hat | Query firewalld for deployed rules"
  command: firewall-cmd --list-all
  register: firewall_rules
  tags:
    - testing_firewall
    - haproxy

- name: "Red Hat | Forward iRODS Traffic to HAProxy"
  shell: firewall-cmd --add-forward-port=port=1247:proto=tcp:toport=1246:toaddr={{ ip4_backend }} --permanent
  when: "'proto=tcp:toport=1246:toaddr=' not in firewall_rules.stdout"
  tags:
    - testing_firewall
    - haproxy
  notify: "Restart firewalld"
```

wellcome sanger institute

Performance Benchmarking

wellcome sanger institute

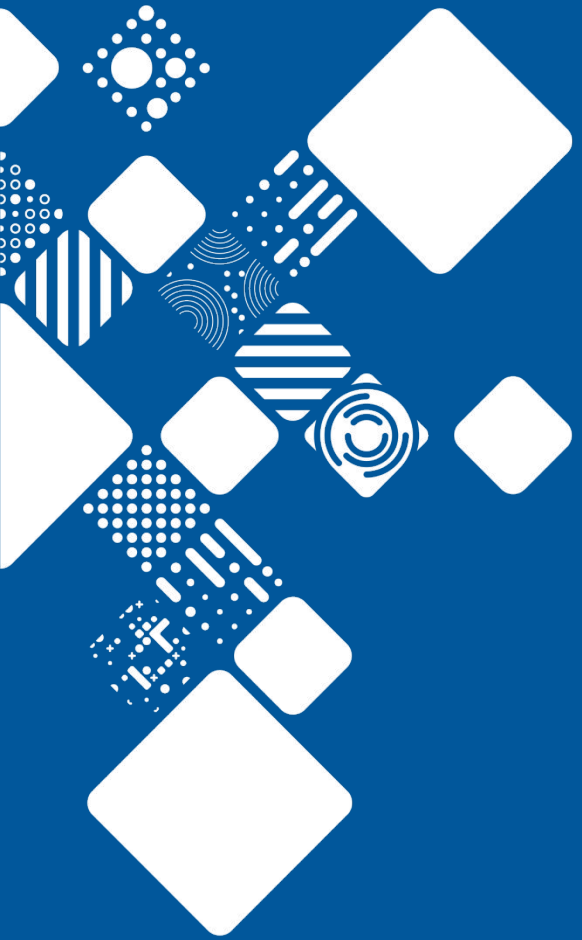

Old Infrastructure;

10G (1 on VM's until recently))

Federation master: 8G RAM, two vCPU's

Largest Zone IES: 8G RAM, 8 CPU

New Infrastructure;

Two servers

front and back end 10G (legacy 25), 520G RAM

SSD

wellcome sanger institute

Using a variant of Terrell's Benchmarking script

https://irods.org/2016/09/irods-4-1-9-networking-performance-whitepaper/

wellcome sanger institute

Unfortunately, this is something of a Work In Progress

Only a small set of test data, from four servers

wellcome sanger institute

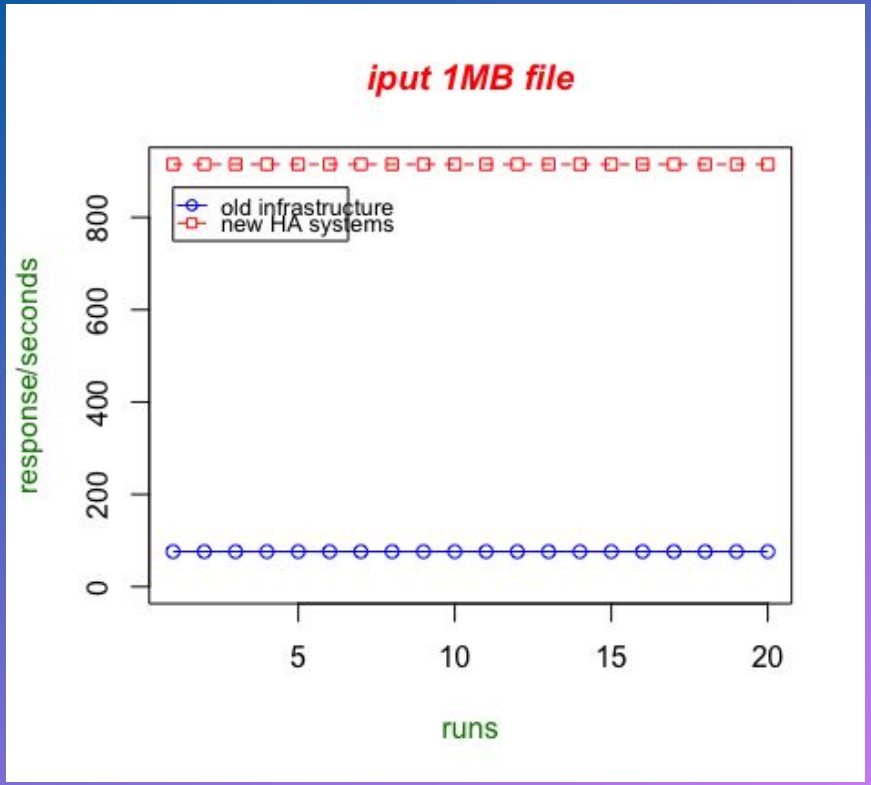
iput 1MB file
response/seconds
0
200
400
600
800
old infrastructure
new HA systems
5
10
15
20
runs

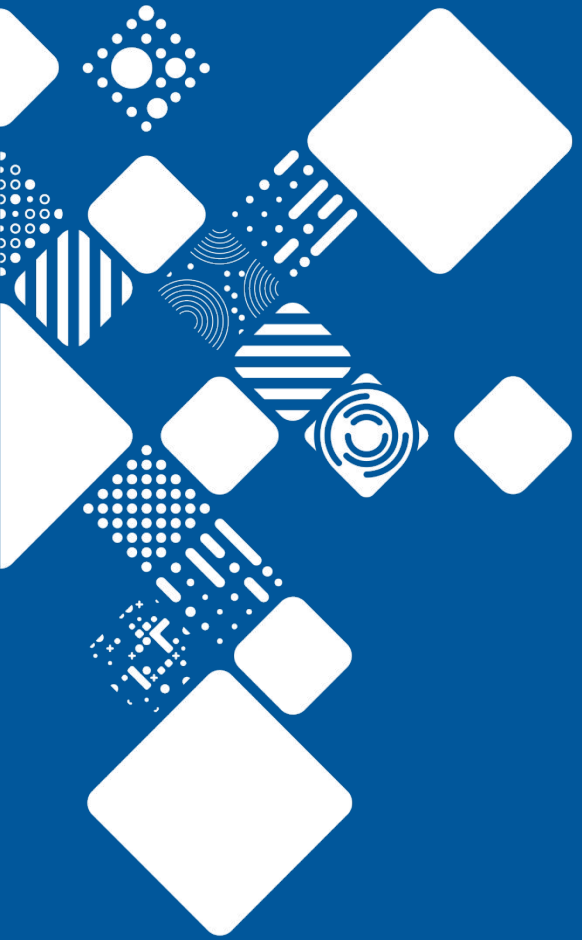

We suspect routing

Or HAProxy.

It’s some kind of misconfiguration on our end, likely.

wellcome sanger institute

To Be Continued?

wellcome sanger institute

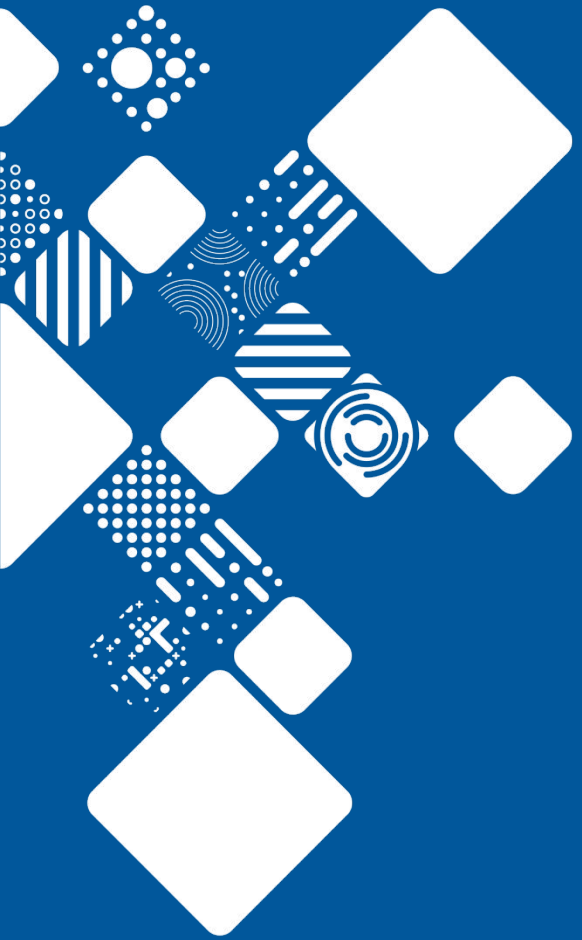

Thank you for listening!

john.constable@sanger.ac.uk

@kript

wellcome sanger institute