
Experience with Xrootd in a small community

Ian Fisk

Community

- ❖ For the last 6 months I have been trying to support data management and access for a community of genomic researchers
- ❖ There are more sources of data than HEP, but the most advanced research sequencing machines are expensive and there are only a few
- ❖ Communities collect data from individuals based on what they are trying to study
 - ❖ Regions, conditions, illnesses, etc.

File Access

- ❖ The machines produce a reasonable standard raw data format called BAM files
- ❖ ~10GB file per person for whole exome (A study of about 1% of your DNA) Mutations here can have severe impact on the rest
- ❖ >100GB file per person for whole genome sequencing. Modern machines can sequence the entire genome
- ❖ A study of a group might be a few hundred individuals
- ❖ Raw data in the few TB range for exome and few hundred TB for full genome

Processing

- ❖ The processing steps are to apply a series of defined steps
- ❖ The code for each step is often written by an independent party and shared
- ❖ Techniques and code improve, groups often go back to the source data and reprocess

Kinds of Processing

- ❖ Align the data
 - ❖ Make sure you have the region you were looking for
- ❖ Variant Calls
 - ❖ Compare your sample to an agreed standard.
Identifying the number of places that your sample varies from some “nominal”
- ❖ Then start looking for correlations
 - ❖ Frequently work is done with groups of related individuals

Differences with HEP

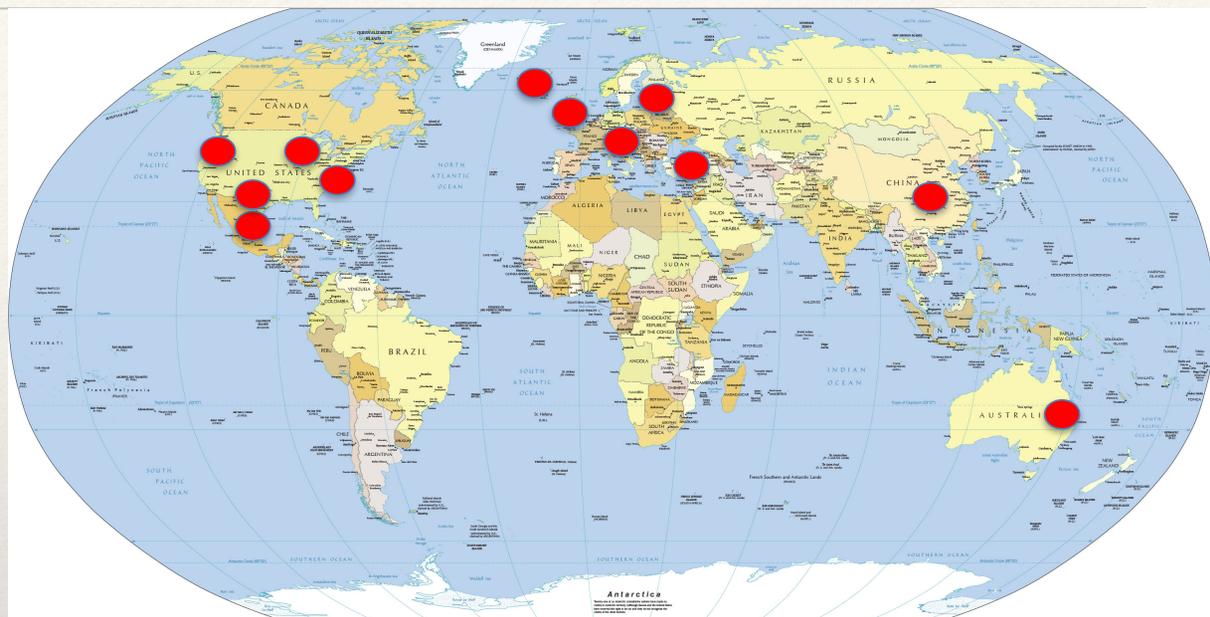
- ❖ Software is written by external people and there is a reluctance even to change where the default install it, let alone build in new IO libraries
- ❖ Solutions that don't require recompilation like `LD_PRELOAD` and `FUSE` may get better adoption
- ❖ Steps are well defined and can often be seen as canned applications
- ❖ The steps are updated so people do a full reprocessing from the raw

Size of the problem

- ❖ Groups were encouraged to host data in Amazon
 - ❖ 0.5PB even in the little local community
 - ❖ Advantages is that it's well understood how to serve to EC2 processors
 - ❖ Disadvantages is that it's expensive to export from Amazon and the government got nervous about storing sequencing data and decided to close the facility down and all data needed to come out

Current Distribution

- ❖ Currently this genomics community uses FNAL as an archival system
- ❖ Recently imported ~400TB of data primarily from S3
- ❖ 2 100TB samples were exported from FNAL using GridFTP to Iceland and Oregon for additional processing



- ❖ The community created and made publicly available 11TB of diversity project data
- ❖ 300 people from all over the planet

The Challenge (1/2)

- ❖ There are about 40 entities that want samples between 10-100TB
- ❖ There is no real infrastructure for data management
 - ❖ File lists are sent with checksums in manifests
- ❖ These are labs with firewalls and data has grown much faster than expertise, so little community knowledge for how to move big samples around
 - ❖ Bare GridFTP is not completely user friendly nor is the entire grid certificate infrastructure

The Challenge (2/2)

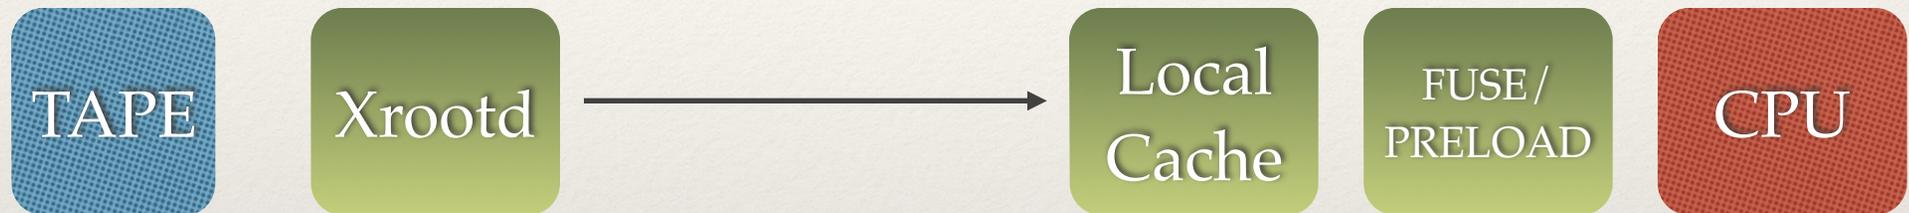
- ❖ Path through the files in question is a semi-pathological
- ❖ Unpacking internal buffers and retrieving objects across large swaths of the file
 - ❖ Access through the file during analysis is not linear and applications know nothing of training or pre-fetching

Hoarding

- ❖ Unlike a big central detector project where everyone is a member of the collaboration with rights to the data, these communities often share samples with detailed written agreements for specific periods of time
- ❖ The convenience of being able to access anything will be weighted against the certainty of having the samples under your control

Deployment

- ❖ During the spring we will begin deploying something very simple



- ❖ Looking forward we have the potential of placing intermediate services for data serving
 - ❖ Currently the fully public samples are 10-15TB and could be replicated

Looking Forward

- ❖ We will be deploying as a poor person's content delivery network
- ❖ The community is smaller than HEP, but equally distributed, more poorly supported for resources, and lacking expertise in handling large data samples
- ❖ Data is growing by a factor of 10 with the next generation of machines and datasets are reaching interesting sizes