Kubernetes Advanced Practical

Batch Processing

David Yuan, Ph.D.

Cloud Bioinformatics Application Architect

Technology and Science Integration

European Bioinformatics Institute, EMBL

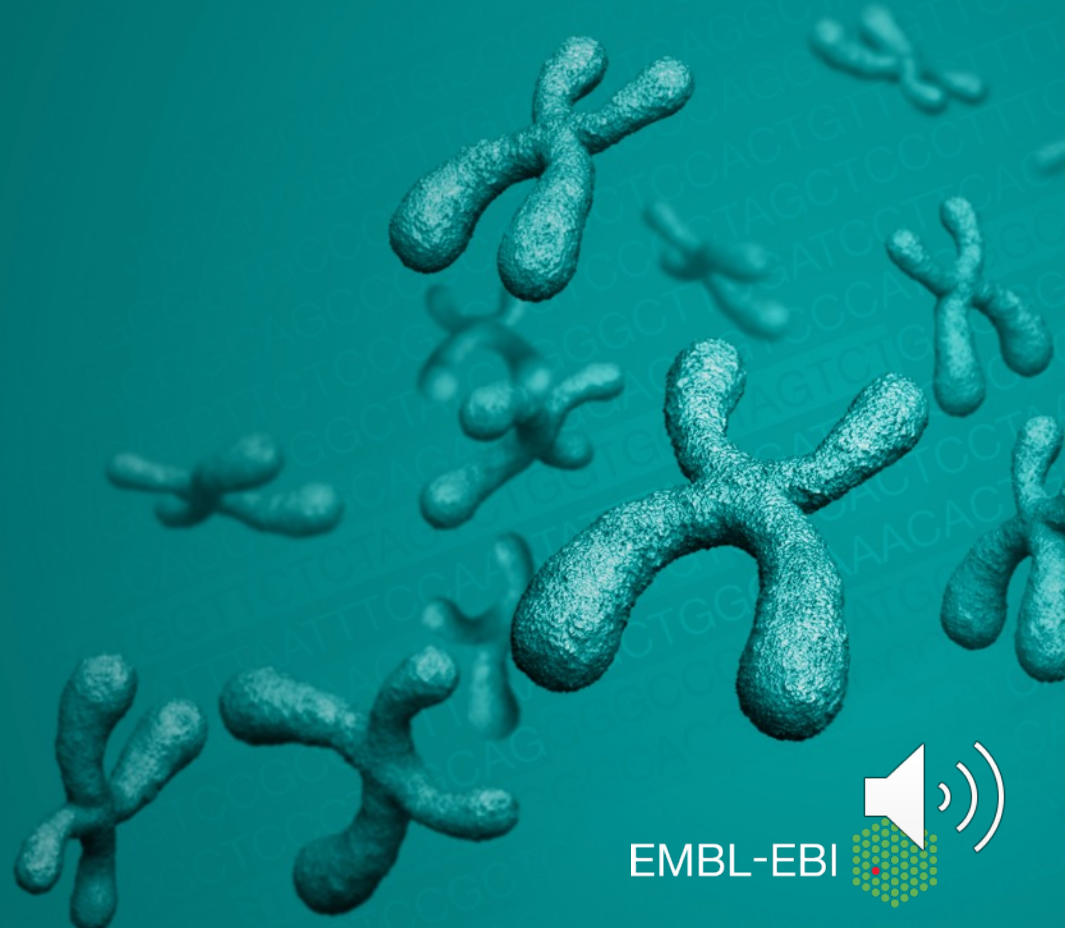

EMBL-EBI

Kubernetes Advanced Practical (batch processing)

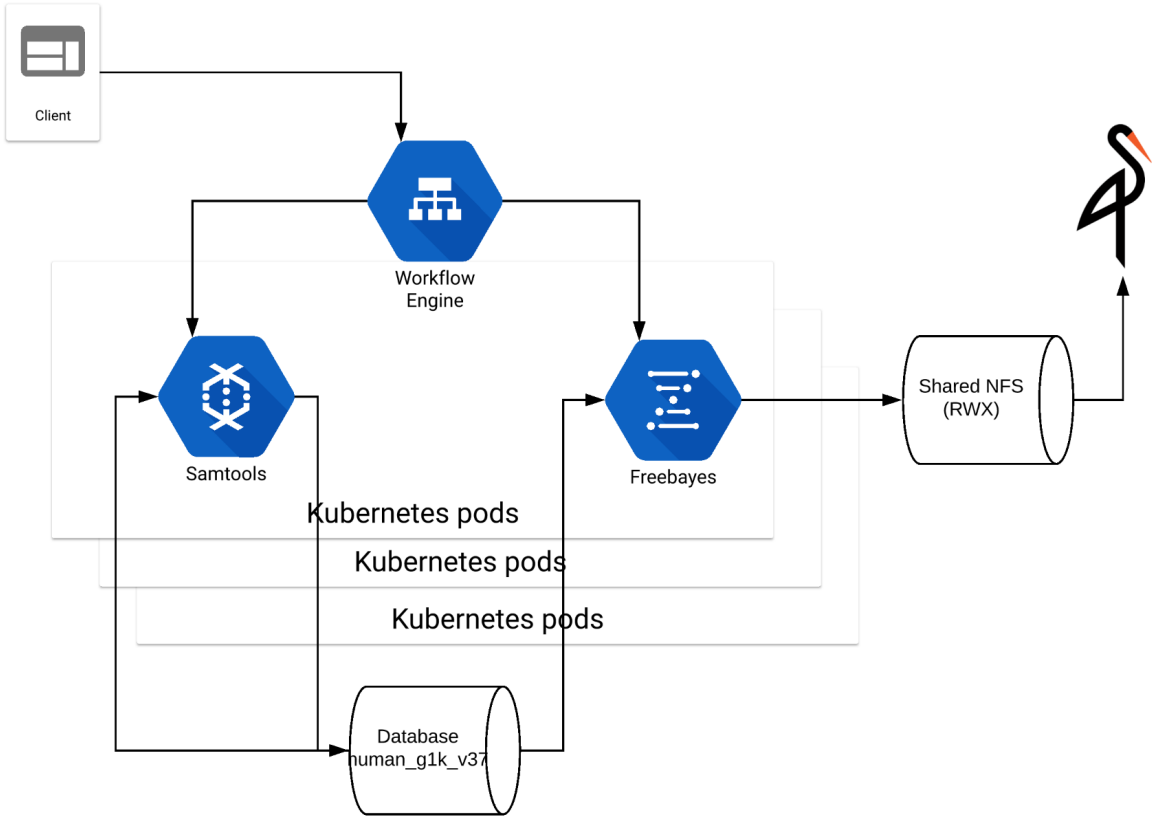

- Reading 0: Preparing Minikube VM to support NFS volumes

Kubernetes Advanced Practical (batch processing)

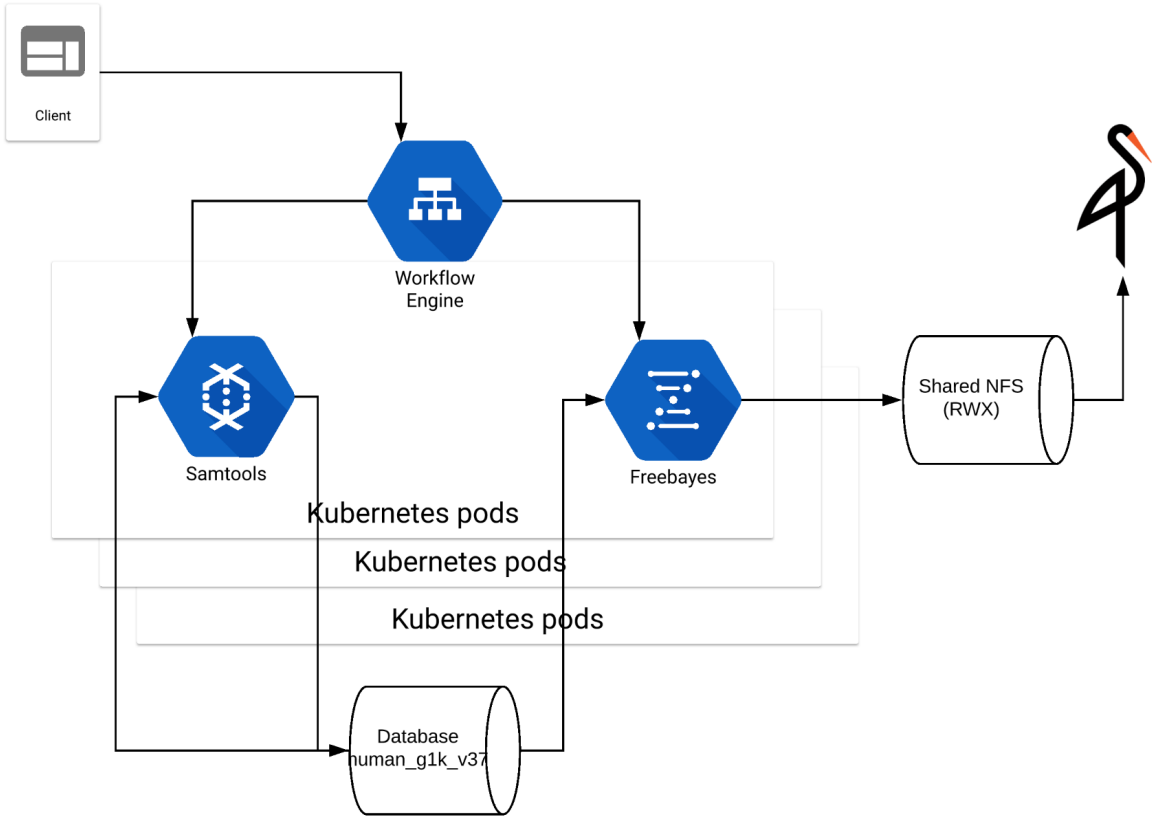

- Reading 0: Preparing Minikube VM to support NFS volumes
- Accessing large amount of data from the source:
  - Exercise 1: ReadWriteMany for shared output
  - Exercise 2: ReadOnlyMany for data source
  - Exercise 3: ReadWriteOnce for private workspace
  - Exercise 4: Initialising persistent volumes
  - Exercise 5: Kubernetes secret & S3 interface

Kubernetes Advanced Practical (batch processing)

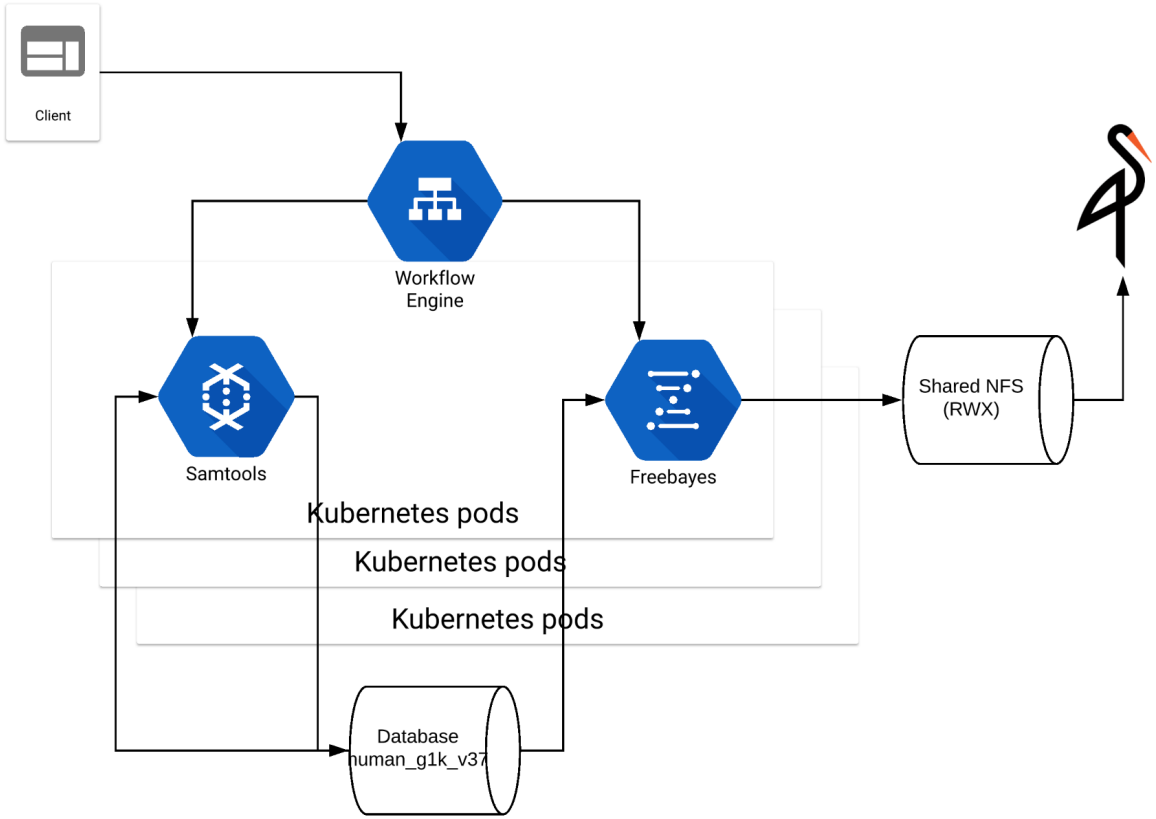

- Reading 0: Preparing Minikube VM to support NFS volumes
- Accessing large amount of data from the source:
  - Exercise 1: ReadWriteMany for shared output
  - Exercise 2: ReadOnlyMany for data source
  - Exercise 3: ReadWriteOnce for private workspace
  - Exercise 4: Initialising persistent volumes
  - Exercise 5: Kubernetes secret & S3 interface
- Scaling up:
  - Exercise 6: Horizontal scaling
  - Reading 1: Vertical scaling
  - Exercise 7: Autoscaling

Kubernetes Advanced Reading

- Reading 1: Bringing compute close to data
- Reading 2: Resource consumption
- Reading 3: Docker builds
- Reading 4: CI/CD toolchain