
Offloading work to a remote HPC using Globus and funcX

Chris Scott (NeSI)
chris.scott@nesi.org.nz



Background

- NeSI (New Zealand eScience Infrastructure)
 - Provides HPC infrastructure and expertise in NZ
- Two clusters (Mahuika and Māui)
- Research Software Engineering Service
 - Available to most NeSI users
 - Embed RSE in research group
 - Short term projects



Use case

- Pharmacology group at the University of Auckland
 - Improving understanding of medicines in humans and improving dosing
 - Want to manage their research from familiar Windows environment
 - Parts of their workflow require significant resources
 - Non-linear mixed effects modelling
 - Fortran+MPI code
 - Bootstrapping

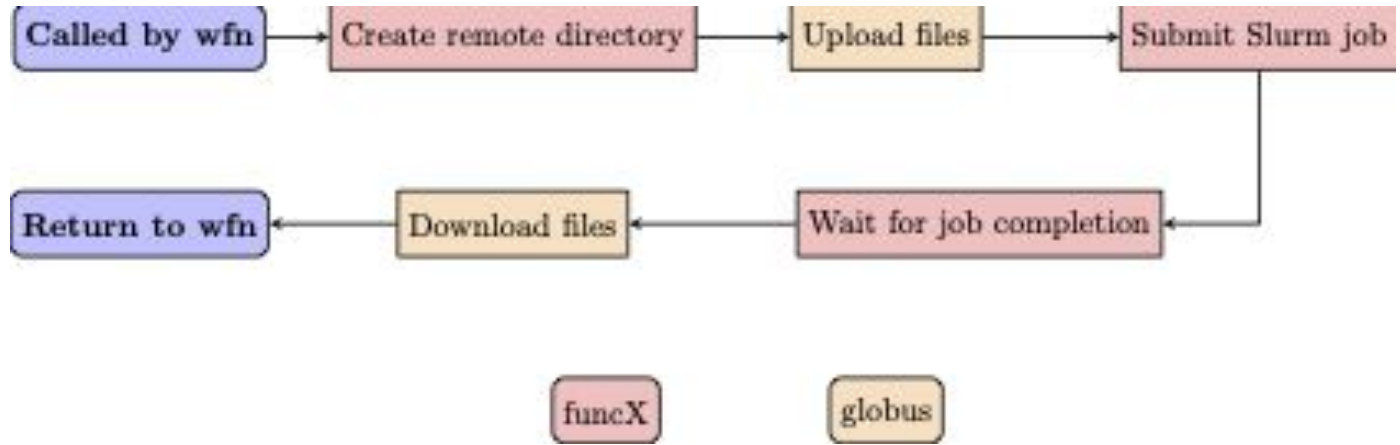


Remote Job Manager tool, 2nd iteration

- Reimplemented in Python using Globus and funcX
 - Globus authentication vs 2FA
- Provides an API that the researchers batch scripts (wfn) interface with
 - File transfer
 - Globus HTTPS transfer to Guest collection
 - Running commands
 - funcX endpoint running on a login node

<https://github.com/chrisdjscott/RemoteJobManager>

How we are using funcX



funcX is mainly used to submit Slurm jobs and check for their completion

Deploying the funcx endpoint (1)

- funcX default endpoint runs on a login node, but...
 - NeSI has two login nodes, sometimes one goes down or is rebooted for maintenance
 - Researchers don't want to have to look after the endpoint themselves
 - No root access, systemd, etc

Deploying the funcx endpoint (2)

- Bash script running as crontab job
 - Connects to both login nodes and checks if an endpoint is running one of them
 - If not, start a new endpoint on the current primary
 - Runs hourly as a small Slurm job

```
#SCRON --time=08:00
#SCRON --job-name=funcxcheck
#SCRON --account=nesi99999
#SCRON --mem=128
@hourly /home/csco212/.funcx-endpoint-persist-nesi.sh
```

Summary

- Using funcX to offload work to a remote Slurm cluster
 - Globus authentication works well
 - Each user having to manage their own endpoint introduced some difficulties
- Research group have been testing this for 6-8 months now and working pretty well
- Other possible use cases
 - Offloading instrument data for processing