CVMFS on Microsoft Azure

September 2022, hugo.meiland@microsoft.com

Who am I?

- Supported HPC and research IT for ~10 years at Dutch University in chemistry and biosciences
- Supported National Research Supercomputer and Meteo (and others) for ~5 years
- At Microsoft from TSP in the field to Azure Engineering HPC
  - Working with customers to unblock/progress
  - Bringing back feedback and learnings on how to improve
  - Educating internal teams on HPC requirements

HPC at Microsoft

- From explaining we actually run Linux to
  - #10 in Top500 (Nov 2021)
  - #26,27,28,29 in Top500 (June 2021)
  - Supercomputers for UK Met Office and Meta/Facebook

- Specialty HPC SKU's
  - 5/6 generations of CPU, 7-9 generations of GPU
  - Nvidia/Mellanox InfiniBand in most of these VM's

- 2 main scenarios:
  - Lift & shift: traditional schedulers Slurm, PBS, LSF & posix filesystems
  - Cloud native: serverless & object/blob storage

CVMFS client

- Client (with EESSI config)
  - Integrated into Az-HOP (HPC Ondemand Platform)
  - Testing in WSL2, lack of automount is still a gap

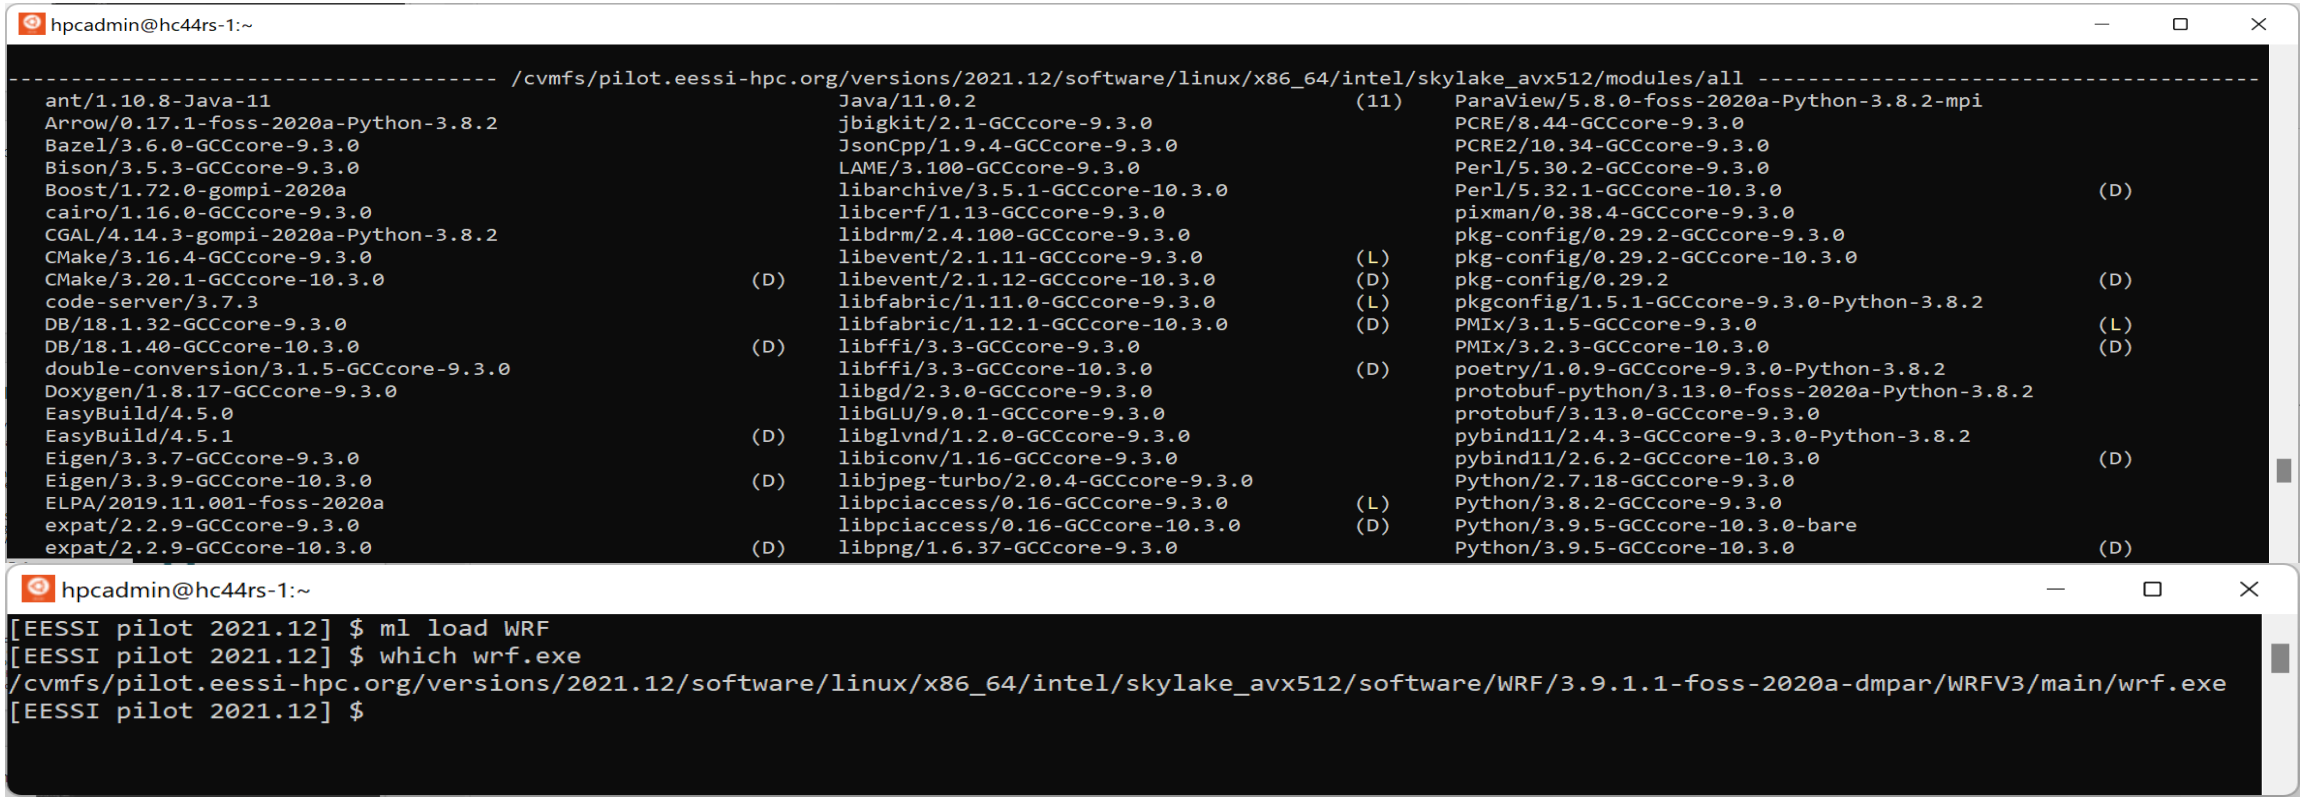

CVMFS repo

- Several tries based on EasyBuild: arch vs skus? which OS? which compilers?
  - Our team can only run “best effort” services
  - Figuring out customer expectations
    - Do you really need compiler XYZ / flag PQR?
    - Or is ml load <application>, mpirun <application> compelling enough?
    - Better to build containers?
- Tracking / (working on) EESSI:
  - Integrated into Az-HOP today
  - Need to double down on GPU and end-2-end runs (upcoming hackathon!)

CVMFS code

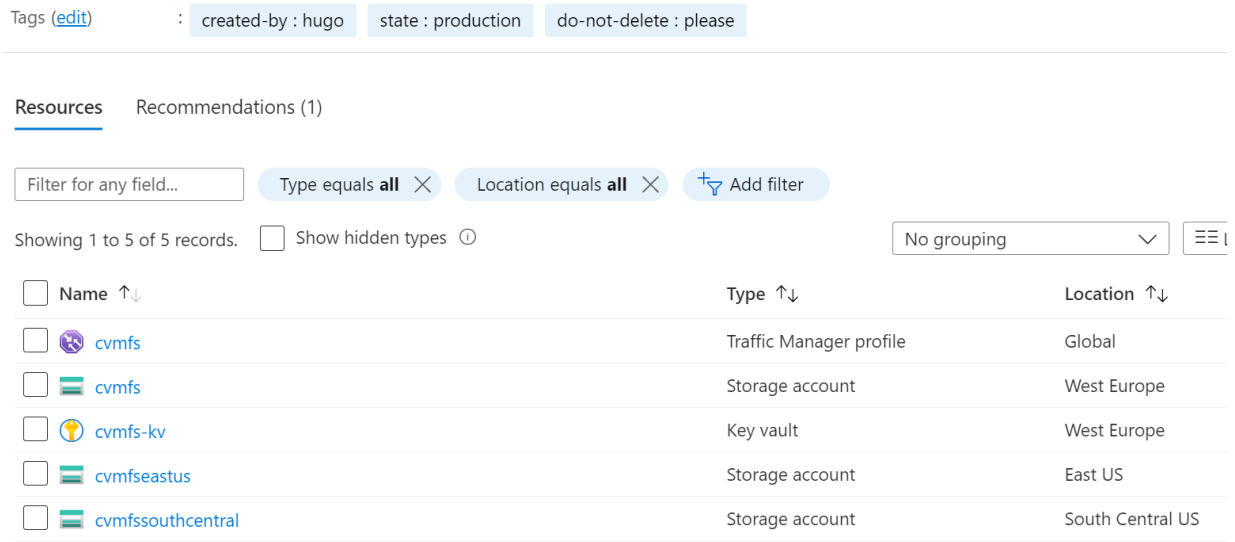

- Added support for Azure Blob next to S3 (okt 2021?)
- Cvmfs_server can build stratum0 directly on Azure Blob
  - With keys in keyvault, build machine is expandable

- Using traffic manager i.s.o. geo-ip
- Sync containers i.s.o. stratum0 -> stratum1

Next steps

- Extending work with/on EESSI with focus on WRF & MD
  - To better support end-to-end in Az-HOP
- Dive deeper and pick up learnings on container deployments
  - As alternative to container registry for HPC
- Happy to discuss/test/extend serverless CVMFS

Thank you!

hugo.meiland@microsoft.com