

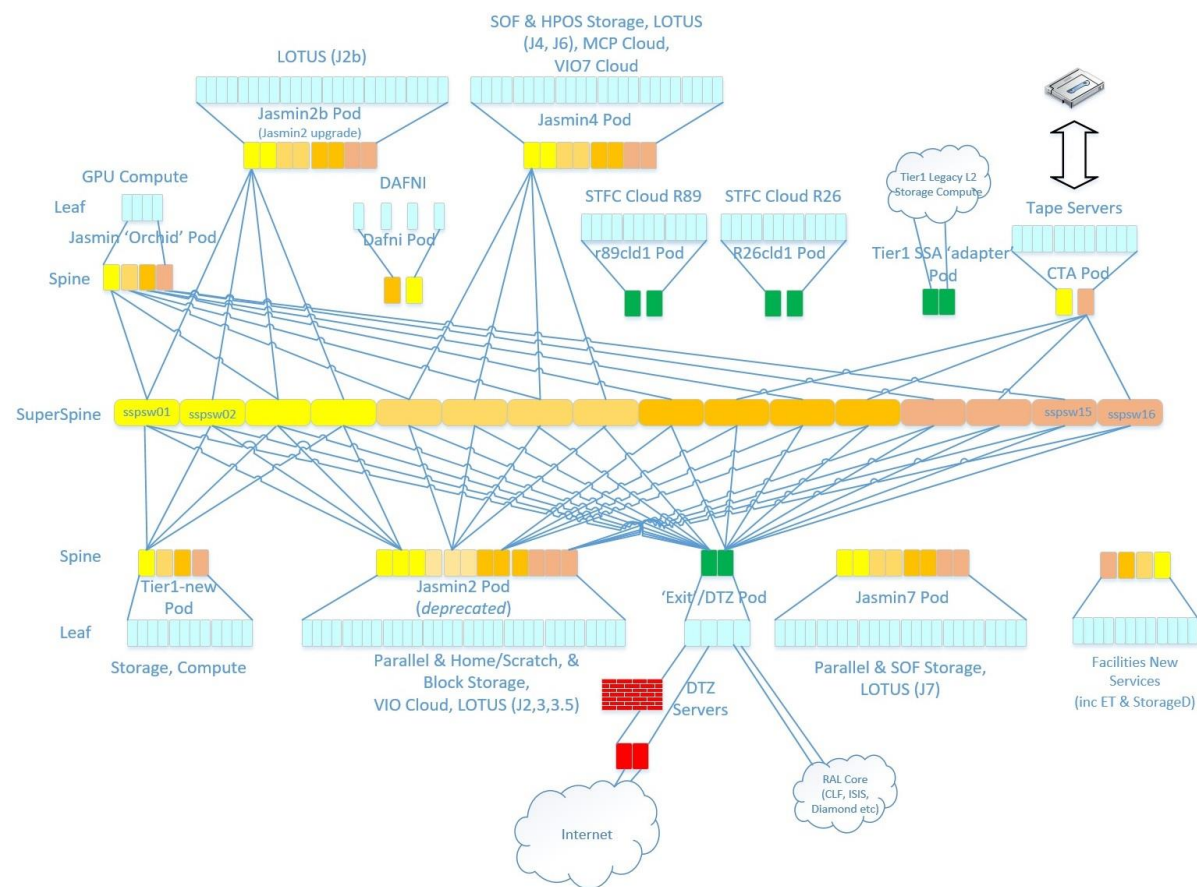
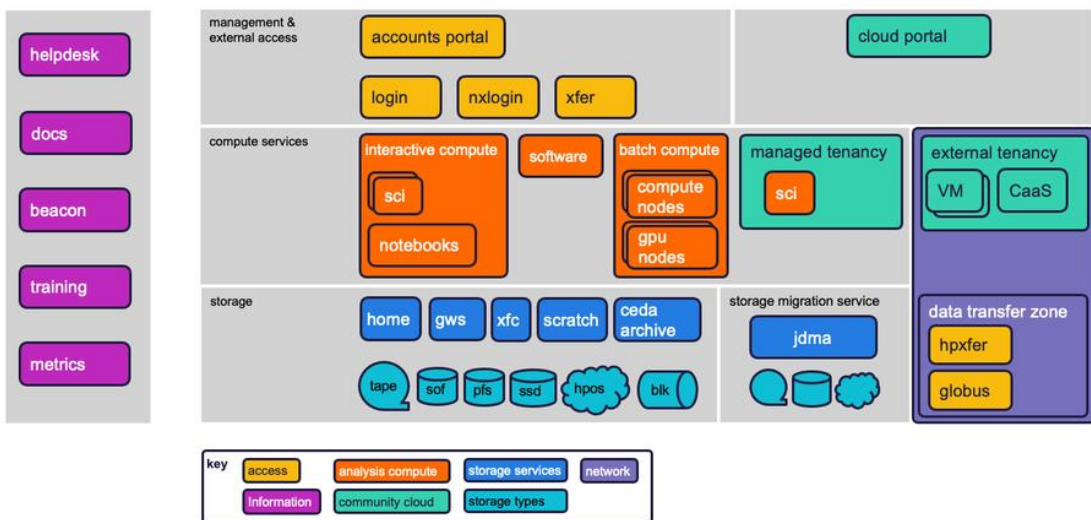
JASMIN

James Hannah – March 2023

Introduction

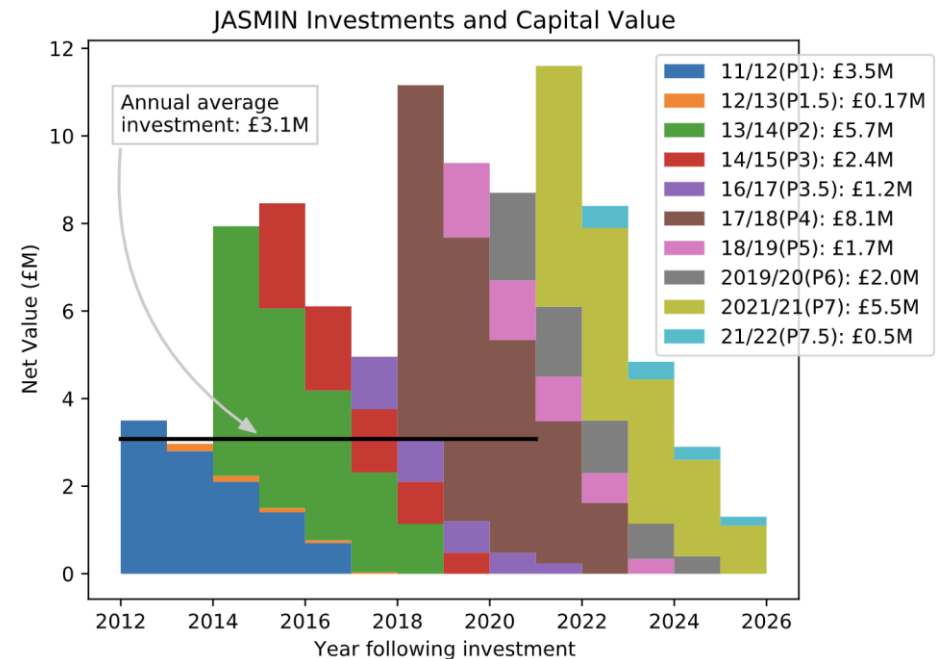
- JASMIN is a data analysis platform for environmental science, funded by NERC and delivered by STFC
- We operate a range of services including a batch cluster, various storage clusters, ancillary services and a very high capacity (10-100Gbps) network
- Approx. 600 servers, mostly in the HTC/HPC (Slurm) cluster (~20k cores)
- Infrastructure team is only 5 people
- ~60PB disk storage, comprising scale-out filesystem, SSD, S3, Parallel file system
- ~50PB tape
- VMWare and OpenStack (VIO) clouds

JASMIN Services & Network



Procurements - past

- JASMIN tends to run large procurements every year to replace tranches of older hardware when capital is available, generally towards the end of the year, averaging around £3.1M annually
- Sometimes alternating between storage/CPU in alternate years
- Huge amounts of network expansion in terms of ports/bandwidth – network is a huge asset



Procurements – future?

- Likely to investigate CPU procurement in 2023
- Likely to investigate adding a new network “pod”
- Users testing PoC GPU cluster (44xNvidia A100), may see increased demand in future
- Starting to think about future storage options post-2023 - for bulk scale-out filesystem layer (approx. 40PB)
- Net-zero is looming on the horizon, power/carbon numbers will be needed at some point including from manufacture
- Interested in automation where possible – APIs & SDKs