

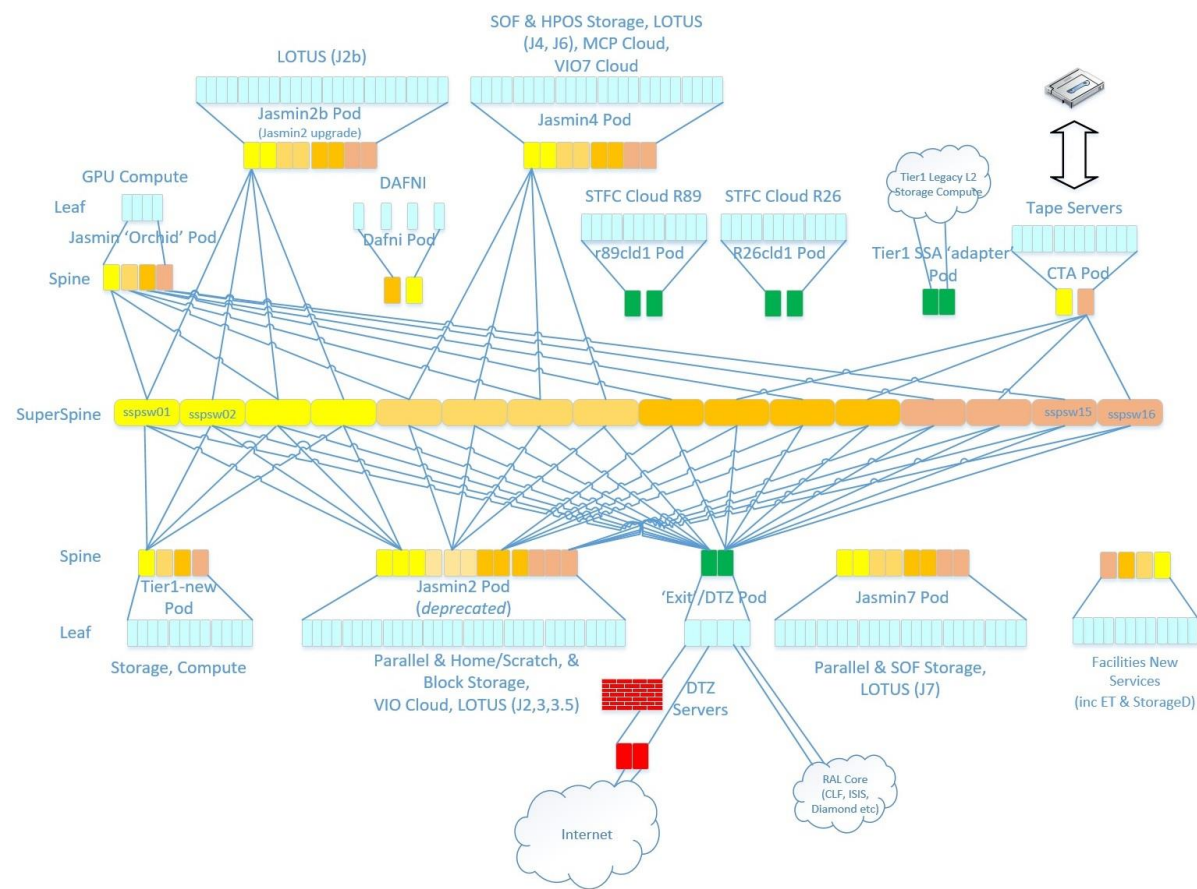
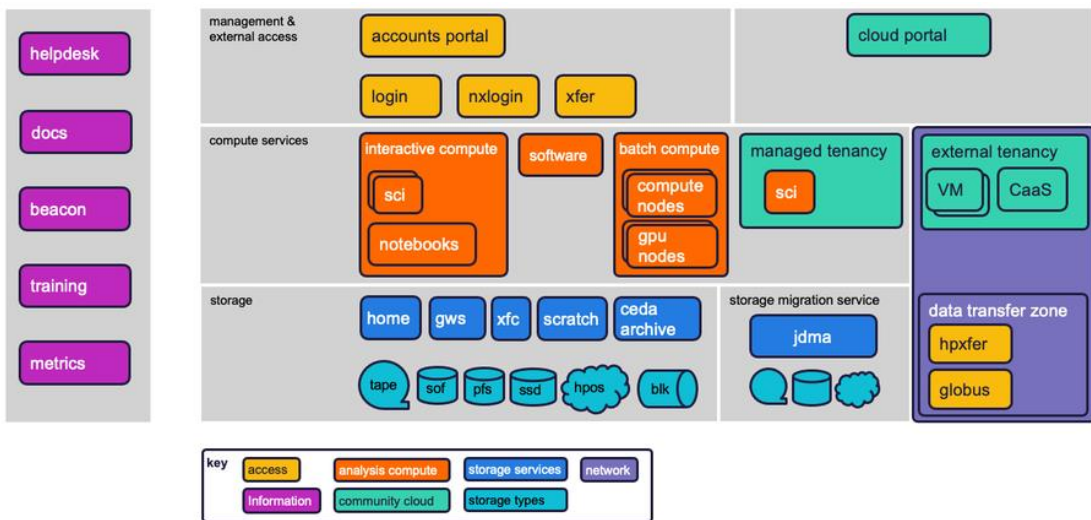
JASMIN

James Hannah – SSSNA Vendor Day – 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. 800 servers, mostly in the Slurm cluster
- ~60PB disk storage, comprising scale-out filesystem, SSD, S3, Parallel file system
- ~50PB tape
- VMWare, bare metal and OpenStack (VIO)

JASMIN Services & Network



Procurements

- JASMIN tends to run large procurements to replace tranches of older hardware when capital is available, generally towards the end of the year.
- Moving towards a more long-term capital spend strategy
- Sometimes alternating between storage/CPU in alternate years
- Networking bandwidth is always a high priority.
- 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

Procurements

- Users testing PoC GPU cluster (44xNvidia A100), may see increased demand in future
- Likely to investigate CPU procurement in 2023
- Likely to investigate adding a new network “pod”
- Starting to think about future storage options post-2023 - for bulk scale-out filesystem layer (approx. 40PB)