Karlsruhe Institute of Technology

Experiences with Lustre, Storage Scale and BeeOND at KIT

Roland Laifer

STEINBUCH CENTRE FOR COMPUTING - SCC

HPC and parallel file systems at KIT

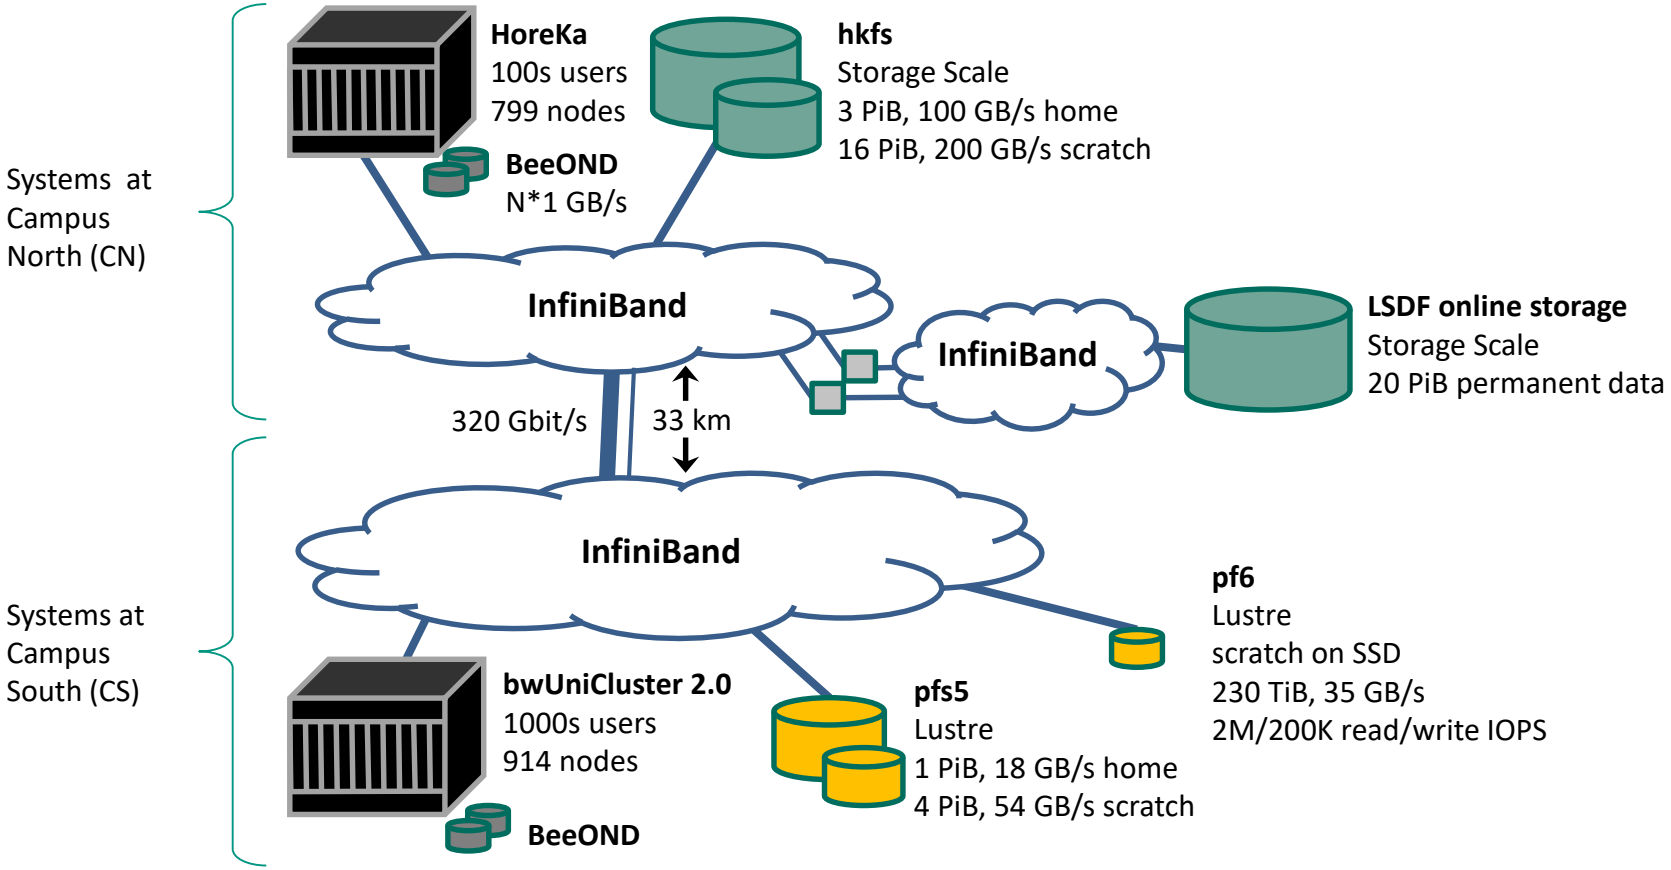

Lustre pros and cons

+ Runs very stable since years with extreme user base
+ Standard features very stable due to huge HPC user base
+ Lustre jobstats allow easy I/O performance monitoring
  + steadily used to find and educate power users
  + allow monitoring of different users on the same host
- Check if RHEL and MOFED version is supported
  - currently needs to be done before every upgrade
- Ability to use snapshots without performance loss
  - reduced metadata performance with ZFS
- Possibly missing features for enterprise usage
  - Windows client
  - commands to rebalance or fix replication

Storage Scale pros and cons

+ Most features stable due to huge HPC / industry user base
  + Snapshots, CIFS/NFS export and HSM available since many years
+ Supports metadata replication
  + Useful for online upgrades/extensions, additional hardware options
+ QoS on commands (e.g. rebalancing) very helpful
+ Many outstanding features
  + Windows client, AFM, multi cluster
- High license costs, frequently changing license policies
- Storage Scale client needs fixed amount of memory
- Normal configuration requires root ssh between all nodes
- Half-dead clients might hang up the complete file system
  - Happened in very rare cases

BeeGFS pros and cons

+ BeeOND very useful to provide on-demand FS for jobs
+ Administration is fairly easy
+ Relatively cheap support (for BeeOND)
- Commits after data is stored in server memory
  - Reason for some good performance rates
  - HA solutions require buddy mirroring, i.e. high hardware costs
- BeeOND normally requires root ssh between all nodes
- Relatively small feature list
  - Some features (e.g. quotas) only available with support contract