

Using Lustre jobstats for automatic user alerts about high I/O usage

Roland Laifer

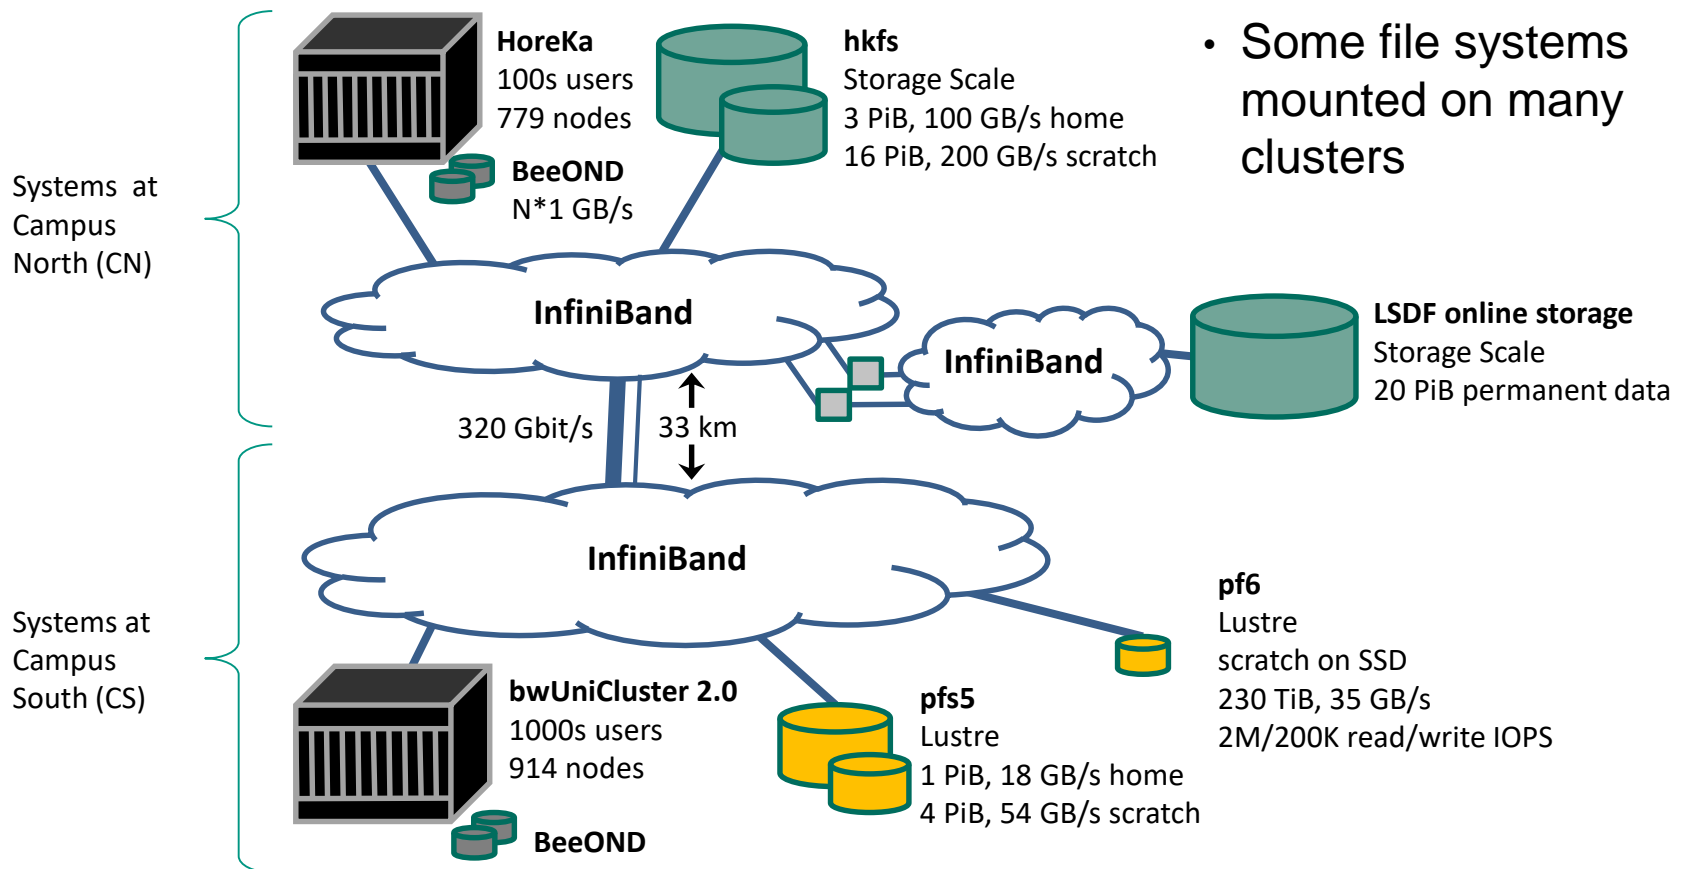
SCIENTIFIC COMPUTING CENTER - SCC



Overview

- Motivation
- How Lustre jobstats work
- Basic ideas
- Implementation
- Experiences
- Conclusion

HPC and parallel file systems at KIT



- Some file systems mounted on many clusters

Motivation

- Administrative experiences on bwUniCluster 2.0 (tier 3)
 - 1000s active users, ~100 logged in, ~1000 running jobs
 - A lot of inefficient or unacceptable Lustre usage
 - Most users are not even aware of their I/O
 - Lustre jobstats can provide statistics per user/job/host
 - Sending emails to power users is very time consuming
 - But this helps a lot for stabilization and speed up
- Goal: Send email alerts to power users automatically

How Lustre jobstats work

1. Define jobstats on MGS

- `lctl set_param -P jobid_var=SLURM_JOB_ID`
- `lctl set_param -P jobid_name=%j@%u@%H`
- `lctl set_param -P`
`*.*.job_cleanup_interval=3600`

2. Distribute jobstat settings

- Done by Lustre

3. Send content of jobid_name

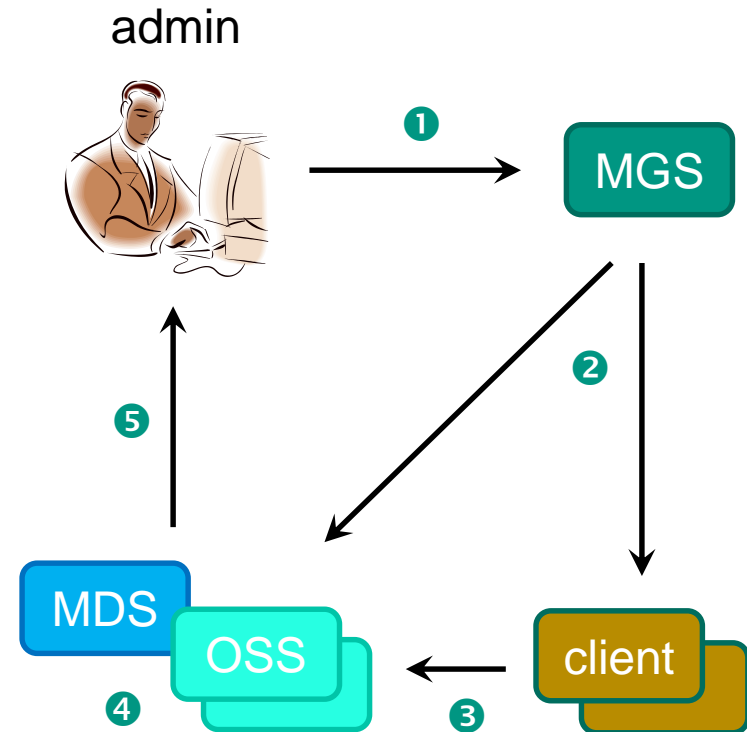
- Done by Lustre with I/O ops

4. Sum up I/O on targets

- For each jobid_name content
- Done by Lustre


5. Collect jobstats from servers

- Sum up I/O of targets
- Done by admin or monitoring application



Basic ideas

- Collect jobstats for all users on servers
 - Per file system sum them up for all MDTs and OSTs
 - Report metadata ops / total IO / IO ops above certain level
- Jobstats contain job ID and corresponding UID
 - No need to go to the batch system to get user for job ID
 - Get email address for UID from LDAP
 - Hence we can report jobs with high IO usage
- Send email alerts for each user only once per week
 - Too many emails will likely be ignored
- Create customized emails
 - Text blocks for high metadata, total read/write, IO operations
 - Same information in German and English



Implementation

- Configuration file with
 - watermarks for metadata ops, total IO and IO ops
 - file system and server names
 - file name for excluded users
 - path for logs and for email text blocks
- Perl script
 - Starts via cron every hour for each file system
 - Collects statistics from all storage targets
 - Creates a user email if any statistic is above watermark
 - Text block (metadata/total IO/IO ops) inserted if above watermark
 - Replace keywords in text block with collected statistics
 - Statistics also include job statistics if above watermark
 - Text blocks include recommended action, e.g. better use local SSDs

Experiences

- Much more users than expected are doing heavy I/O
- Pretty few users respond or start recommended actions
 - Another manual email usually helps
 - Anyway makes work easier since problem is already described
- Tool helped to find possible Lustre bug, see next slide
 - Many alerts concerning > 20 million reads on home file system
 - File system includes installed software packages
 - Some users were even doing more than 1 billion reads
 - Some users claimed to only use software on that file system

Possible Lustre bug

- Huge amount of read operations on home file system
- Discussed the issue for months with DDN
 - Client stats showed many page faults
 - Client cache was nearly empty while high read rates appeared
 - Shared libraries apparently were the main reason for page faults
 - Corresponding OSTs showed higher read statistics
 - Problem was hardly reproducible
 - For similar job problem could appear or not
 - Main game changer after we had reproducer on our test system
 - DDN developer was investigating the issue there
- Possible reason:
 - OSC page cache shrinker selects unevictable pages
 - Public part: <https://jira.whamcloud.com/browse/LU-17463>

Summary

- Lustre jobstats are exceptionally helpful
 - Just a few scripts needed to create alerting system about high I/O usage
 - Allow IO monitoring of different users on the same host
 - No comparable feature of other parallel file systems
- Check out if the mentioned Lustre bug exists on your systems
 - High read rates might be an indication
 - We've seen multiple GB/s created by the issue
- Credits for help with the development of our tool
 - Begatim Bytyqi, begatim.bytyqi@kit.edu
- My talks about Lustre
 - <http://www.scc.kit.edu/produkte/lustre.php>
 - roland.laifer@kit.edu