
Experiences & Performance of SFS/Lustre Cluster File System in Production

Roland Laifer

**Computing Centre (SSCK)
University of Karlsruhe**

Laifer@rz.uni-karlsruhe.de

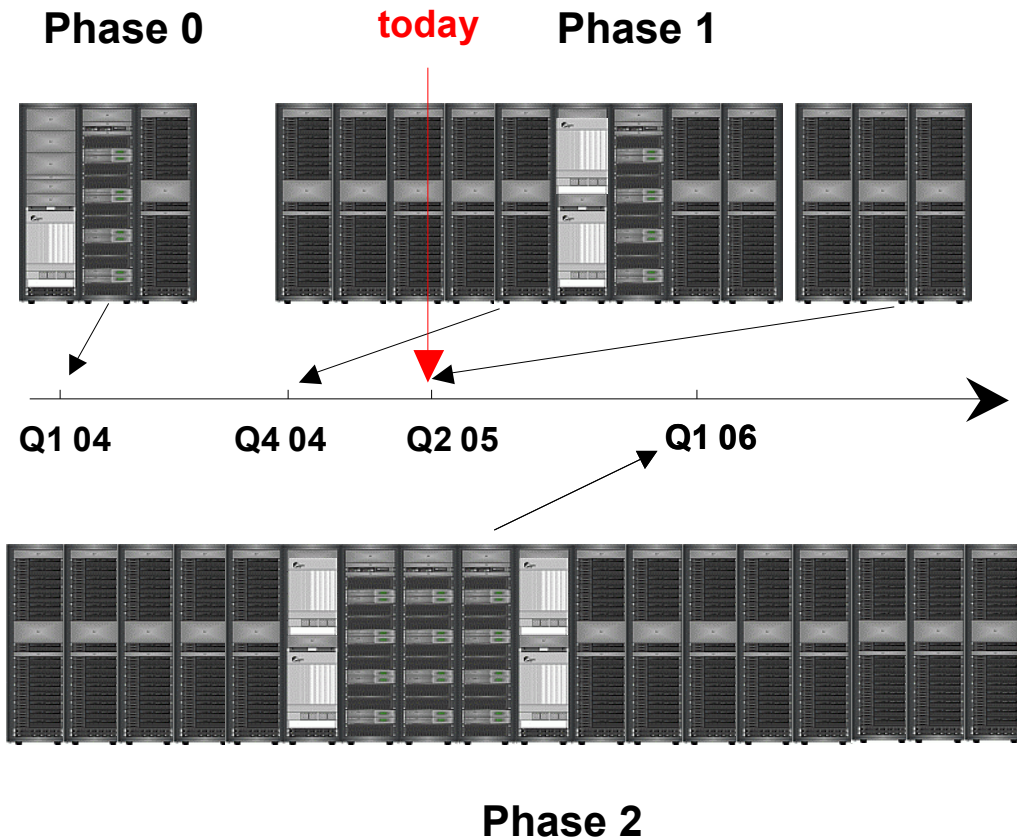


Outline

- » **Underlying hardware at SSCK**
- » **Performance measurements with HP SFS 1.1-0**
- » **Experiences with the HP SFS 1.1-0**
- » **Bugs**
- » **Wishlist**



HP XC 6000 Cluster installation schedule at SSK



Phase 0 (Q1 2004), Development

- » 16 two-way nodes
 - 12 Integrity rx2600
 - 4 ProLiant DL360 G3
 - Single rail QSNet II
- » 2 TB storage system

Phase 1 (Q4 2004), Production

- » 116 two-way nodes
 - 108 Integrity rx2620
 - 8 ProLiant DL360 G3
 - Single rail QSNet II
- » 11 TB storage system

Phase 1 (Q2 2005), Production

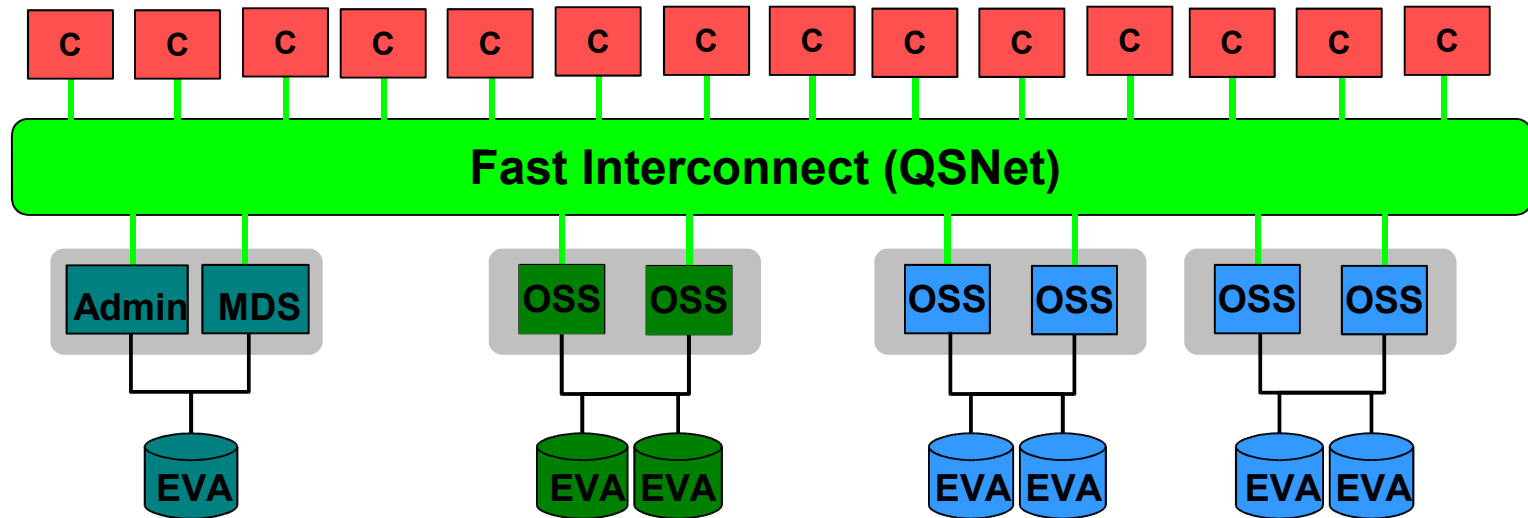
- » 12 8-way nodes
 - 6 Integrity rx8620, two partitions
 - Single rail QSNet II

Phase 2 (Q1 2006), Production

- » 218 four-way nodes
 - Two sockets
 - Dual core Montecito
 - Single or dual rail QSNet II
- » 30 TB storage system



HP SFS on SSCK's HP XC6000



**MDS and Admin for
\$HOME and \$WORK**

- allows > 50 million files

\$HOME

- 3.8 TB storage

\$WORK

- 7.6 TB storage

Legend

Admin: Administration Server

MDS: Metadata Server

OSS: Object Storage Server

EVA: EVA5000 storage array

C: Client



Performance measurement environment

» Used HP SFS software version was 1.1-0 (GA)

- Is based on CFS Lustre version 1.2.6

» Underlying HW

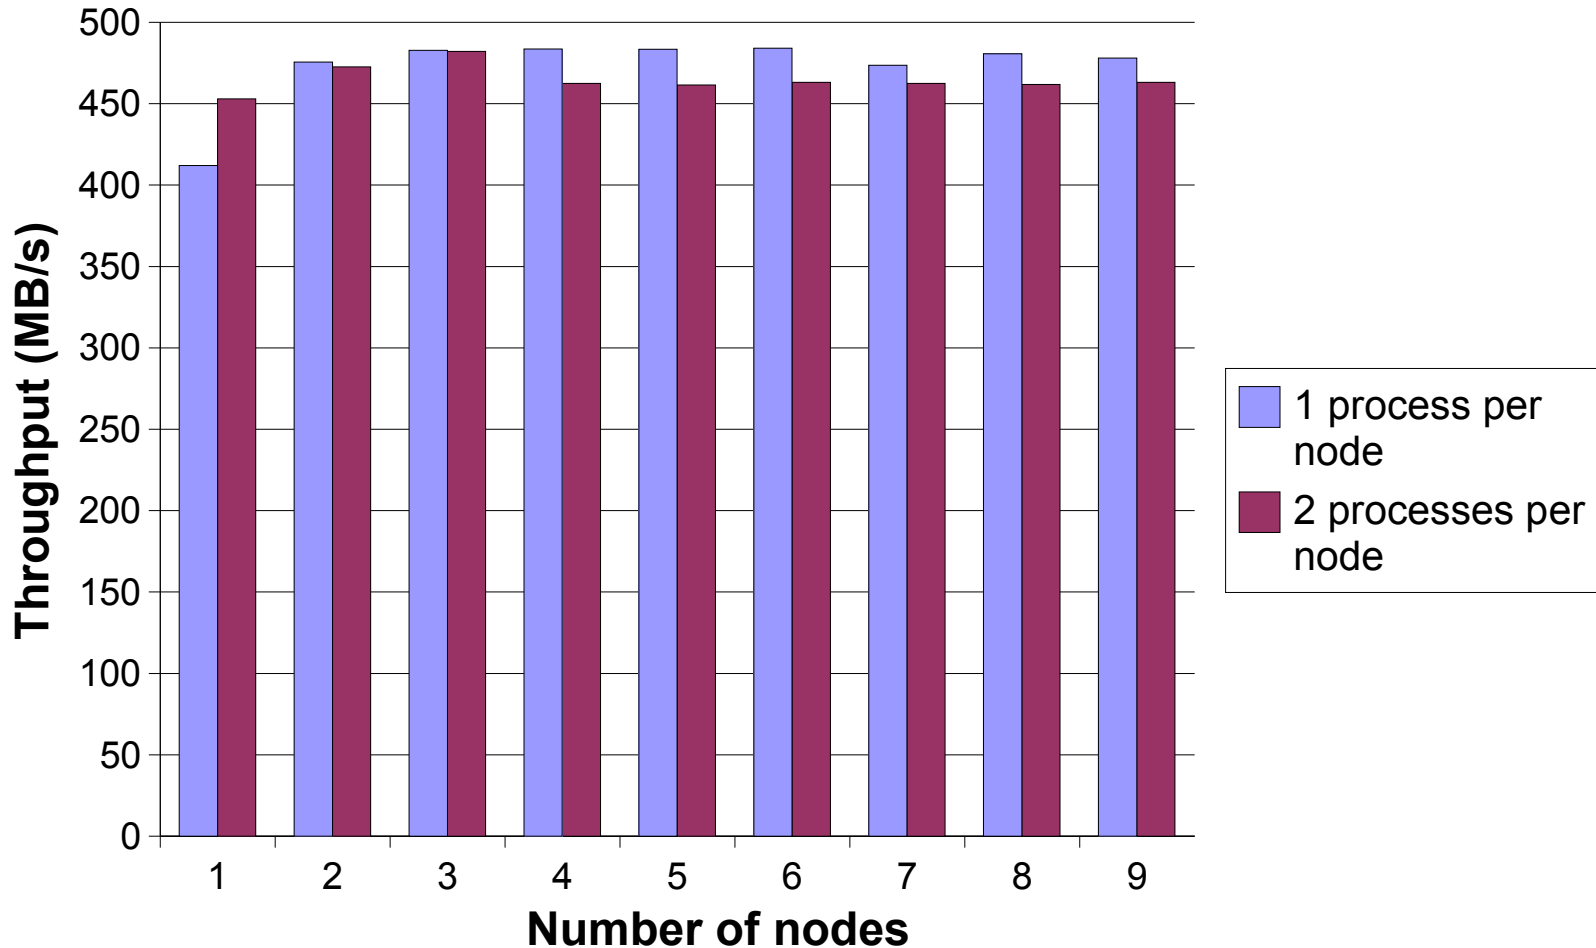
- Clients are IA64 systems (rx2600, 1.5 GHz, 2 CPUs, 6 GB memory)
- Quadrics QNet-2 (Elan4) interconnect
- EVA5000 (not EVA3000) storage systems with 2 controllers
 - OSS disks are 146 GB 10K, MDS disks are 72 GB 15K
- Servers are IA32 systems (DL360 G3, 3.2 GHz, 2 CPUs, 4/2 GB memory)
 - One file system (\$HOME) with 2 OSS and 128 KB stripe size
 - One file system (\$WORK) with 4 OSS and 1 MB stripe size

» Performance measurement details


- Measurements were done in parallel to production
 - \$HOME file system was used which also reduced the MDS performance
- Benchmarking software was bonnie++




Sequential block write performance with 4 OSS




Block vs character write performance with 2 OSS




Sequential block read performance with 4 OSS




Block vs character read performance with 2 OSS



File creation performance with 4 OSS



File creation vs file deletion performance with 2 OSS



Performance measurement results

- » **RAW lun performance using 2 controllers (1 EVA) in parallel:**
 - `raw_lun_check.bash` showed 107 MB/s for writes and 192 MB/s for reads
- » **EVA seems to be bottleneck for writes**
 - Bottleneck seems to be the mirrored cache or the RAID controller
- » **FC adapter seems to be bottleneck for reads**
 - `portPerfShow` on FC switches shows that only one path is used !
- » **Main benchmarking results**
 - Write performance is about 115 MB/s per OSS
 - Read performance can reach 190 MB/s per OSS
 - File creation performance can reach 5000 creates/s
 - Character-wise operations decrease throughput on clients only
 - CPU on clients is 100% used
 - Overall throughput on servers is same as for block operations



Experiences with HP SFS 1.1-0

- » **Works pretty stable when everything is up and running**
 - **Production server system was running > 4 weeks without any problem**
 - **more problems seem to arise when the system is up for a long time**
- » **Understanding Lustre error messages is important**
 - **in order to separate critical from normal error messages**
- » **After dumps check local disk space**
 - **Filesystem /local on OSS is hidden and not visible by df**
- » **Be patient:**
 - **sfsmgr commands may succeed when you think they are hanging**
 - **Otherwise they usually run into timeouts**
 - **Status of services may be reported with a delay**



Experiences with the HP SFS 1.1-0 (cont.)

- » **While one problem is repaired, often additional problems arise**
 - e.g. a server dumps while it is rebooted
 - plan enough time for maintenance windows

- » **Same sfsmgr commands may supply different results:**
 1. **Timing has an influence**
 - e.g. takeover only happens if replicating server is up for more than 10 minutes
 2. **Status of clients can have influence on servers**
 - e.g. filesystem start is faster if all clients can be reached
 3. **Bugs in software components**

- » **Filesystem operations continue after a problem is repaired**
 - Usually batch jobs continue to run
 - However they may run into the batch system's job time limit



Bugs in HP SFS 1.1-0

- » **MDS stops working and node needs to be manually rebooted**
 - Occured 4 times in 18 weeks
 - Dumps are still under investigation

- » **rx8620 nodes with 8 CPUs / 64 GB memory crash while copying files**
 - Is reason for delay of rx8620 integration into production system
 - Problem with the virtual memory cache, fix is available

- » **Filesystem hangs because OSS and client lost Quadrics connection**
 - Happened twice while OST services were unbalanced
 - Dumps are still under investigation

- » **Open vi on file while deleting subdirectory crashes client node**
 - Solved with a new kernel

- » **Invalid inode during MDS recovery causes dump**
 - Usually happens after MDS takeover; ASSERTION() failed messages appear
 - partly solved in HP SFS 1.1-1



Wishlist for enhancements

- » **Improve problem reporting system at www.itrc.hp.com**
 - Allow direct communication with higher level support
 - Most actions and the current status should be visible
 - Well formatted ASCII text output is required

- » **Improve HP SFS integration into XC**
 - Supply HP SFS start scripts and start it before other services
 - This would allow other services like SLURM to use Lustre

- » **HP should consider additional hardware support in HP SFS**
 - especially for storage subsystem and servers
 - CEA has demonstrated more than 2 GB/s Lustre throughput per OSS



Conclusion

- » **We still expect a hard time to reach a highly reliable system**
 - **Parallel file systems are very complex**
 - Hence it is normal to have critical software bugs with new file systems
 - **We further need excellent support**

- » **HP SFS has the most important features of a parallel file system**
 - **Performance, resilience, scalability, and ease of administration**
 - **Additional features are needed for using file systems from two clusters**
 - e.g. support for different HP SFS/Lustre versions between clients and servers

- » **HP SFS is a very interesting and promising product**
 - **It works and is heavily used at SSCK's production system**
 - Now it's the right time to start using it !

