
Latest Experiences with HP SFS

Roland Laifer

**Scientific Supercomputing Centre (SSCK)
University of Karlsruhe
Germany**

Laifer@rz.uni-karlsruhe.de



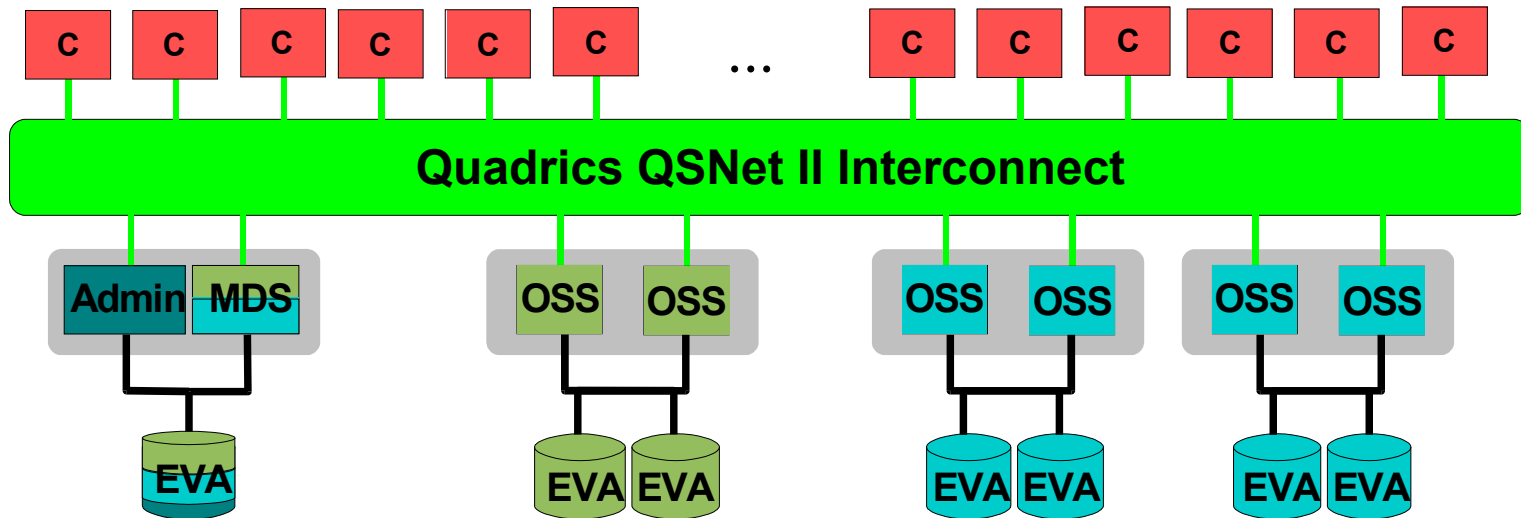
Outline

- » **SSCK's SFS production systems**
- » **Performance measurements**
- » **Performance monitoring graphs**
- » **New operational experiences**
- » **Some open problems**



Itanium production system (xc1)

120 clients (Itanium)



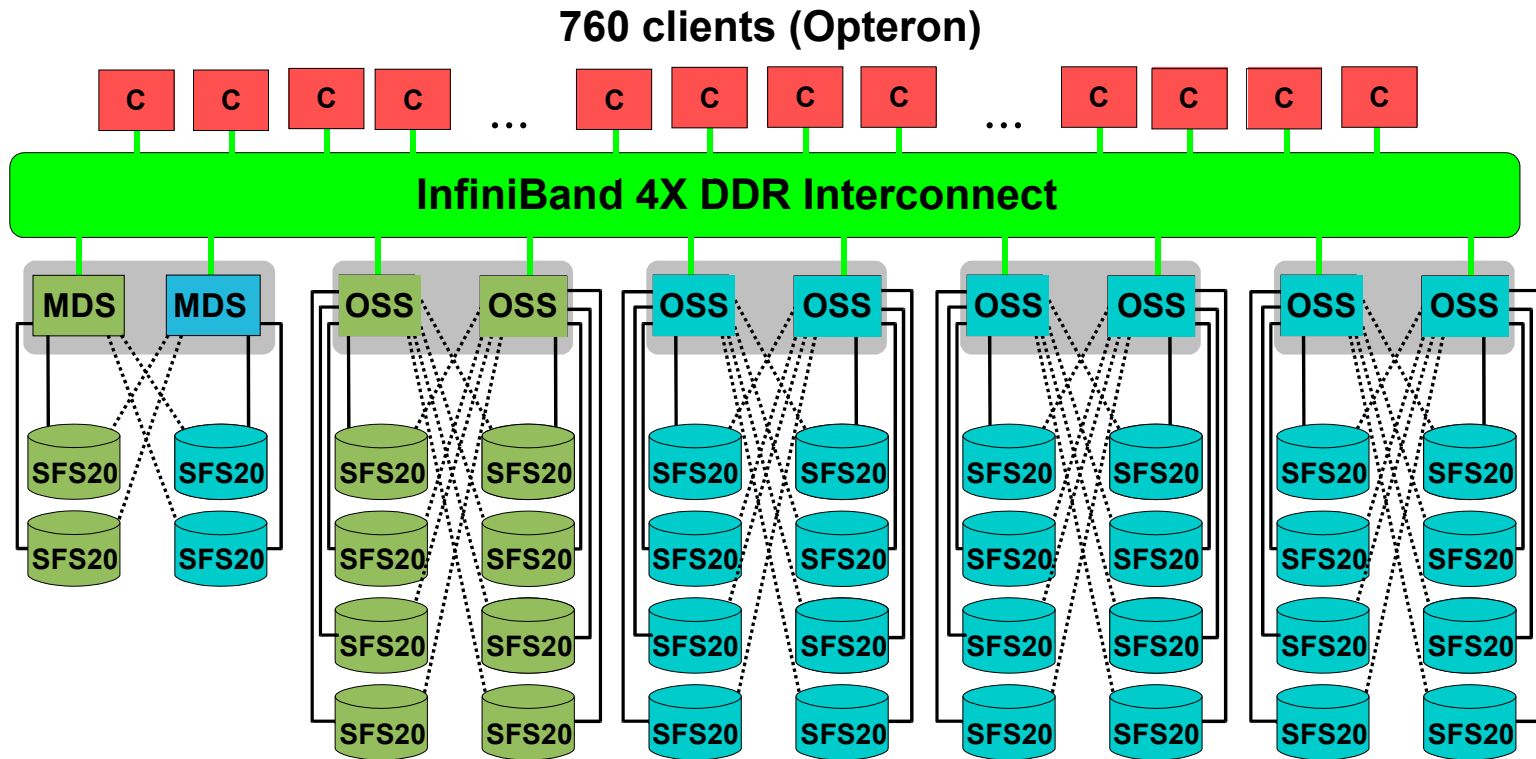
	\$HOME	\$WORK
Capacity	3.8 TB	7.6 TB
Write performance	240 MB/s	480 MB/s
Read performance	380 MB/s	760 MB/s

Notes:

- Fragmentation issues solved by recreating both file systems
- Saved and restored subtrees on many clients with tar files stored on the other file system



Opteron production system (xc2)



	\$HOME	\$WORK
Capacity	8 TB	48 TB
Write performance	360 MB/s	2100 MB/s
Read performance	800 MB/s	3000 MB/s


Notes:

- \$HOME file system uses mirrored OST luns
- Lower performance if other jobs create congestion on InfiniBand




Performance monitoring graphs for one OSS

» On xc1:




» On xc2:




Performance monitoring graphs for one MDS

» On MDS for \$HOME on xc2:




» On MDS for \$WORK on xc2:



Performance measurement with parallel dd on xc1

Write and read performance of file system data on xc1




Note:

- Omit caching effects by reading a file which was created on another node




Performance measurement with parallel dd on xc2

Write performance of file system work on xc2



Notes:

- Each process used its own file with stripe count 4
- Peaks when all 24 OST luns were used



New operational experiences (1)

» Disk failure rates

- FC failure rates decreased
 - During first year possibly many broken disks due to infant mortality
- SATA failure rates are normal
 - Lost 4 of 432 disks in 6 months

» Congestion on InfiniBand might have impact on I/O performance

- Congestion on switches possibly caused by other jobs

» Check return code of `ptlrpc_check_status()` Lustre error messages

- See `/usr/include/asm/errno.h` on Itanium
- See `/usr/include/asm-x86_64/errno.h` on Opteron
 - Return code -122 means quotas are exceeded

» Lustre works pretty good on an unstable network

- During initial system test InfiniBand on xc2 was pretty unstable



New operational experiences (2)

» Investigation of Lustre timeouts

- Evlview shows:
 - Apr 30 09:42:42 xc1-ls5-adm ... LustreError: ... lock callback timer expired: evicting client ... nid 367@elan ...
- /var/log/messages on corresponding client shows:
 - Apr 30 09:43:11 xc1n112 ... LustreError: This client was evicted by xc1-ls-ost29 ...
- Check how many seconds timeout has been missed:
 - Mar 12 18:36:39 ... evicting ... nid 10.22.1.130 ... lock: 43efbc80/0x5e38434477ce8fcc ...
 - Mar 12 18:36:51 ... received cancel for unknown lock cookie 0x5e38434477ce8fcc ...
 - The 2 lines above show that the timeout was missed by 12 seconds.

» Possible reason for Lustre timeouts

- Each task of a parallel application uses small regions in one large file
 - e.g. 512 tasks do multiple 1 KB writes to different locations on one large file



Some open problems

- » **Performance degradation on mirrored OST luns with SFSv2.2-0**
 - Write performance drops to 50 % during concurrent writes
 - Problem is under investigation

- » **Performance degradation on xc2**
 - After 6 months maximum throughput shows 40 % lower peaks
 - Problem is under investigation

- » **Used quotas not decreased after deleting files**
 - Happens sometimes after quota limit is reached
 - Problem is under investigation



Summary

- » Lustre provides a scalable and stable parallel file system
- » HP SFS supplies additional features which make it a real product
- » Investigation of performance problems is not easy
- » Further experiences with HP SFS:
 - <http://www.rz.uni-karlsruhe.de/dienste/lustretalks>

