



science + computing

| A Bull Group Company



Lustre-based Storage Solutions in the Automotive Industry

Daniel Kobras

science + computing ag

IT-Dienstleistungen und Software für anspruchsvolle Rechnernetze

Tübingen | München | Berlin | Düsseldorf


Who are we?

Founded in
Offices

1989
Tuebingen
Munich
Berlin
Duesseldorf

Employees
Shareholder
Turnover 07/08

270
Bull S.A. (100%)
26 Mio. Euro



Portfolio

IT Service for Complex Computing Environments
Complete solutions for Linux- and Windows-based HPC

scVENUS System management software for efficient administration of homogeneous and heterogeneous environments

Where are we?

- Automotive Industry
- Plant/Mechanical Engineering
- Aerospace Industry
- Microelectronics
- Chemical/Pharmaceutical Industry
- Biotechnology
- Public Sector
- IT Service Provider




What are we doing (with Lustre)?

- Operating HPC installations in the automotive industry (aerodynamics, aero accoustic, crash simulation, ...)
- Introduced Lustre as a scalable storage solution for HPC sites at several environments
 - Operating Lustre sites since 2007 (early 1.6)
 - Typically small-scale storage (~100 TB, 2-4 OSS)
- Most installations used as cluster scratch space
- Installations of group-level central storage (scratch and long-term), providing faster job turnaround times and simplified workflows

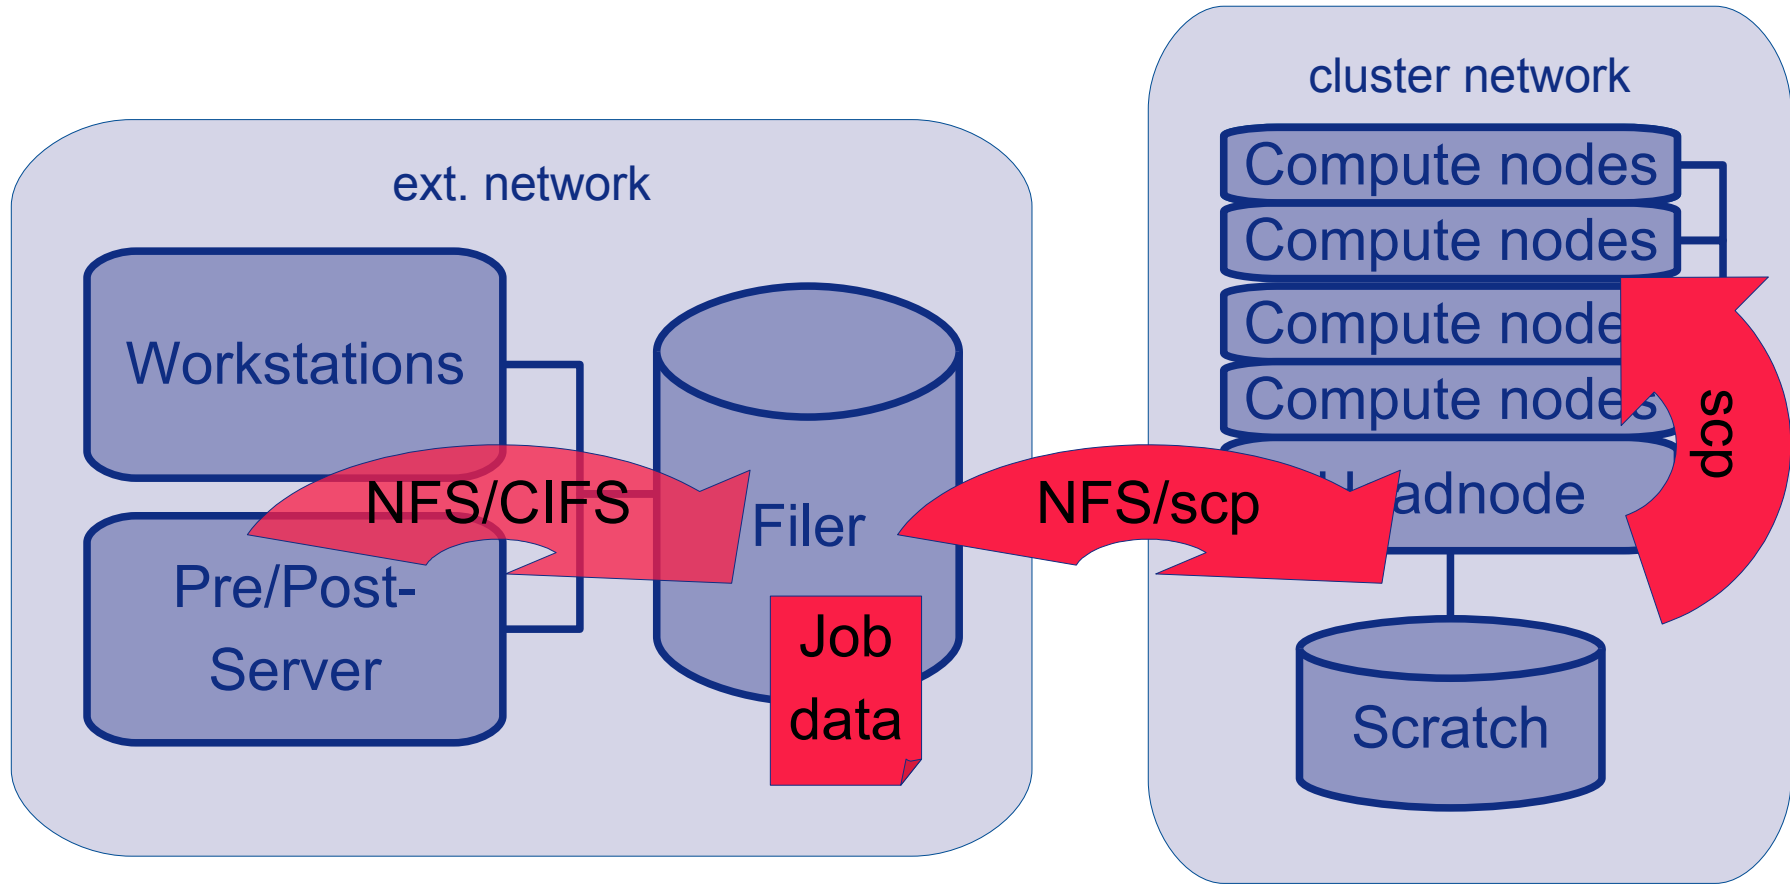
Why did we choose Lustre?

- Obvious reasons:
 - High I/O throughput for massively parallel compute jobs
 - High efficiency by leveraging existing RDMA interconnects
 - Well-established in HPC community
 - Free, no vendor lock-in
- Allows easy filesystem growth beyond partition boundaries
 - Namespace layout no longer limited by hardware constraints
 - TB-sized files don't cause administrative burden
- Allows centralised storage without hampering performance
 - Simplified workflow
 - No duplication of data
 - No more data loss from compute node failures

Traditional engineering workflow




Traditional engineering workflow

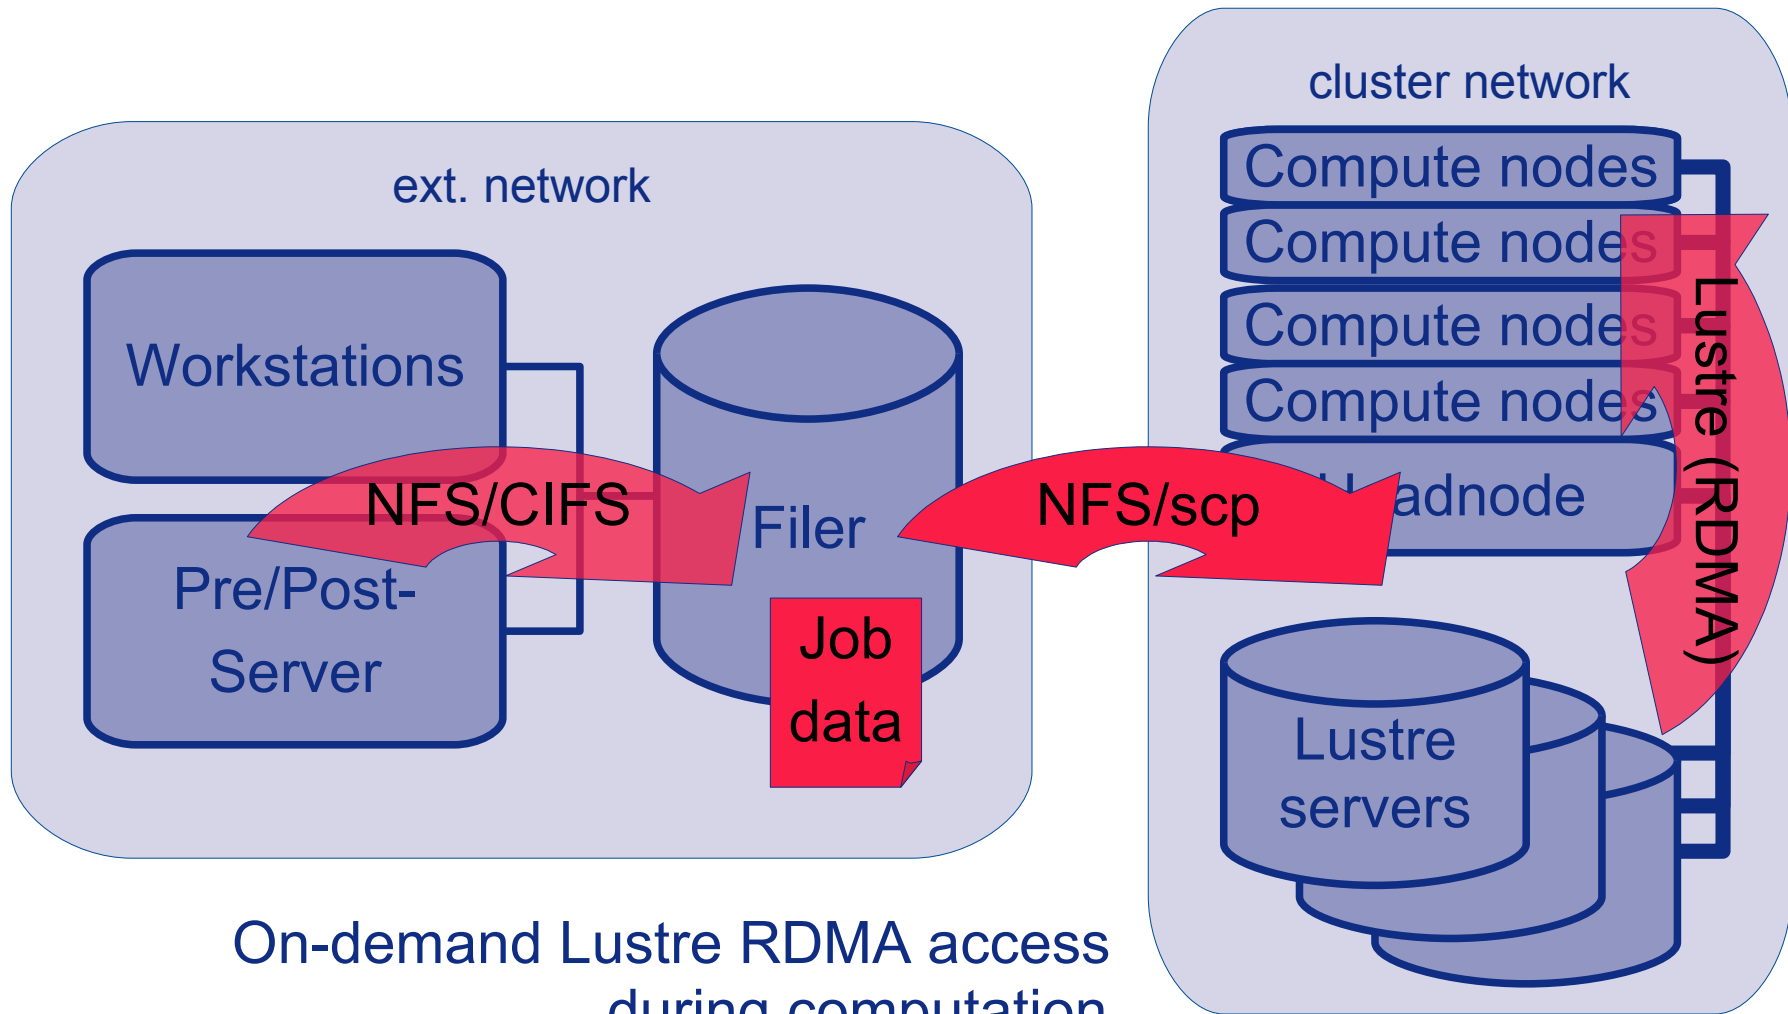


Data movement **serializes** workflow

Lustre as cluster scratch

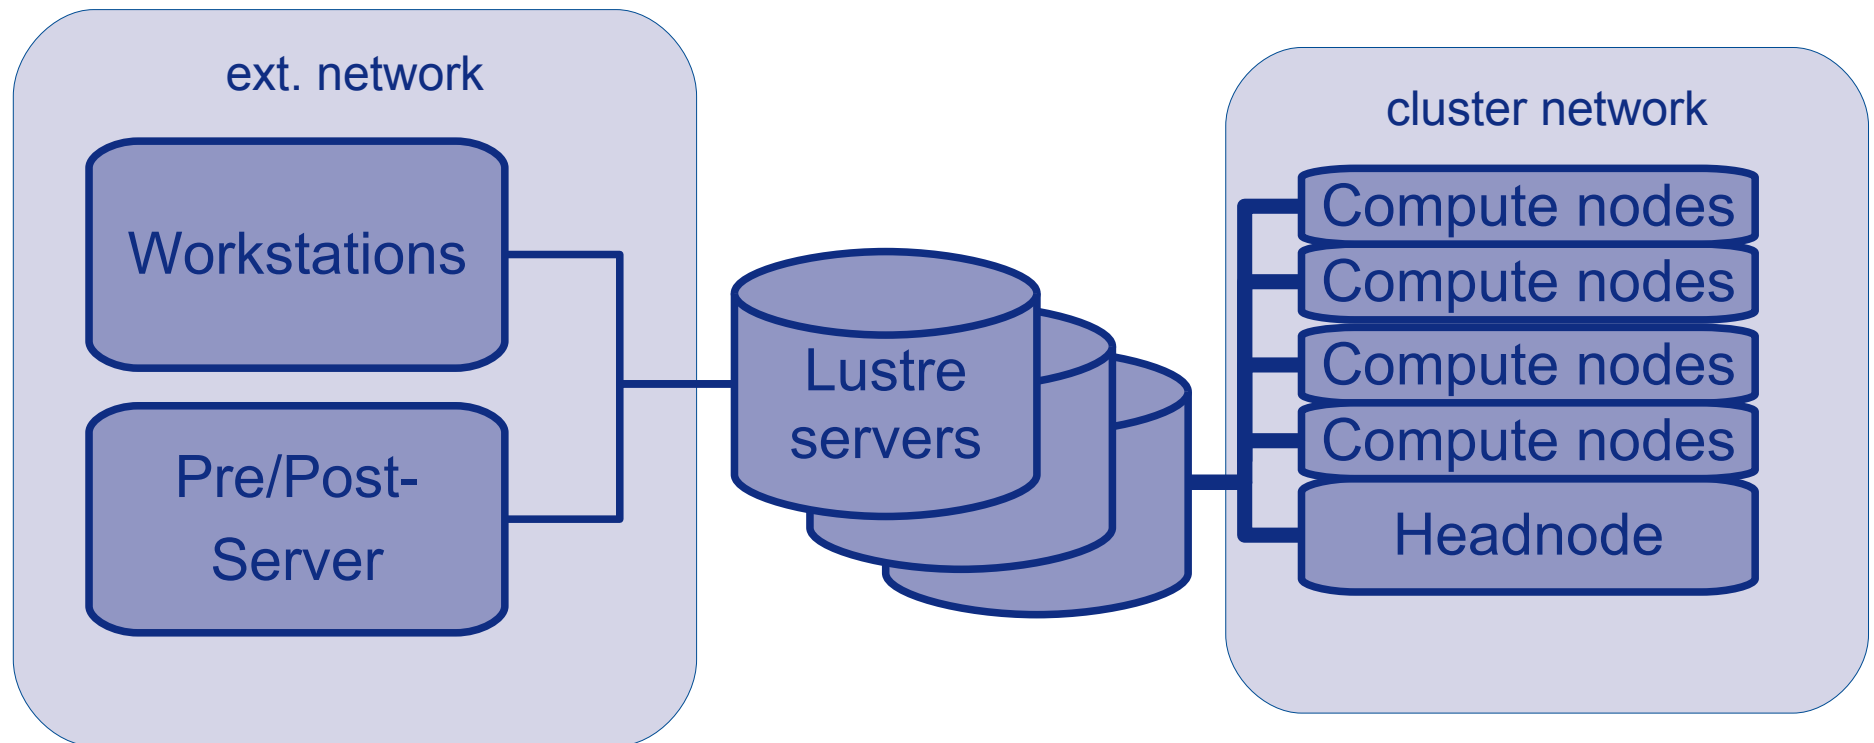


Lustre as cluster scratch




On-demand Lustre RDMA access during computation

Lustre as central storage



Lustre as central storage

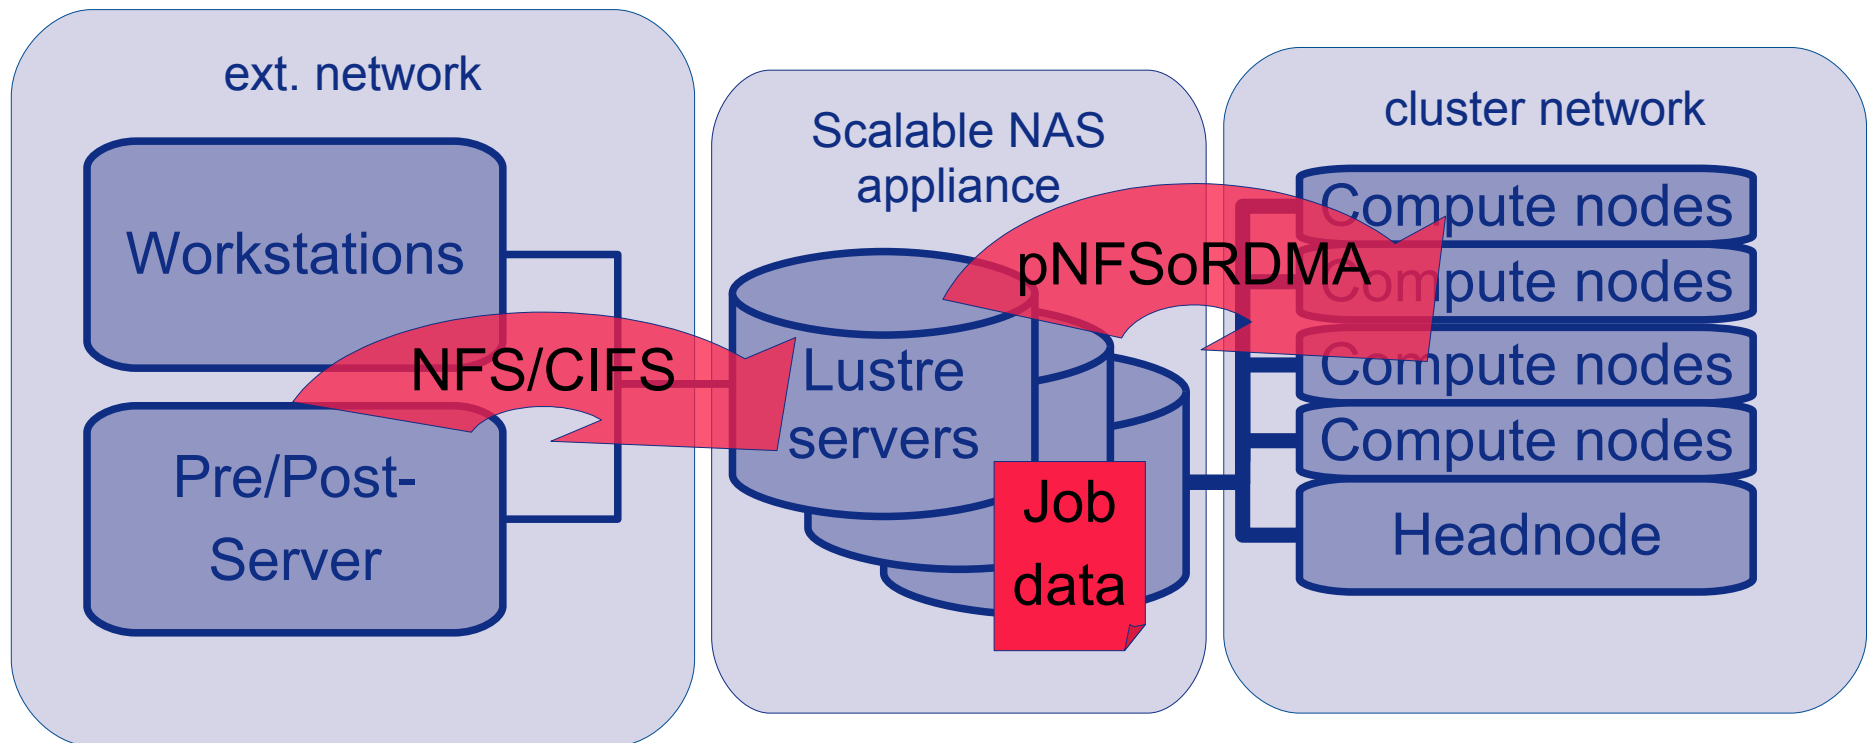


Speed-up due to elimination of extraneous data movement

Workflow enhancements

- When using Lustre as a cluster scratch, time is still wasted copying data
- Lustre's inherent scalability makes it possible to use it as a central storage location without compromising on performance
- Benefits:
 - Eliminates duplication of data
 - Eliminates copy processes between different types of storage
 - Facilitates data management (administrators and end users)
- Downsides:
 - Lustre system as a whole becomes single point of failure
 - Ubiquitous roll-out of Lustre client software necessary (unless using NFS/CIFS re-exports)

NAS appliance, powered by Lustre (tbd)



Standard (parallel) access protocols dispense with need for client-side software installations. No vendor lock-in.

Case study

- Car manufacturer
- Aerodynamics/aeroacoustics division
- Simulation of fluid dynamics using commercial CFD codes (StarCD, StarCCM+)
- Massively parallel jobs, scaling up to 1000 cores; typically 200-500 cores/job
- Bound by CPU and memory bandwidth
- Input and output data typically from 100 GB to 1 TB
- I/O: regular checkpointing from dedicated nodes
- High increase of data volume from transient jobs
- Interactive and automated pre/post-processing
- Small user base (~10 users)

- Several Linux clusters (64 - ~1000 cores), mostly Sun Fire X4100M2
- Interconnects: Myrinet, InfiniBand (SDR, DDR)
- Queueing: Sun GridEngine 6.2
- Several SMP nodes for pre/post-processing (Sun Fire X4600)
- Visualisation, modelling, testing on local workstations

Challenges

- High data volume in general (output data ~100 GB–1TB/file)
- Data volume of individual projects exceeds typical partition sizes
- Moving job data to/from temporary storage on cluster nodes contributes significantly to overall job runtime
- Inevitable data loss when individual nodes go down
- Idle nodes during copy processes
- **Central** storage accessed from cluster nodes and interactively from workstations desirable

Components

- Lustre 1.6 (currently upgrading to 1.8.1)
- MDS
 - 2 Sun Fire X4100M2, SE3320, 6x140 GB, Raid10
 - CentOS 5
- OSS
 - Initially: 2 Sun Fire X4500, 48x 500 GB SATA each
 - Expansion: 1 Sun Fire X4540, 48x 1 TB SATA
 - 4 OSTs per OSS, Raid 6, CentOS 5
 - **No OSS failover possible**
 - ~66 TB total capacity, ~80 per cent used
- InfiniBand 4x SDR links to cluster nodes and main servers
- Ethernet links to workstations

- In production since October 2007, ~260 clients
- Central archive for all project data
- 1-2h faster turnaround for typical jobs, mostly due to lack of copy overhead
- Lack of redundancy in data servers does not pose significant problems
 - Redundancy of hardware components sufficient for most problems
 - Short reboots can be scheduled outside busy hours
- Backup strategy
 - TSM incremental backups of all data
 - Backup different subtrees via 3 (initially 1) clients
 - Still problems with large (~ 1 TB) or many small files

Summary

- Lustre can provide more than just scratch space
- Benefits from using central storage on top of performance improvements
- Proves to be reliable enough for interactive use even without full high-availability setup
- But: It still can give an admin headaches...

Day-to-day fun with Lustre

Administrative experiences

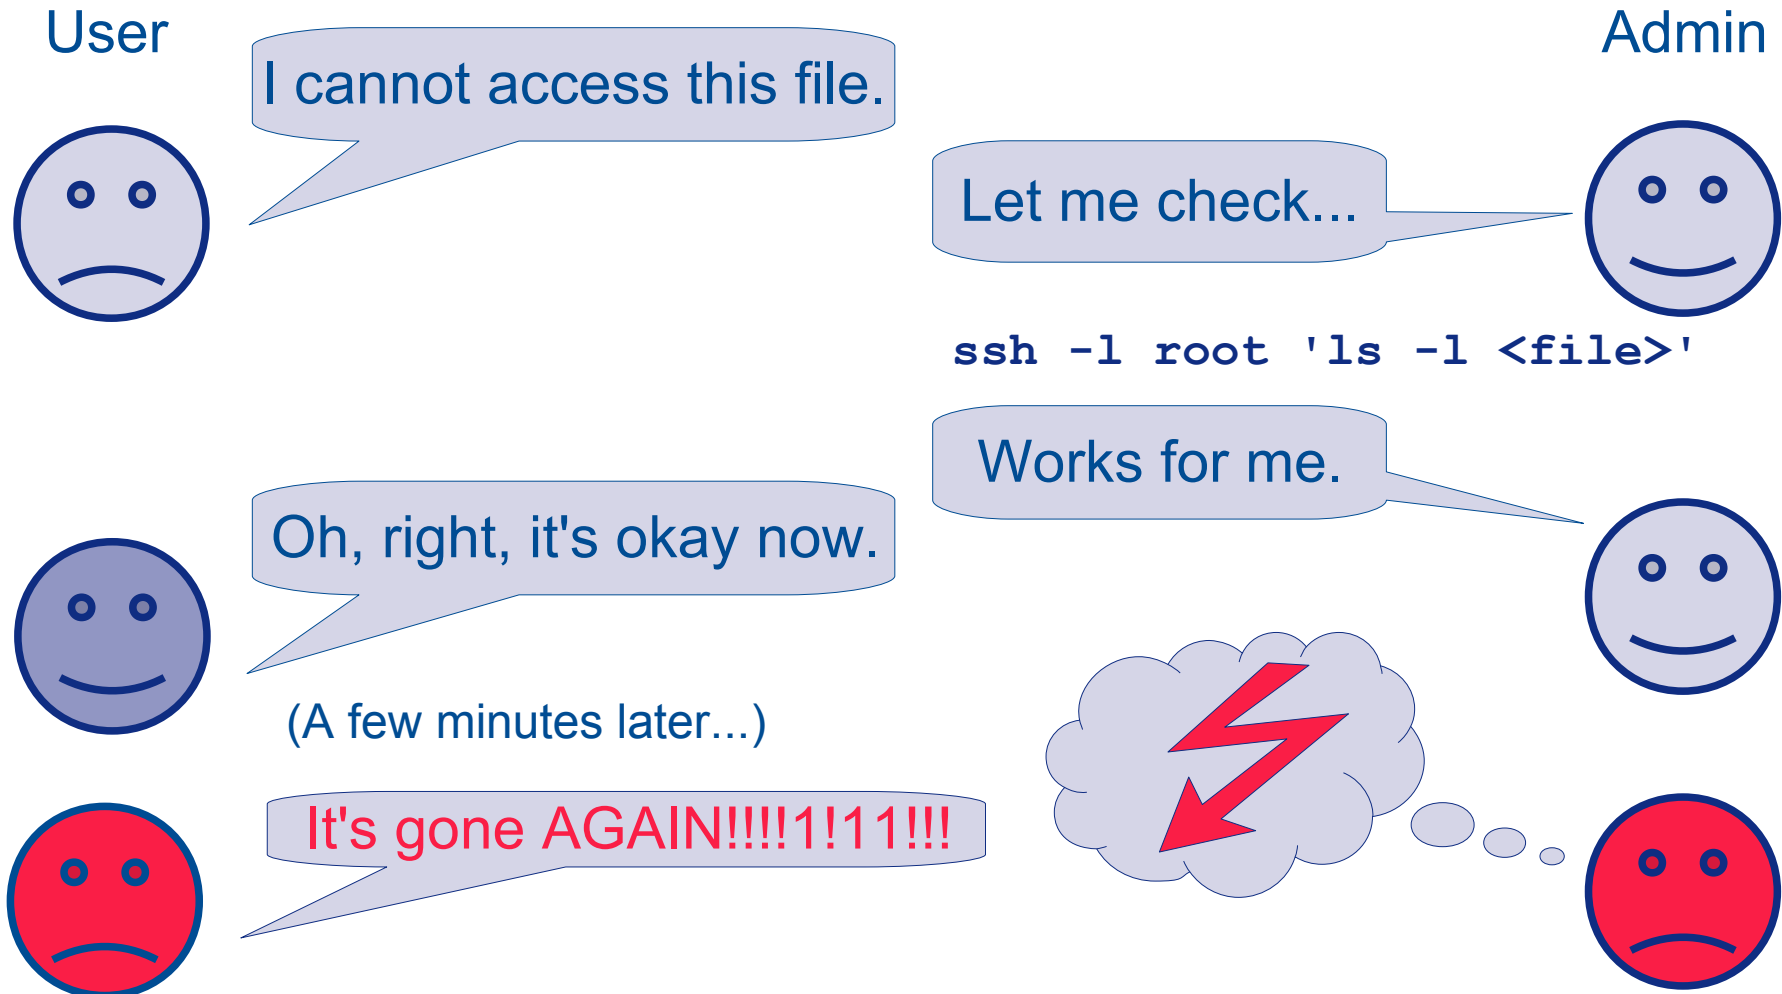
- Lustre error messages rarely give an immediate hint to the root cause of the problem
- Massive stability problems in early production that didn't occur in testing phase, finally tracked down to instability of IB subnet manager
- Configuration changes require filesystem downtimes (eg. network reconfiguration, configuration roll-back)
- Extension process is fragile because new OSSes immediate become operational (in 1.6)
- MGS interaction is opaque in general
- Lustre can break assumptions in other software where you'd least expect them

Ex 1: Server-side authorisation – or not

- Group membership of a user accessing a file in Lustre is evaluated on the MDS
- Can be configured via upcall (group upcall nowadays configured by default)
- MDS authorises user access based on this information
- **EXCEPT:** client authorises user access for cached metadata, group membership is evaluated locally

Ex 1: Server-side authorisation – or not

- Consider a defect in group resolution on MDS:



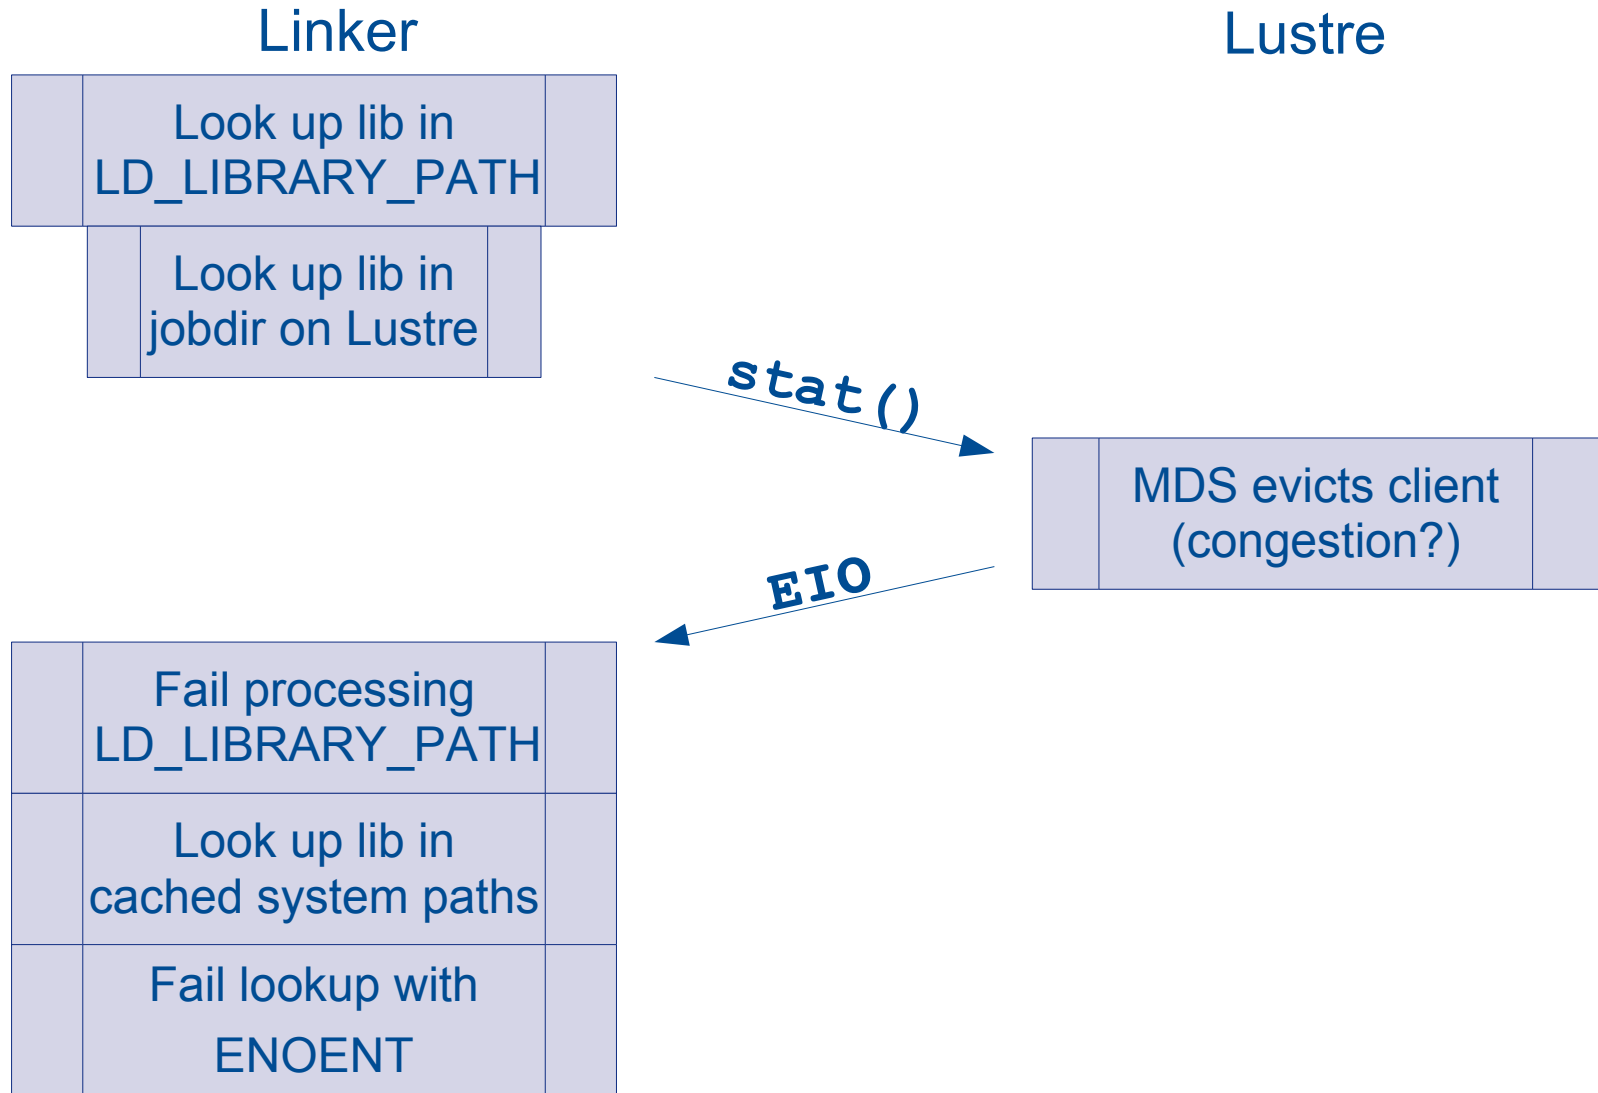
Ex 2: With Lustre, EIO can be transient

- Parallel application fails to start because arbitrary nodes fail to find a private shared lib on the local filesystem:
`/usr/local/myapp_v1.2.3/bin/myapp: error while loading shared libraries: libapp.so: cannot open shared object file: No such file or directory`
- Cannot possibly be a Lustre problem?

Ex 2: With Lustre, EIO can be transient

- Parallel application fails to start because arbitrary nodes fail to find a private shared lib on the local filesystem:
`/usr/local/myapp_v1.2.3/bin/myapp: error while loading shared libraries: libapp.so: cannot open shared object file: No such file or directory`
- Cannot possibly be a Lustre problem? – It is!
- `libapp.so` resides in `/usr/local/myapp_v1.2.3/lib`
- Application adjusts `LD_LIBRARY_PATH` to contain
`/lustre/my/jobdir/plugins:/usr/local/myapp_v1.2.3/lib`
- Syslog reveals that client was evicted from MDS, but immediately reconnected (so no long-term malfunctions were visible)

Ex 2: Flow of events





science + computing

| A Bull Group Company



Thank you for your attention!

Daniel Kobras

science + computing ag

www.science-computing.de

Telefon 07071 9457-0

info@science-computing.de