


Parallel File System Concept developed by NEC Fast and reliable data storage in Computer-aided Automotive Engineering



→ Challenges

Whether the need is to reduce weight, or to develop quieter and especially more fuel efficient and low-emission vehicles - in the automotive industry, numerical simulation methods for product development and optimization are essential. High-performance clusters of computers based on standard components are capable of providing the computational resources for running software tools such as Nastran, Abaqus and LS-Dyna. However, with increasing model complexity and spatial resolution, the amount of data is rapidly increasing and very often the limited I/O bandwidth of the filesystem proves to be the determining factor of overall performance for the computer system. A mature concept has been developed by NEC for a global parallel file system, which eliminates this bottleneck and significantly increases the availability.

→ Solution


All Linux HPC Clusters (e.g NEC LX series 2000-4000) can now be equipped with NEC LXFS, this being a high performance parallel file system based on Lustre technology. NEC has developed its own solution which complements the performance of Lustre by adding outstanding reliability. By utilising a building-block approach, it is very easy to handle and to scale. Moreover, a procedure was developed to considerably simplify installation and commissioning. This solution was jointly developed with a customer in the automotive industry and is now being used in the daily production process.

The complete **NEC LXFS** solution includes both the software and associated hardware components, such as Metadata Server and Object Storage Server which provide the disk resources and is fully integration into **NEC OSCAR-Pro**.

→ Concept and Realization

NEC LXFS consists of one Metadata Block and at least one Object Store Block. The servers in both blocks are bound together in a failover configuration and are able to cope with failures of hardware components. All components have multi-layer redundancies at the software and the hardware level. The Metadata Block is responsible for metadata services and it also takes the management role of the file server cluster.

The Object Store Block is the basic element of the scalability provided by NEC's approach. The integrated file servers in active-active failover configuration manage multiple Object Storage Targets. The bandwidth, as well as capacity, of the parallel filesystem scales with the number of Object Store Blocks.






NEC LXFS

- ✓ Modular
- ✓ Scalable
- ✓ Highly available

PROMISE[®]
TECHNOLOGY, INC.



Modular and scalable Design

→ Delivered Solution

An NEC customer deploys hundreds of compute nodes based on AMD and Intel quad-core processors connected into a InfiniBand fabric. One metadata block and three Object Store Blocks with 36 TB capacity each were attached to the existing Infiniband Fabric. At the core of the Object Store blocks are RAID systems delivered by Promise Technology Inc. The VTrak Series E310s can deploy SAS or SATA drives, being wireless and fully redundant. All system components are constantly monitored to detect possible malfunctions and to automatically start countermeasures. Each RAID system houses two RAID controllers in an active-active configuration, which further increases the availability. NEC selected Promise RAID systems in a RAID6 configuration because of their proven success for NEC LXFS. NEC supplied the parallel file system as a turnkey solution, including user and administrator training as part of the overall package.

→ Solution Highlights

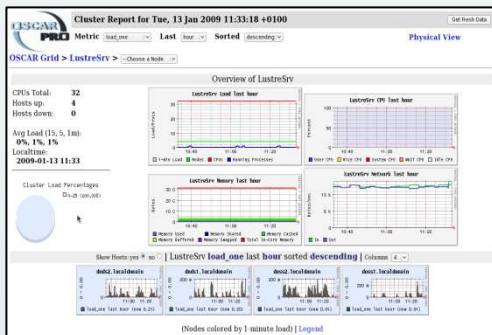
NEC provided analysis, consulting, optimization and training for the efficient use of NEC LXFS and also wider, ongoing support after the installation phase.

Scalability and modularity

NEC LXFS offers the greatest flexibility in terms of bandwidth and capacity. The Object Store Blocks may vary according to customer requirements. Just as the storage capacity of the parallel file system scales with the number of Object Store Targets, the I/O bandwidth also scales with the number of Object Store Targets.

High Availability

NEC has always focused on providing high system availability and reliability, hence the overall design of NEC LXFS ensures continuous operation. Multiple layers of redundancy both at the hardware and software level are implemented, such as RAID6, multiple controllers and failover servers. NEC also carefully selected all system components. The monitoring of the system is transparent, flexible and adjustable and an automatic notification system informs about critical events.

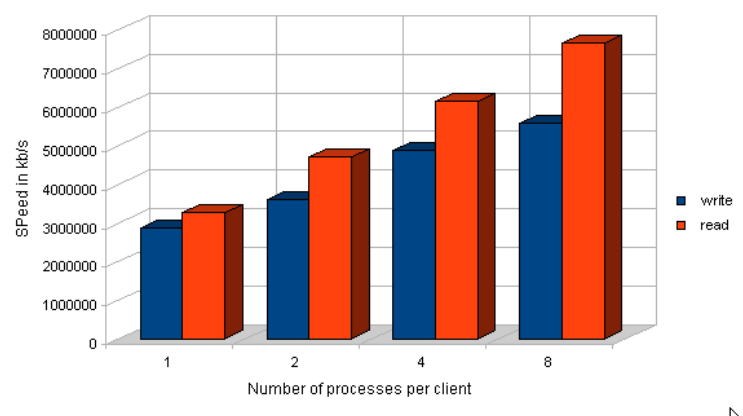


NEC Oscar-Pro GUI Monitoring Tool

→ Results

Designed using Object Store Blocks with a high throughput, in this configuration we achieved an outstanding I/O bandwidth measured with the IOZONE benchmark. We used a filesize of 16 GB in order to eliminate caching effects. With the installed NEC LXFS configuration the bandwidth for writing is more than 5 GB/s and more than 7 GB/s when reading.

An even higher overall bandwidth and a higher capacity can be realized simply by adding extra Object Store Blocks



→ Contact

Dr. Andreas Findling
info@hpce.nec.com