

# LISA Compute Cluster

## Why LISA Compute Cluster?

For access to a high-performance compute cluster that combines processing power with user friendliness. LISA Compute Cluster is a solution designed for calculation intense applications and for those seeking a centrally managed Linux cluster that is easy to work with. You simply upscale your computations to a higher level, or use it to handle jobs requiring more advanced features like accelerators and high memory capabilities.

## What is LISA Compute Cluster?



LISA Compute Cluster consists of several hundreds of multi-core nodes, running Linux. It is ideal for large-scale computations and suitable for both high-throughput computing as well as moderate parallel applications. User jobs are allocated to compute resources on node-level and automatically executed in a series of computer jobs (batch). Users are given a login, which enables immediate, 24/7 access to the service.

The LISA Compute Cluster supports:

- A large collection of utilities, compilers and libraries
- Unlimited data transfers for up- and downloads
- Sharing data and analyses with other users
- Optional accelerators nodes to execute specific functions much faster than normal CPUs
- Optional Infiniband nodes for applications that require high I/O
- Optional high-memory node for analyses that require large datasets to be loaded in memory
- Own workspace with 200GB default home directory

- Installation of specific software packages and high-level support concerning the optimal use of the service

## Additional services

The following additional services are available on request.

### Customized LISA Compute Cluster workspace

The customized workspace can provide for example:

- An exclusive fileserver
- An exclusive login node with customized environment
- A customer driven tailored workspace

## Consultancy

SURFsara's advisors consult you on optimizing code for performance improvements.

## Visualization

Calculations often result in large amounts of data. Visualization can be of use in the interpretation of the result. We can support you on those visualizations for demonstrations and presentations or for easier research interpretation.

## Storage, Backup & Data Archive

Additional storage and optional backup for usage above 200 GB is available upon request. The LISA Compute Cluster is connected to SURFsara's Data Archive, for long-term and secure storage of your data which is also available upon request.

## Contact & Support

We are pleased to help you with gaining access to the LISA Compute Cluster, answer your questions or assist you to specific requests you may have about the service and SURFsara. Please just contact us by email via [info@surfsara.nl](mailto:info@surfsara.nl).

Specific details on obtaining accounts by affiliates of one of the Dutch Universities or Grand Technology Institutes can be found on our website. Detailed information on the LISA Compute Cluster is provided in the Service Level Specification.

Support for LISA Compute Cluster users may be requested via email: [helpdesk@surfsara.nl](mailto:helpdesk@surfsara.nl) or phone 020 - 800 1400. Support is available during office hours, i.e. between 9 a.m. and 5 p.m.

more information  
[www.surfsara.nl](http://www.surfsara.nl)

visit us  
Science Park 140  
1098 XG Amsterdam


contact us  
020-8001300  
[info@surfsara.nl](mailto:info@surfsara.nl)



## Use Cases

Here are a few examples to illustrate what the LISA Compute Cluster is used for. More references can be found on the SURFsara website.


### Psychiatric genetics, VU University Amsterdam



### Large-scale datasets

The VU researchers have been collaborating in international consortia as a part of the genetic research. These consortia share large-scale datasets: “The largest dataset is the one for schizophrenia, which includes data on 70,000 patients and 110,000 control group subjects (healthy individuals). We measured a total of one million genetic variants for each control group subject. The subsequent analysis obviously required a lot of computing power. We use the LISA Compute Cluster for statistical analysis: we assess each measured variant to determine whether it's more common in patients than the control group subjects. We also use the compute cluster to conduct advanced analyses, such as the interaction between various genetic variants and the effects of multiple variants from the same biological pathway.”

### Tata Steel's magnetic field research



### Magnetic field research

Little is known about the effects of magnetic fields. For example: what is the ideal strength of the magnet? And how should the poles be oriented? Van Vliet performs these computations on SURFsara's LISA Compute Cluster. “This is our first research project with them. Previously we calculated this kind of Computational Fluid Dynamics using commercial software on our own computers. But that became too expensive, which was part of the reason why we switched to the open-source software package ‘OpenFoam’.

### LISA Compute Cluster technical specifications

Below is a summary of the technical specifications (July 2015). It illustrates the quality and quantity of the LISA Compute Cluster computing capacity.

Operating System	Debian Linux
Full system	7.856 cores; 149 TFlop/s (peak performance); 26 TB memory
Mellanox nodes	32 nodes with InfiniBand network FDR: 56 Gbit/sec Latency FDR: 1.3 µsec
Accelerator nodes	Mona cluster; 8 nodes 16 Intel Xeon Phi 7120 accelerators
High Memory nodes	Node with 1TB RAM, 48 processor cores, 13TB scratch space
Memory	32 to 64 GB memory per node
Disk space	Separate file servers Default home file system: 200GB 750 to 870 GB of scratch space per node

more information  
[www.surfsara.nl](http://www.surfsara.nl)

visit us  
Science Park 140  
1098 XG Amsterdam

contact us  
020-8001300  
[info@surfsara.nl](mailto:info@surfsara.nl)

