Centre for Scientific and Technical Computing

Progress report - LUNARC 2016

Table of contents

Introduction

2016 was an exciting year for LUNARC. After a year of writing applications, procurement documents, and preparing our computer room, our new resource Aurora was installed in the beginning of the year. Aurora replaced our previous production resource, Alarik, which served us since 2011. In addition to the funding from SNIC, Aurora was co-funded by several research groups at LU (Astronomy, High Energy Physics, and Medicine) and the e@LU project provided funding for LU specific hardware, such as accelerator and large-memory nodes.

With Aurora, we introduced many new technologies for making high performance computing more accessible and easy to use, such as:

- LUNARC HPC Desktop - Remote desktop environment with visualisation capabilities.
- EasyBuild – Build and installation framework for maintaining scientific software on High Performance Computing (HPC) systems.
- Read the Docs – System for creating and hosting documentation, applied to a user documentation, which went through a complete overhaul.

During 2016 LUNARC helped many research groups with their scientific workflows (HUMLAB, LBIC, and Medicine). Our ongoing collaboration with MAX IV also intensified and a message of understanding (MoU) was agreed upon in the context of computation, storage, and workflows. Similarly, a closer relationship with High Energy Physics was established due to their private partition in the Aurora resource. LUNARC provided valuable knowledge about run-time compatibility issues and enabled them to run existing legacy code on the new hardware. In 2016, e@LU application experts started working within the areas of metadata ontologies and sensitive data management and workshops were held within the e@LU code-optimisation project.

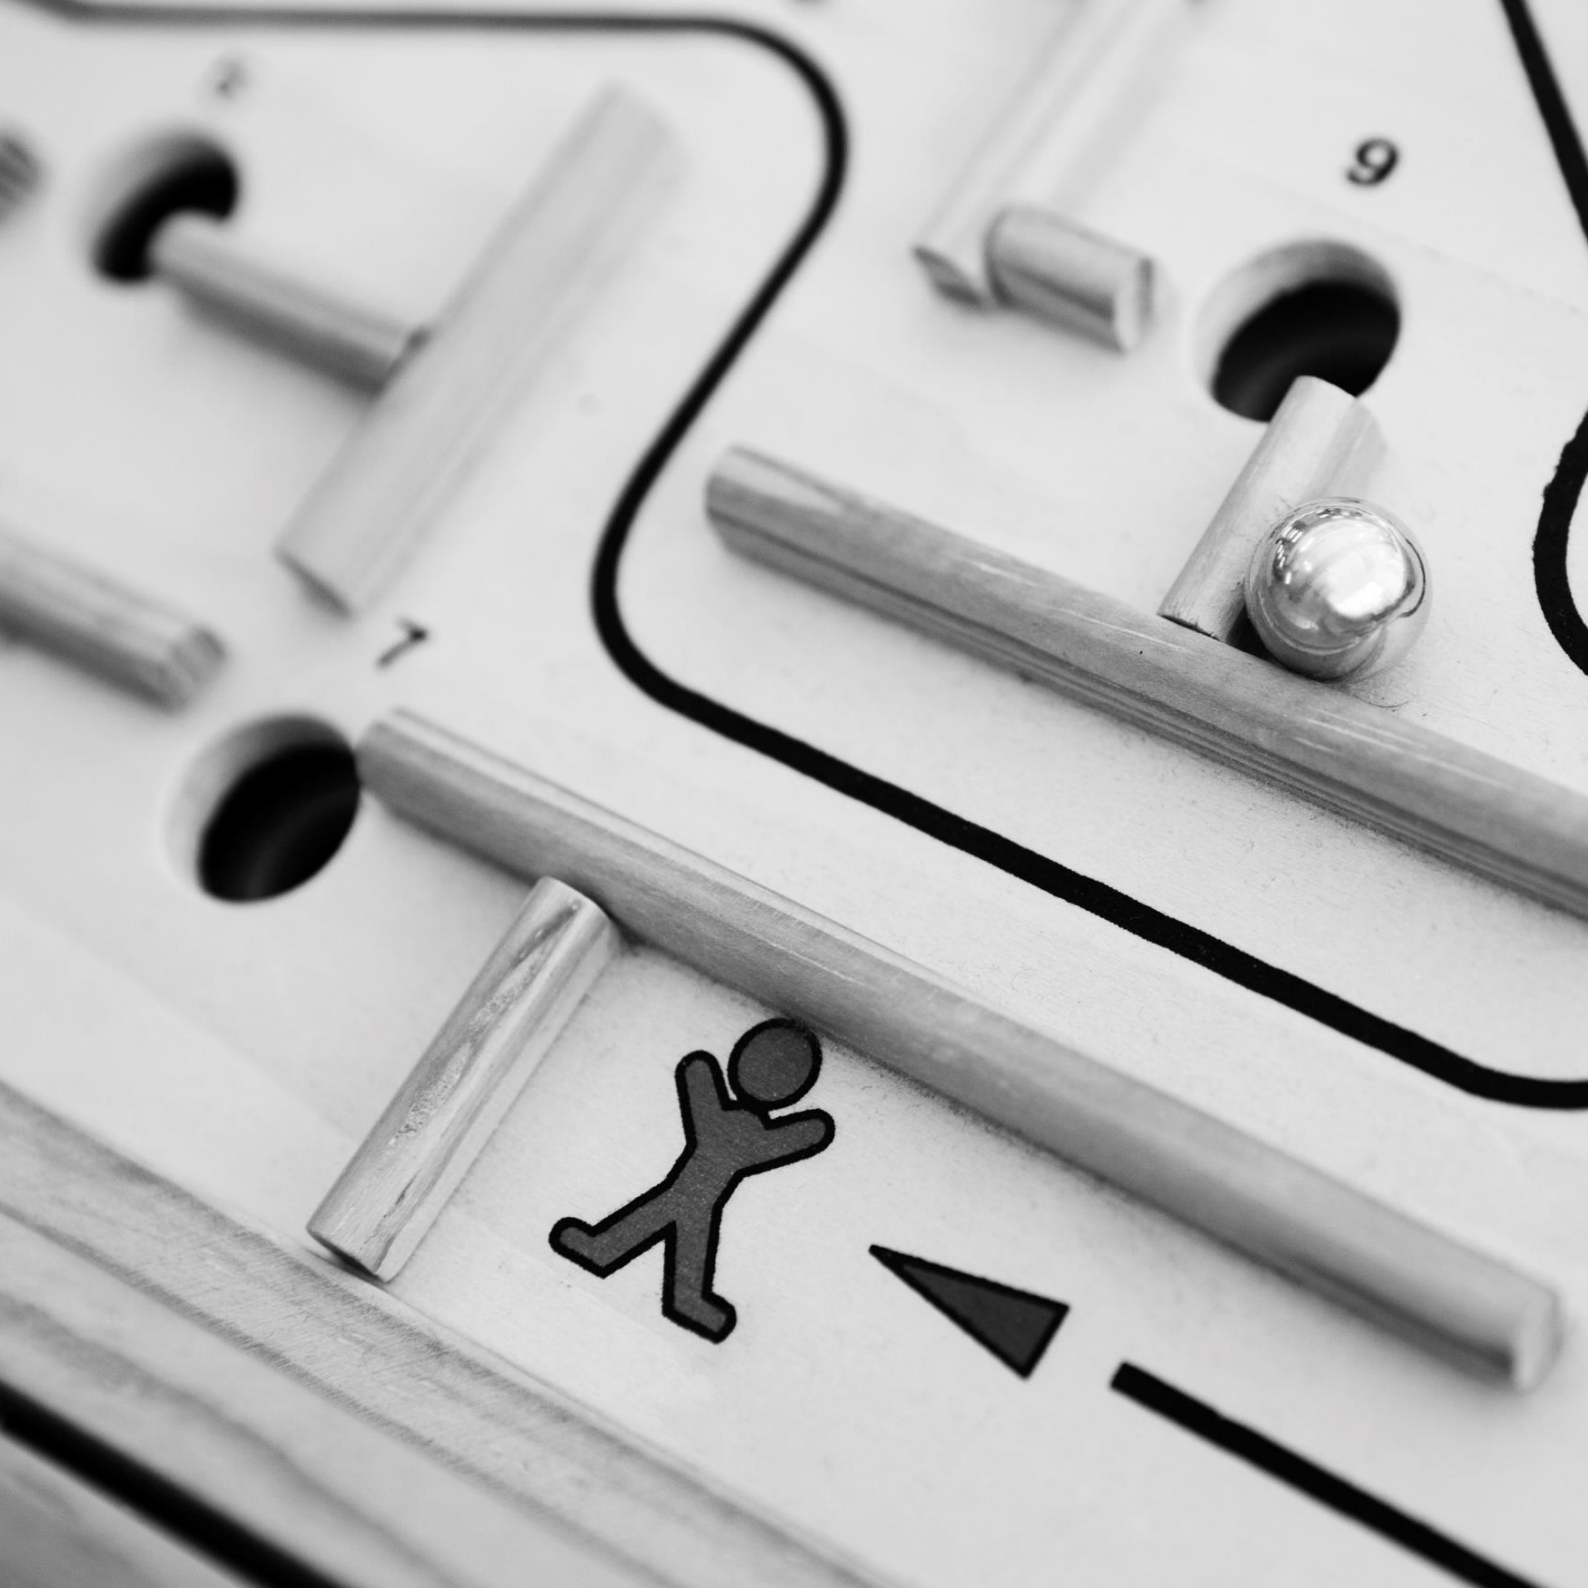
9
7

History

LUNARC, began operations on 20 November 1986 with the inauguration of an IBM 3090. In the early years, it was an informal organisation with three groups from the fields of computational chemistry and computational mechanics at the core, but the family grew with more and more research groups and in November 1996, LUNARC was reorganized as a formal centre, the Centre for Technical and Scientific Computing at Lund University.

In 2003, the Swedish National Infrastructure for Computing (SNIC) was formed as a metacentre for coordination and collaboration between HPC centres under the Swedish Research Council. LUNARC joined as a node together with five other university centres. In 2012, SNIC was reorganised as an independent national infrastructure hosted by Uppsala University, but still mainly funded by the Swedish Research Council.

As a university centre, LUNARC was originally linked to the Science Faculty, but on 2 June 2016, LUNARC was moved to LTH and earned a more distinct place in the organisation by becoming a division at the Department of Construction Sciences.

Mission

LUNARC should initiate and support activities and cooperation within e-science by providing computational resources and services for research and education. The operation includes resources and services in scientific computing, especially those calculations that require extensive computing capacity, analysis, and processing of large volumes of stored data. Furthermore, LUNARC should make efforts to develop and

adapt services and resources to the needs of different researchers and make resources easily accessible for researchers at Lund University and other universities.

Operational goals

The LUNARC Board shall, within LUNARCs activity:

- monitor progress and push for long-term capacity building;
- encourage external training, research and equipment resources supplied to the university;
- be the university's liaison with external stakeholders both nationally and internationally;
- develop the university's strategy for the provision of computing and storage resources for large research facilities MAX IV and ESS.

Current topics

During 2016 LUNARC's Board should:

- effectuate organizational move from the Faculty of Science;
- ensure that the role of LUNARC becomes more clear within the organization of LU / LTH and that LUNARC becomes more visible at the websites of LU and LTH.
- using existing projects, promote good skills in LU terms of comprehensive scientific computing and data storage;
- implement a long-term cooperation with MAX IV concerning the computation and storage;
- work towards LUNARC setting a clear and strong role within SNIC;
- work towards LUNARC and LDC relocating to the Science Village Scandinavia.

The LUNARC Experience

At LUNARC, we strive to provide an easy-to-use, well documented and efficient computational environment for both new and experienced users. In addition, we offer expertise for implementing customized workflows for computing and storage in connection to the existing LUNARC resources. The following sections describes activities directed towards improving the LUNARC experience.

Training

The key to facilitating the use of both basic functions and advanced features is a good training and educational program. For new users, we offer introductory training, which are half-day training events with lectures and hands-on exercises. The introductory training events are also offered in 2 parts, so that experienced users can take the lecture part only. These training events are very important to show users the most efficient way of using our resources.

In addition to the introductory training LUNARC also offers more advanced intensive 2-4 day courses in special topics such as:

- Parallel programming in OpenMP and MPI.
- Parallel programming in MATLAB and Scientific Python (Numpy)
- Debugging and optimization of scientific codes.

LUNARC is also involved in the COMPUTE research school at the Science Faculty, where we give longer courses in scientific programming topics.

Support

When using our resources, users often have questions and requests, these are handled through LUNARC support. LUNARC support is the first contact point for users and it is key to providing a good service to our users.

Support requests can be handled either by calling to a special service number or by sending an email to support@lunarc.lu.se. When sending a mail to LUNARC it will automatically be sent to the support system, where it will be given a ticket-number, so that it is tracked and not forgotten. A support request through the ticket system will be visible to all LUNARC staff. Anders Sjöström coordinates and distributes the tickets to the right person and also makes sure that they are handled and closed.

Sometimes a support request is difficult to handle by email and phone. In those cases, we schedule a meeting with the user either at the LUNARC offices or using Skype and TeamViewer.

The LUNARC web-pages also provide information on what happens at LUNARC and important information on the LUNARC service.

Documentation

Another important key in providing a good service is documentation. In 2016 we completely rewrote the documentation for using the LUNARC resources and presented it through Read The Docs. To be able to quickly update the documentation it is written in a simple text format (Markdown) and maintained in a GitHub source repository. When an update to the documentation has been edited, GitHub automatically initiates an update of the web pages.

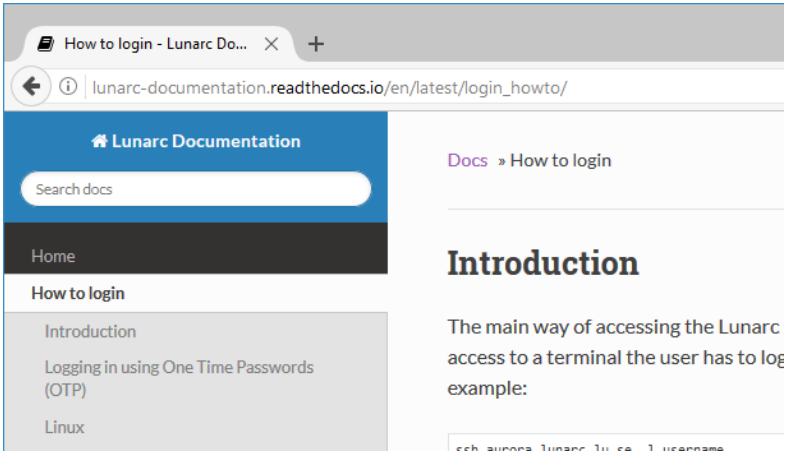

Figure 1 - LUNARC documentation pages

The documentation has the following sections:

- Quick start guides – Shorter guides to get users quickly acquainted with a topic.
- FAQ – Answers to the most common questions asked in the support.
- Longer comprehensive guides on how to use the resources with complete examples and templates that can be used.
- Special topics.
- Reference documentation.

The LUNARC documentation is regularly updated.

Software environment

LUNARC has 150-200 active users. They all use some form of scientific application, which can be:

- User-compiled packages.
- LUNARC-provided packages built from source.
- Commercial software provided as binary installs (MATLAB, ABAQUS, LS-DYNA and others).

To handle the amount of software packages and their dependencies on the underlying environment, a special system for compiling and maintaining software packages, EasyBuild, was deployed in 2016. This system combines a module system for easily selecting software packages, with a build system. EasyBuild comes with

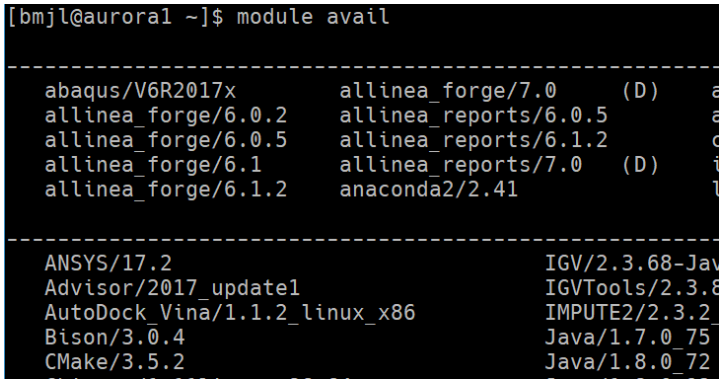

Figure 2 - Available software on Aurora

over 600 predefined recipes for building efficient software for scientific use. By deploying EasyBuild, build times for user-requested software has been reduced significantly and our staff can get more time to focus on other tasks. It has also enabled us to provide a much wider software portfolio for our users. Now, almost no software request is denied due to lack of staff time. Instead we can focus on providing application support and training to our users.

Remote desktop and visualization

Traditionally, HPC resources have been accessed through a terminal. This is an efficient way for many users to accomplish their computational tasks. However, the amount of data generated in many simulations are often extremely large and it is often difficult and takes a long time to download the data to a personal workstation for visualization or processing. Processing power available to the users at their home institutions can also be limited. To solve this LUNARC has pioneered the use of remote desktop services. During 2013 we deployed a test environment for Alarik. This test environment quickly became very popular with our users. At the time, we had 30-50 users continuously using the desktop environment. In 2016 we deployed a production desktop environment for our new resource Aurora, using a completely new architecture and backend.

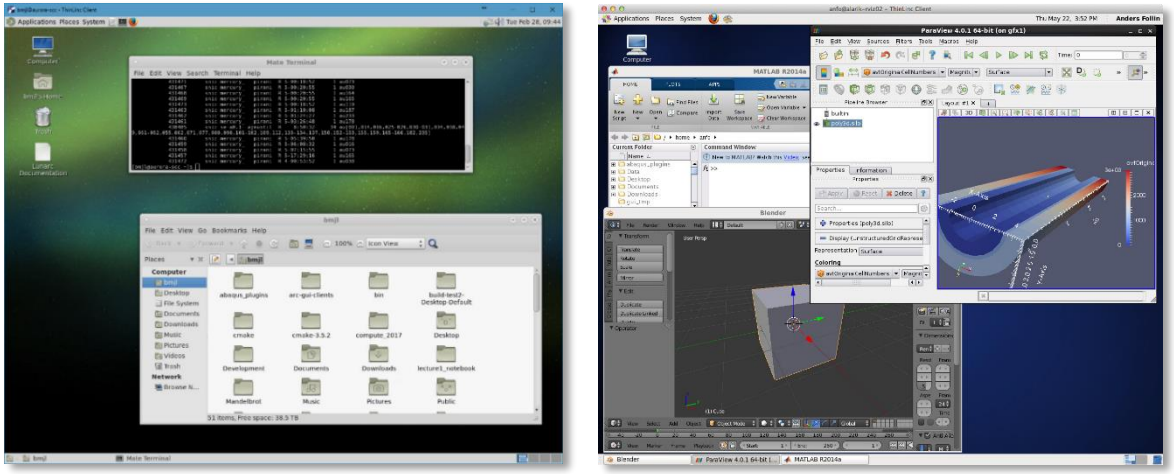

Figure 3 - LUNARC HPC Desktop

The new desktop environment has been successfully used by all categories of users, from novices to the more experienced. It has also been a valuable tool in the LUNARC training program, where we use the desktop as the standard environment for all training exercises. This removes the need to support different user environments for each course, giving the participants a consistent environment when taking part in our training and educational program.

Tailored workflow solutions

One area where LUNARC stands out as an HPC center is in designing and implementing customized HPC workflows for scientific user communities. LUNARC's project team consisting of experienced solution architects and software developers most often performs an initial in-depth assessment study in close collaboration with the end users. The solution is tailored for the specific user group, often consisting of researchers with no or very limited skills in high performance computing. Many user groups are not aware of the possibilities offered by LUNARC and how LUNARC's resources can benefit their research. The majority of the projects are related to data management in one way or another ranging from rather straightforward solutions to very complex end-to-end implementations involving custom software development and associated computing hardware.

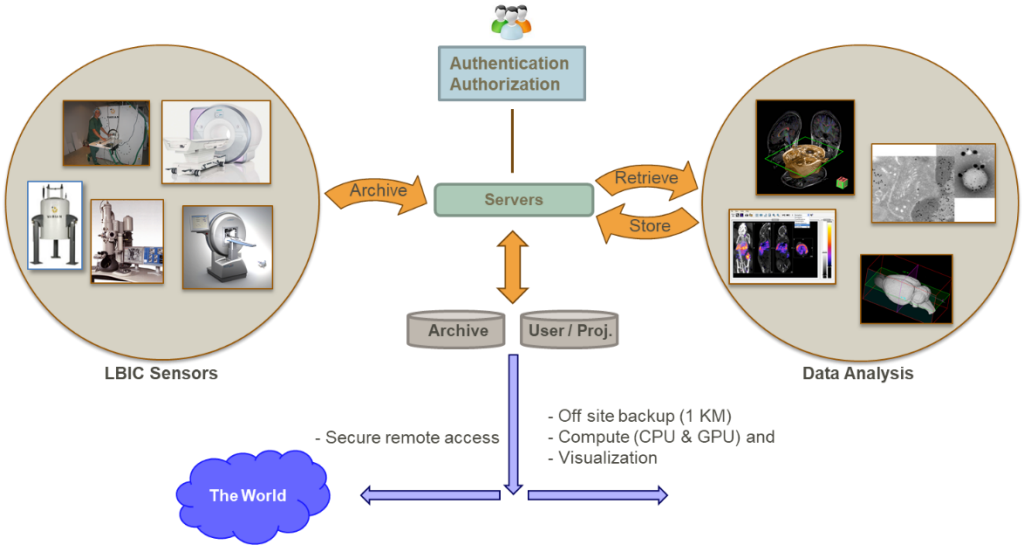


AURORA

Organisation

Board

Lunarc is lead by a board that has the overall responsibility for the activities at the center. The board also sets the guidelines on how LUNARC should operate, approves the yearly budget and development plan and issues annual reports. The development plan should ensure that LUNARC is a resource available for all of Lund University.

The composition of the board reflects the different activities at Lund University, requiring the services the Lunarc provide. The board consists of members from Economics, Humanistic and Theological faculties, the Faculty of Science, Faculty of Medicine and LTH. The chairman is appointed by the dean of LTH.

In 2016 the board consisted of:

Chairman	Erik Swietlicki
	Aylin Ahadi
	Christofer Edling
	Darren Spruce
	Jette Guldborg Petersen
	Kirk Scott
	Marianne Gullberg
	Marie Skepö
	Mauno Vihinen
	Melvyn B Davies

Director

The daily operation of the center is led by a director appointed for a three-year period by the dean of LTH. The director is responsible for ensuring the ongoing activities of the center is in line with the strategies and guidelines provided by the board.

The current director is Jonas Lindemann.

Assistant director

The board has also appointed a deputy director, assisting the director in the ongoing activities of the center.

Anders Follin is currently the deputy director.

PDC Center for
High Performance Computing

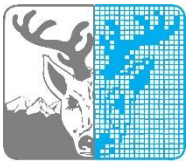

LUNARC within SNIC

LUNARC is one of the 6 SNIC centers in Sweden. SNIC funds and coordinates computational and storage resources among these centers.

SNIC Resources at LUNARC

In 2016 LUNARC operated 3 HPC clusters for SNIC:

- Aurora – 180 node cluster
- Aurora-Grid – 20 node cluster for WLCG operations
- Erik – A 68 GPU HPC cluster

In addition to computational resources, LUNARC operates ~ 2 PB of dCache storage pool for SNIC National Storage and a Lustre base storage system from DDN, which is used as a center storage system.

Training

LUNARC is actively participating in the SNIC coordinated training program, which our staff member Joachim Hein is heading. Please refer to the separate progress report for this activity for full details on the training provided by LUNARC.

Outreach

LUNARC organised and hosted the SNIC user forum 2016

LUNARC participates in LU events that are related to LUNARC and SNIC activities. LUNARC also actively takes part in the SeSE research school giving courses in scientific programming. These courses provide an excellent opportunity to reach new users and PhD students that potentially can become users of SNIC and LUNARC resources.

NeIC

LUNARC is also collaborating around training in the nordic countries through the NeIC initiative.

SNIC Future Architectures

Within SNIC LUNARC coordinates the research and development program SNIC Future Architectures program. Within this program centers can apply funding for evaluating developing hardware for upcoming computational and storage resources. Currently the program have evaluated:

- ARM based servers and switched for HPC use.
- Power-based system with GPU:s and NVlink interconnect.
- Hardware and virtualization solutions for remote desktop architectures.


REGIA · ACADEMIA · CAROLINA

LUNARC within Lund University

Outreach

In 2016, LUNARC was actively involved in the advanced study group INTEGRATE at Lund University (http://www.astro.lu.se/INTEGRATE/). The INTEGRATE ASG brought together researchers utilising computers as an integral part of their efforts to answer a broad range of research questions. Research areas included the physical sciences, medical imaging, large facilities such as MAX IV and CERN, environmental modelling, and the humanities. There were four one day events during HT 2016, where several LUNARC staff made contributions. Anders Follin of LUNARC served in the INTEGRATE steering group.

In addition to his participation of LUNARC's involvement with INTEGRATE, e@LU and him leading LUNARC's EasyBuild efforts, Joachim Hein provided in-depth support and outreach to the following users and groups

- Discussing software roll out on the Max IV cluster via EasyBuild with the Maxlab application experts.
- Utilising EasyBuild to deploy software on SNIC installations. Discussion group with staff from C3SE, HPC2N and LUNARC
- Presentation of HPC for the PhD students in Mathematics at LU
- Enabling George Masterton (Theoretical Philosophy) to run his multi-agent modelling software in the job schedulter on Aurora.

Anders Sjöström has also provided in-depth support and outreach in:

- Parallel computations using MATLAB within the application experts collaboration.
- Database calculations using MATLAB.

- Discussed with Sebastian Wasserstrom from the Lund Biomedical Centre (BMC) on the use of parallel toolbox for calculation speedup in datamanagement and image recognition.

Training

To be able to use the LUNARC resources effectively, it is important that users are given relevant training, so they can use resources at LUNARC effectively. LUNARC has developed a significant training program for Lund University users that includes courses developed by LUNARC as well as inviting external trainers when required.

During 2016 LUNARC provided training in the following areas:

- Introductory HPC courses (LUNARC material and instructors)
- Parallel computing using shared memory and OpenMP message passing with MPI (LUNARC courses and instructors)
- Using MATLAB in a HPC environment. (Mathworks instructors)
- Using the Intel Compiler (Intel instructors)
- Using the DDT debugger (Alinea instructors)

We also provide training in Fortran and Python.

COMPUTE School

LUNARC actively participates in the COMPUTE research school, both in the steering group (Magnus Ullner and Joachim Hein) as well as a course provider (Joachim Hein, Jonas Lindemann)

Steering group. Magnus/Joachim. Attending meetings. Planning for courses in 2017

Code consultancy

LUNARC also provides guidance on code optimization by offering consulatancy on how to optimize and parallise code, to better take advantage of the LUNARC resources.

Ongoing collaborations

MAX IV

The Scientific Data Management project (SDM) is now going into production (Developed with funding from SNIC). This system will consist of an analysis desktop and a compute and storage infrastructure. An agreement between MAX IV and LUNARC now exists on the operation of the MAX IV online and offline analysis and storage systems (funded by LU and MAX IV).

LBIC

LBIC Imaging data management project: In this project a customised system is designed with purpose to increase the throughput of the imaging modality (i.e. maximize the ROI) and add scalable compute and post processing functionality to an advanced optical imaging platform. The microscopes (confocal, STORM and Super HiRes) are integrated into a customized scientific data management workflow including customized software components in combination with a hardware accelerated imaging backend. The solution combines batch-oriented data processing with interactive remote visualization. LUNARC SHIP is used for conveinent and high performance parallel data transfer from the biomedical modalities to LUNARC HPC center. The same data management framework can be utilised for other imaging platforms such as MR or NM.

Coordination and architectural work for the national 7T MRI archive including local and remote data access and radiology imaging. The system was put into production autumn 2015 and is serving medical researchers from all over Sweden. LUNARC is still actively engaged in the project with an advisory- and coordinating role.

CCI Sahlgrenska

This project is similar to the LBIC project but also includes some specific add-ons. CCI 's solution will include more local imaging functionality and interface to a resource pool hosted by LUNARC. Currently we are implementing a new software solution, LUNARC Ship, for automatically transferring data from CCI to the resource pool.

HUMLAB

LUNARC has now completed the project to improve the HUMLAB laboratory infrastructure to increase security of laboratory workflows.

Archeology

In addition to this project we have also involved in a project with archeology to provide remote HPC facilities and services, so that they can do quick assessments in the field when they are on site for 3D scanning with drones. A end-to-end solution has been implemented utilizing a cluster aware version of the 3D reconstruction application. Both CPU- and GPU-based interactive as well as batch oriented workflows has been implemented accessible from the LUNARC HPC Desktop.

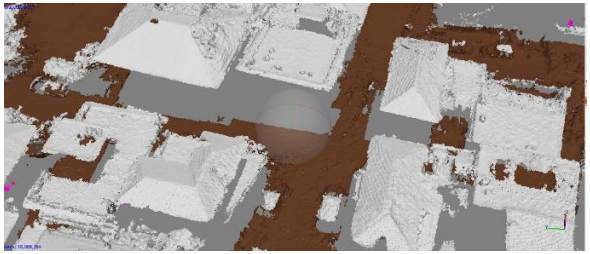

High Energy Physics

High Energy Physics has procured a large partition of 26 nodes in the Aurora cluster. Initially the usage was low and they contacted us on how we could improve the environment, with regards to workflows and scientific software environment. One of the biggest issues was how to be able to run the CERN scientific software environment on the Aurora resource. The CERN software is compiled for Scientific Linux 5.x and 6.x. The aurora resource is based on CentOS 7, which have limited backwards compatibility with Scientific Linux 5.x and 6.x. To solve this, we developed a special module, based on the Singularity project, that provide a compatible runtime environment, enabling the CERN software to run on Aurora without recompilation and the use of virtual machines.

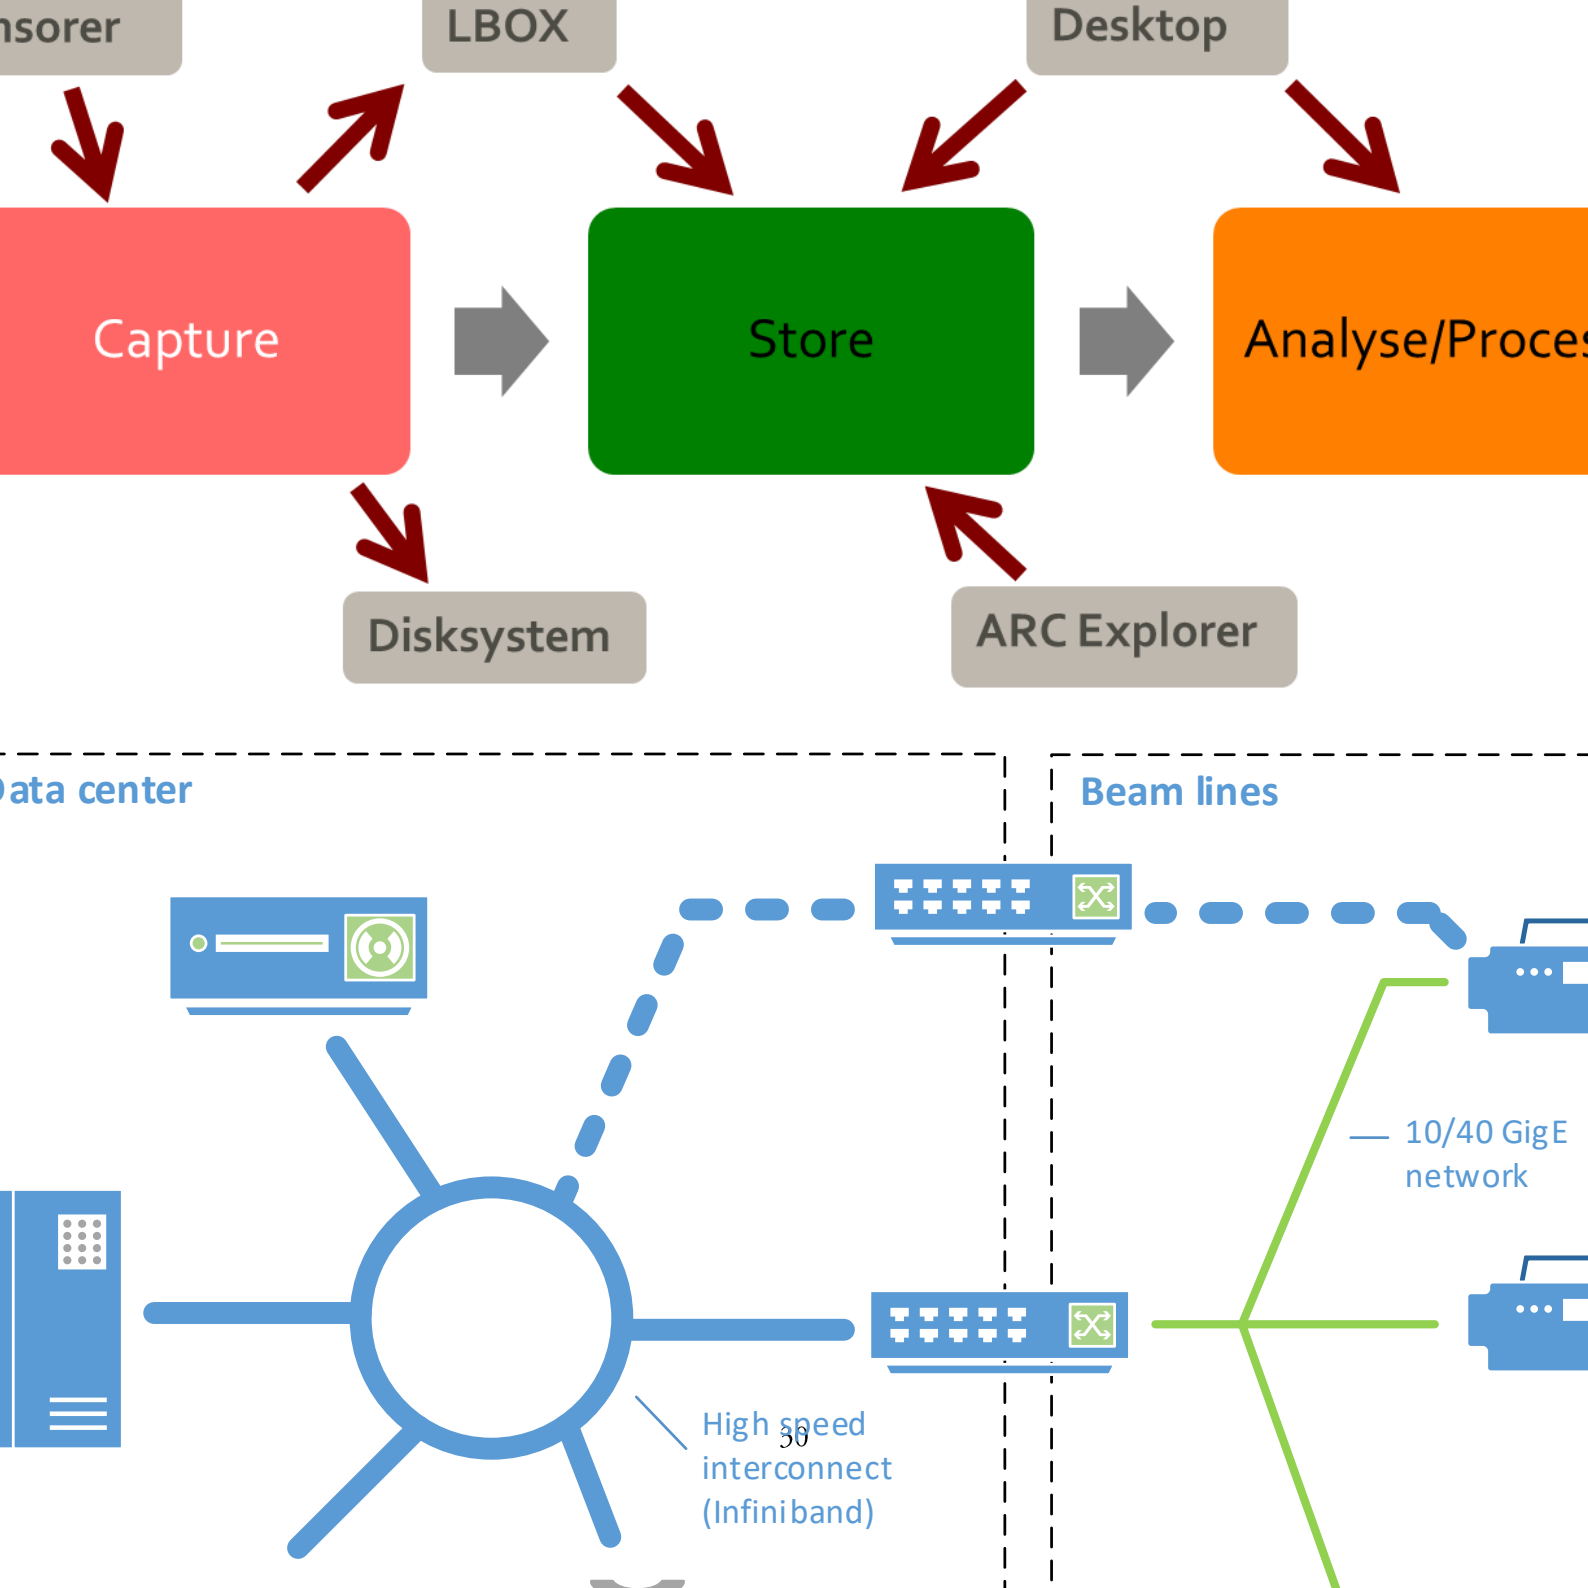
nsorer
LBOX
Desktop
Capture
Store
Analyse/Proces
Disksystem
ARC Explorer
ata center
Beam lines
10/40 GigE
network
High speed
interconnect
(Infiniband)

e@LU

e@LU is a collaborative project within the University relating to e-infrastructure, with particular focus on the facilities and organizations within e-science, suchj, MAX IV, BMC, LBIC and HUMLAB. These have similar needs in terms of data management solutions, but also analysis data and modeling and simulation. The project includes hardware and several application experts, which will provide support to different areas.

e@LU administered by LUNARC and governed by the LUNARC board.

Application experts

An important part of the project is the application experts to help research groups with questions that they themselves do not have the time or knowledge to manage. Since May 2nd, e@LU has 4 applications experts in the following areas:

Metadata, ontologies and data provenance

Application experts: Monica Lassi, Jörgen Eriksson, Maria Johnsson
Project leader: Monica Lassi / Alex Vermuelen

Science is all about data and now Science is turning into eScience, where all data is transferred to and exchanged in digital form. This opens up vast possibilities but also presents some major challenges. In order for humans and computers to understand the data outside of the context of the individual or group of scientists that created the data, and to preserve the usability over time, the issue of enriching the digital data with digital metadata including data provenance is becoming an urgent matter. This has led to a strongly developing and dynamic field where semantic interoperability is the keyword. This is the consistency of meaning and understanding within and between the large amount of data produces and data users. As the scientific data is complex by nature and requires understanding of the context of the science that produces and consumes the data this is far from trivial but it is essential for almost all sciences in the future.

The problem is relatively large in environmental and also social/medical sciences, where definitions are sometimes trivial and/or widely different among communities that study the same parameters from different viewpoints or locations. Just one example is the definition of what is "a forest in land-use classification". ´

Data management with special attention to safety aspects with regard to human data

Project leader: Kirk Scott
Application experts: Gustav Öberg

The discussion rather quickly converged on what all participants see as an urgent need at the university level. The manipulation, coordination, and storage of personally identifiable data are topics that concern researchers at multiple faculties and departments, but are completely absent at the university level. These topics are very broad, and beyond the reach of a single half-time employee. The group felt strongly that all topics are vital to the competitiveness of the university and should be accounted for in the medium to long term. For the purposes of the pilot program e@LU, it was decided that a single task of designing a common strategy for data security is the most pressing.

It is the conclusion of this group that one of the positions funded by e@LU should be tasked with an evaluation of data security at the university. Within this task, it is envisioned that this person would perform an inventory of existing facilities and practices throughout the university, and assess how these practices correspond to our legal obligations. As a part of this inventory, we also believe that the person employed should serve as a key person in identifying research similarities and potential synergies – a subtask which will hopefully lead to potential collaboration across disciplines and faculties. Given this subtask, the person selected must have a good insight into the research process.

As a long-term follow-up, we envision that this person could be the seed for a new center aimed at supporting and facilitating research using personally identifiable data at the university.

Parallelization of TEMPO

Project leader: Philipp Birken
Application expert: Peter Meisrimel

Goal: TEMPO is a C++ library for numerical time integration of Initial Value Problems. Large parts of it have been coded by Philipp Birken and Peter Meisrimel. It has been partly parallelized using MPI. This should be completed, the documentation finished and the library published under an open source license. This code is so far used by the numerics group in Kassel, as well as the numerics group in Lund. The Lund Observatory (Chao-Chin Yang) has expressed interest to use this code in their astrophysical simulations.

Multiphysics multicore multirate methods

Application expert: Azahar Monge
Project leader: Philipp Birken

Goal: We are working on time integration for multiphysics problems. Current methods are either parallel or time adaptive and we have an idea to get the best of both worlds. The mathematical development is part of the PhD thesis of Azahar Monge. Within e@LU, we would bring this method in a software form that allows use by others as well. In particular, the Department of Physical Geography and Ecosystem Science (Mikhail Mishurov) is interested to apply this methodology to their code RCCS-GUESS, which couples the codes LPJ-GUESS and RCA to get couple a dynamic vegetation model with a meteorology code.

Code optimization group

Application expert: Joachim Hein, Anders Sjöström
Project leader: Jonas Lindemann

In the original proposal of e@LU there was an application expert position for accelerators. Instead of hiring this expert it was decided to create group consisting of the experts at LUNARC and the groups with this need. During 2016 we started a series of workshops where we discuss different topics related to algorithms and performance optmisiation. For every workshop one of the participants gives a short present a problem from their field and we have a discussion around the topic on how to optmise the algorithm and how it should be implemented to be as efficient as possible.

Collaboration projects

The e@LU project also aims to provide support and guidance on how to implement scientific workflows. During 2016 LUNARC has developed this collaboration projects to more long-term collaborations. Some of these collaborations are described in the following sections.

MAX IV

The collaboration with MAX IV has developed from the SNIC funded project to a long-term collaboration on storage and workflow. During 2016 LUNARC helped MAX IV design and procure a small HPC resource with a connected GPFS storage system, which is jointly managed and operated. Some of the storage equipment is also mirrored in the LUNARC computer room. In 2017 the MAX IV storage will also be made available on the Aurora system, so that MAX IV scientists can do offline analysis with beamline data using the same software environment as used in the online analysis at MAX IV.

BMC

In 2016 BMC contacted us on how to move from their existing computing environment to instead using LUNARC resources. Several meetings where held, which resulted in a storage and compute consolidation. One of the groups now have their own desktop frontend and storage integrated in Aurora. During the process, several groups expressed the interest to be able to do computation on sensitive data. This led to the development of the L-SENS project described in a following section.

Development projects

e@LU has also funded the development of software and services that is required by the different collaboration projects. Development is strongly driven by the requirements and needs of the research groups.

File transfer service – L-SHIP

In discussions with several sensor-based research groups such as and biomedical materials research voltage, it has been a strong need for an effective and user data transmission service. The reasons are several:

- Releasing the sensor as soon as possible for new data collection. First, to increase the utilization and thus exchange of investment (ROI) and partly to create a more efficient workflow with higher throughput rate.
- Enabling a high performance, safe, cost effective and highly flexible workflow utilizing University common resources of supercomputer center (HPC center) of the provision of a graphical remote desktop solution. The researcher is thanks to this tech-nical solution not bound to a specific workstation (computer) when the central resources are to-reach-from arbitrary place with either wire or wireless network connection.
- Centralization of hardware for storing, processing and post-processing (imaging) creates a more cost effective environment not only from a hardware perspective but also from Softwrulicenser. More users can be in a so-called floating license jointly utilize a common lithium censpool.

To solve these problems, LUNARC utecklat L-SHIP, a service that manages data transfer in the background and frees up the equipment for use.

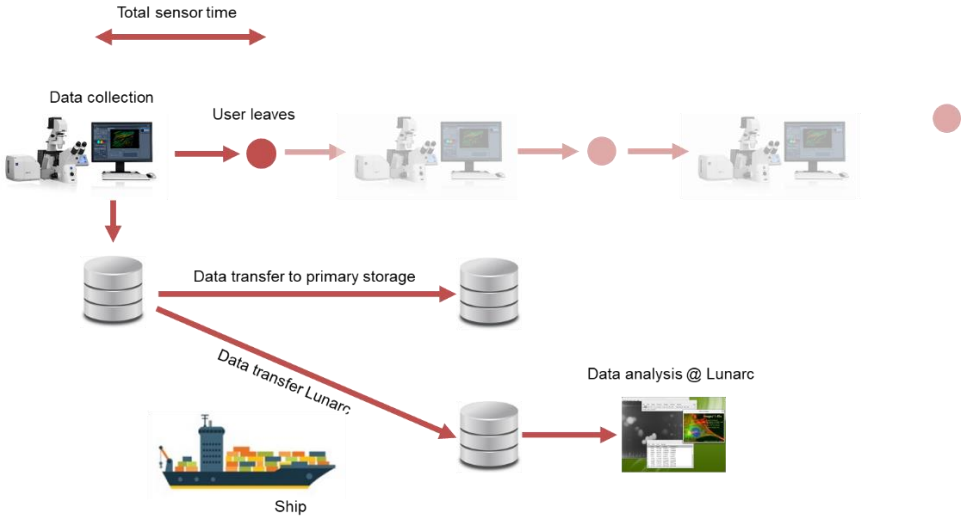

The L-SHIP concept

The ship system was developed by Roger Larsson and the application group at LUNARC during 2016 and will be put into production during 2017Q1.

Aurora – L-SENS

Aurora L-SENS is an initiative to provide a secure high performance computing facility for doing analysis on sensitive data. Currently 4 research groups are involved in this project. Each group is provided their own secure computing silo, with their remote desktop frontend, compute nodes and storage. Each silo is physically separated from each other, to provide the maximum security. Authentication is done using the same 2-factor authentication system as on other LUNARC resources. Physical access is also limited by secured doors in the compute racks.

Project was initiated during 2016 and will be installed during 2017.

Hardware resources for LU users

The e@LU project also finances additional dedicated hardware for the Aurora system consisting of large memory nodes and GPU-nodes.


LUNARC
AURORA
A
B

Resources

Aurora

Aurora consists out of 180 compute nodes for SNIC use and over 50 compute nodes funded by research groups at Lund University. Each node has two Intel Xeon E5-2650 v3 processors (Haswell), offering 20 compute cores per node. The nodes have 64 GB of DDR4 ram installed.

System information

Hostname	aurora.LUNARC.lu.se
Queueing system	SLURM
Home space	/home NFS mounted, available on all nodes
Work space	/LUNARC/nobackup 640 TB DDN EXAScaler Lustre filesystem.
Linux distribution	CentOS 7.2 x86_64 (RHEL7 compatible)
Software	organised in a hierachical module system

Node information

CPU	2 Intel Xeon E5-2650 v3 (2.3 Ghz, 10-core)
Memory	64 GB (3.2 GB/core)
Local disk	1.7 TB, temporary directory given by $SNIC_TMP or $TMPDIR
Interconnect	4xFDR InfiniBand

Aurora-Grid

Aurora-Grid consists out of 20 compute nodes for SNIC-WLCG. Each node has two Intel Xeon E5-2650 v3 processors (Haswell), offering 20 compute cores per node. The nodes have 128 GB of DDR4 ram installed.

System informaation

Hostname	aurora.LUNARC.lu.se
Queueing system	SLURM
Linux distribution	CentOS 7.2 x86_64 (RHEL7 compatible)

Node information

CPU	2 Intel Xeon E5-2650 v3 (2.3 Ghz, 10-core)
Memory	64 GB (3.2 GB/core)
Local disk	1.7 TB, temporary directory given by $SNIC_TMP or $TMPDIR
Interconnect	4xFDR InfiniBand

Erik

Erik is based on HP Proliant Gen8 servers with 16 SL250 nodes and 8 SL270 nodes each with dual 64-bit, 8-core Intel Xeon E5-2650 2.00 GHz processors, 384 processor cores in total. The system includes 68 Nvidia Tesla K20m GPU cards (2 in each SL250 node and 4 in each SL270 node). The system interconnect is FDR Infiniband and Gigabit Ethernet.

In addition to the GPU equipped nodes Erik includes an additional node with 2 Intel Xeon Phi accelerator cards in a separate partition.

System information

Hostname	erik.LUNARC.lu.se
Queue system	SLURM
Scheduling policy	Fairshare
Global filesystem	/LUNARC/nobackup HP X9000 Parallel file system, available on all nodes.
Home	/home NFS mounted, available on all nodes.
Compilers	GNU, Intel, PGI
Software	Module-based software handling

Node information

Thin nodes

CPU	2 E5-2650 (2.0 Ghz, 8-core)
GPU	2 Nvidia K20
Memory	64 Gb
Linux distribution	Scientific Linux 6 x86_64 (RHEL6 compatible)
Local disk	250 Gb, temporary directory given by $SNIC_TMP or $TMPDIR

Fat nodes

CPU	2 E5-2650 (2.0 Ghz, 8-core)
GPU	2 Nvidia K20
Memory	64 Gb
Linux distribution	Scientific Linux 6 x86_64 (RHEL6 compatible)
Local disk	250 Gb, temporary directory given by $SNIC_TMP or $TMPDIR

Aurora-L-SENS

During fall 2016, there have been discussions with a number of stakeholders in the medical faculty regarding the storage and calculation of sensitive data on LUNARC's resources. This has led to the conclusion that we needed to develop a technically acceptable solution to handle this on LUNARC. The proposal describes a system that is at a good level to the requirements that may be required for the purpose.

The system will be operational during spring 2017.

Hardware is equivalent to Aurora.

Källa: Science Village Scandinavia

Future Vision

LUNARC has during the last years grown in scope and increased its involvement in many research groups and infrastructures such as MAX IV and BMC. In 5-10 years, we foresee that LUNARC will grow both in increased organization and increased infrastructure needs. To accommodate this growth, LUNARC will move to Science Village Scandinavia.

A new more efficient data center will be part of the move to Science Village Scandinavia. The new center will be designed to be very environmentally friendly, use excess energy from MAX IV and ESS for cooling as well as recycling heat from the data center. By placing the data center close to the MAX IV and ESS, it can be placed close to high speed data links going to and from these infrastructures, enabling storage and computing resources to work directly on produced data.

A key to providing excellent services to LU users is to have experienced application experts in both domain specific and generic areas. In the near future, we see an increase in the size of the application expert group at LUNARC to accommodate this. LUNARC's new premises will include not only a state of the art data-center and office space but also an education center with collaborative seminar- and meeting rooms to enable an even better knowledge transfer and teaching environment.

LUNARC will also continue to develop user-friendly computing and storage resources, enabling these resources for more users at LU. In addition to the current offering of computational resources, LUNARC will also provide a user-friendly Cloud platform providing a richer service offering for our user needs.

LUNARC
Institutionen för Byggvetenskaper
Lunds Univeristet
V-huset, våning 1
John Ericssons väg 1
221 00 Lund

Hemsida:
www.lunarc.lu.se

Kontaktperson:
jonas.lindemann@lunarc.lu.se