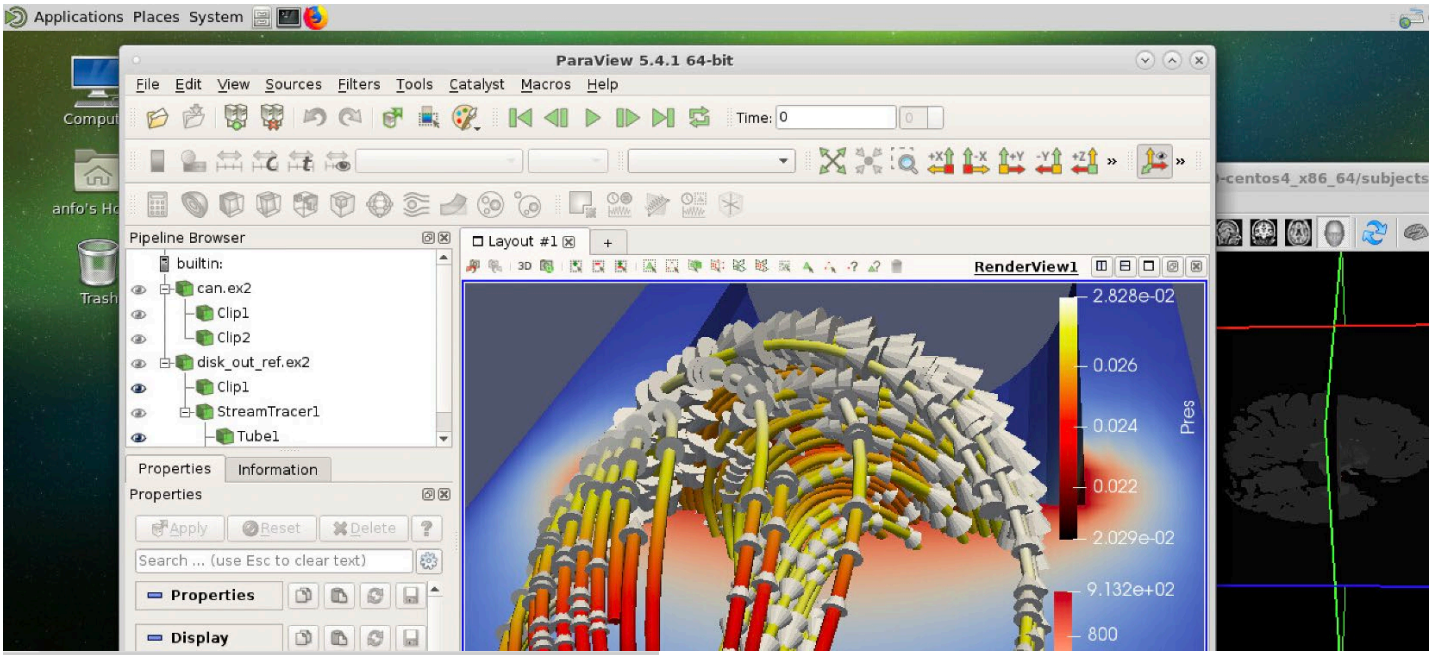


```
In [1]: import numpy as np

In [2]: import matplotlib.pylab as plt

In [7]: x = np.linspace(-2*np.pi, 2*np.pi, 100)
        y = np.sin(x)

In [8]: plt.plot(x, y)

Out[8]: [<matplotlib.lines.Line2D at 0x7efec725bdd0>
```

LUNARC

Progress Report 2017 – 2019

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

Hemsida:
www.lunarc.lu.se

Kontaktpersoner:
erik.swietlicki@nuclear.lu.se
jonas.lindemann@lunarc.lu.se
anders.follin@lunarc.lu.se

Introduction

2017-2019 has been exciting years for LUNARC. After the introduction of the new resource Aurora in 2016 the resource became very popular both for LU and SNIC users. During 2017 and forward the resource also established itself as a important base infrastructure for specialised resources ranging from high energy physics, medicine, archaeology to astronomy. LUNARC's datacenter does also play an important role in the university wide effort to consolidate energy consuming compute clusters in one optimized location both from an environmental and system support perspective.

Some examples:

- A large amount (25+) compute nodes financed by Astronomy and High Energy Physics
- Custom high-end graphic frontends for Physical Geography and Ecosystem Science
- Compute and acceleration cluster for Bioinformatics
- Research databases and compute portals for Integrated Carbon Observation System (ICOS)

Several of the projects that were initiated with Aurora has also been developed further since 2016.

- LUNARC HPC Desktop – Has been developed into a production platform for interactive HPC and Visualisation supporting hardware acceleration and Windows based applications
- The build and installation framework, EasyBuild, has been used to install over 800 scientific application packages enabling LUNARC to support a wide group of users over many research areas.
- The documentation has added many user guides and best practice documents for multiple usage scenarios.

Since 2016 LUNARC helped many research groups with their scientific workflows (HUMLAB, LBIC, LUBI, Region Skåne and BMC). Our ongoing collaboration with MAX IV continues. Beamline data can easily be made available directly on the Aurora resource, so that researchers does not have to move data to Aurora before performing analysis tasks.

During 2017 work started on the development of computational resources for sensitive data. This resulted in the L-SENS resources at LUNARC in 2018. Currently LUNARC operates 6 clusters in the L-SENS project and an additional standalone resource for the LUBI-project.

Following the success of the L-SENS project we were offered to develop and procure a similar resource for Region Skåne, RS-SENS, which is currently in production.

During 2020 we will also co-host resources for sensitive data from LUCC - Lund University Cancer Centre.

Since 2016 the LUNARC user base has grown to over 700 active users from a wide range of scientific areas. A recent trend is that we see many new users from research areas previously not using HPC resources. An example of this is the archaeology department which is using LUNARC resources for 3D scanning using hardware accelerated graphics and GPUs.

LUNARC has also expanded it training and educational program significantly during 2017-2019. LUNARC provides training for new users both för LU as well as for SNIC users. In addition to training, LUNARC also participates in research schools such as COMPUTE and SeSE where we provide training in parallel programming, C++, Fortran and Scientific Python.

Since 2018, LUNARC has added expertise in the field of research data strategies and taken part in the development of the research data strategy of LU.

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 to work 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 2017-2019 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.

Organisation

LUNARC is led 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 and decides on a 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 2019 the board consisted of:

Chairman	Erik Swietlicki
	Aylin Ahadi
	Darren Spruce (Adjunct to the board)
	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 technical director and deputy director.

Usage of LUNARC resources

LUNARC has about 600 active users of around 1200 users in total. These users are working in around 180 active projects with 166 PI:s. Most of the users (80-90%) are from Lund University. Lund University usage of SNIC resources is around 5%. A large part of this is from the Aurora resource.

Figure 1 shows the distribution of users in 2019, which good representation of usage since 2016. The largest users of Aurora are from N-fak at 70% followed by LTH at 16,6%. Biggest external user is Linné University 1t 7,2% followed by KTH at 1,7%. MAX IV usage is around 1%, but we expect to see this number to grow when they move their offline analysis from MAX IV resources to LUNARC resources.

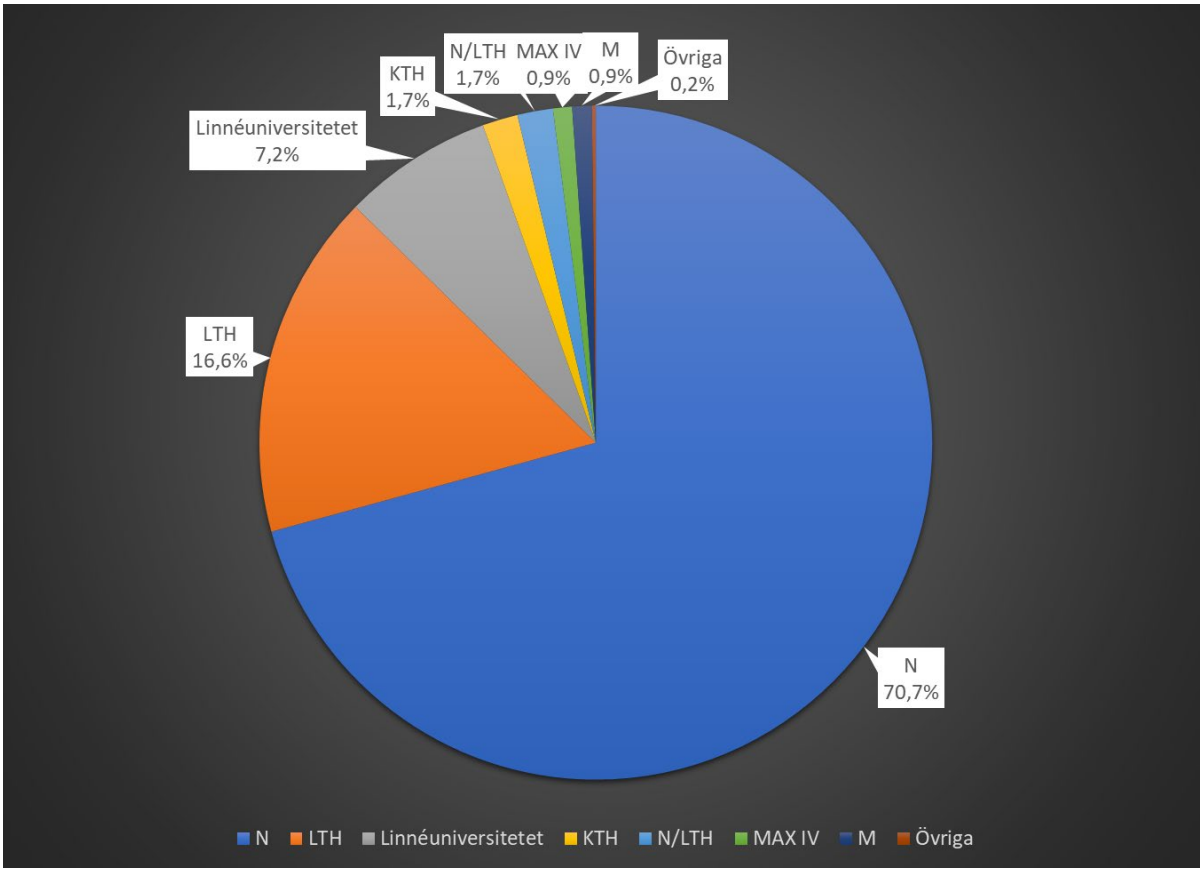

Figure 1 - Total usage of Aurora in 2019

A positive development is the increase of usage from the humanities. This increase has mainly been focused around interactive HPC and visualization using the L-VIS system. The main use is on 3D reconstruction of image data acquired by drone photography.

Another significant increase is the faculty of medicines usage of LUNARC resources. This has been facilitated by the provision of special resources for analysis of sensitive data. The resources do not make a large usage share of the total resource use at LU, but they enable research on data that can't be processed on more general-purpose resources.

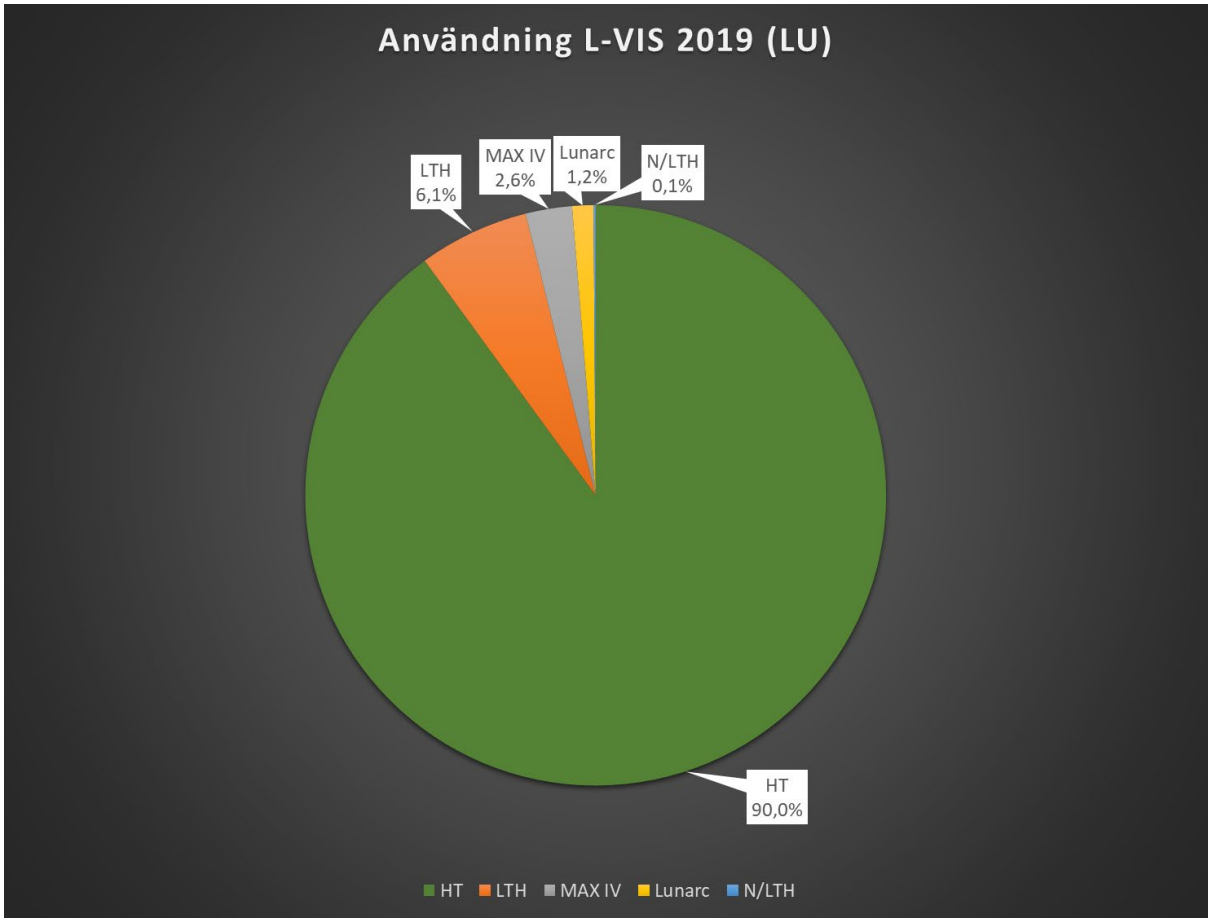

Figure 2 - Usage of L-VIS in 2019

LUNARC within LU

Outreach

LUNARC has continuously been working with research groups to facilitate and optimize their workflows around computational, visualization and storage resources. Outreach efforts are done in multiple ways. Some of these are listed below:

- Meetings with new research groups that need help in getting their workflows on to the resources provided by LUNARC. Our staff then guides them on how to get access to the resources.
- Application experts works within their fields to give directed guidance to researchers to enable efficient use of the LUNARC resources
- Workshops and training events are given regularly to educate users in different topics around high performance computing (HPC), visualization and storage.

LUNARC has also added a twitter channel (@LUNARC_LU) in which we tweet what is happening at LUNARC and retweet interesting news on HPC, visualization and other relevant topics that can be of general interest to our users.

Training and education

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. Coordinated by Joachim Hein, 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.

Training and edcucational activities in 2017

- An introduction to the LUNARC HPC Desktop, 19 December 2017
- Hierarchical modules and software selection, 5 December 2017

- Using Matlab in an HPC environment, 7 November 2017
- Transferring data to and from an HPC system, 28 November 2017
- An introduction to parallel programming using Message Passing with MPI, 24 - 26 October 2017
- Cluster architecture and job submission, 31 October 2017
- An introduction to UNIX/LINUX, 17 October 2017
- Cluster architecture and job submission, 11 May 2017
- Using Matlab in an HPC environment, 18 May 2017

Training and educational activities in 2018

- An introduction to parallel programming using Message Passing with MPI, 21 - 23 November 2018
- Making your research data fit for a future of open science and open data, 13 November 2018
- Using Matlab in an HPC environment, 7 November 2018
- Hierarchical modules and software selection, 23 October 2018
- Cluster architecture and job submission, 17 October 2018
- Transferring data to and from an HPC system, 2 October 2018
- An introduction to UNIX/LINUX, 18 September 2018
- An introduction to shared memory parallel programming using OpenMP, 29-30 May 2018
- Using Matlab in an HPC environment, 10 April 2018
- Visualisation and interactivity in HPC - The LUNARC HPC Desktop, 7 March 2019
- CodeRefinery Workshop, hosted by LUNARC, 15-17 May 2018
- Hierarchical modules and software selection, postponed
- Software Carpentry Workshop, 15-16 March 2018
- An introduction to UNIX/LINUX, 6 February 2018
- Cluster architecture and job submission, 13 February 2018

Training and educational activities in 2019

- Cluster architecture and job submission, 17 October 2019
- Introduction to data handling using R in a modern software environment, 13 -14 November 2019
- SNIC and SNAC: Handling of large-scale computer time allocations in Sweden, 30 September
- Debugging of HPC applications, 26 September 2019
- An introduction to UNIX/LINUX, 12 September 2019
- Artificial Intelligence Workshop with IBM, 21 March 2019
- Making your research data fit for a future of open science and open data, 25 April 2019
- An introduction to UNIX/LINUX, 7 February 2019
- Cluster architecture and job submission, 14 February 2019

- Transferring data to and from an HPC system, 21 February 2019
- Visualisation and interactivity in HPC - The LUNARC HPC Desktop, 7 March 20190

Application experts

Application experts are people with expert knowledge in a particular field. At LUNARC we have application experts in parallel programming, MATLAB, image processing, visualization and bioinformatics. The application experts are the link between researcher and the means of utilizing the resources at the center in an efficient and easy way.

CIPA Application expert (Anders Sjöström)

CIPA is a Lund University financed infrastructure providing access to researchers and application experts in image processing and analysis to serve users from Lund University as well as to the largest possible extent also users from elsewhere including industry and other private companies. Users can be found in four faculties: the Medical Faculty, the Faculty of Science, the joint Faculties of Humanities and Theology as well as the Faculty of Engineering. Anders provides expertise in HPC computing and MATLAB usage on large compute resources. Anders also provides support in writing resource applications for compute resources through SNIC.

LUBI Application expert (Gerard Schaafsma)

During 2018 and 2019 LUNARC received funding for an application expert in the field of bioinformatics, Gerard Schaafsma. He has been working for BMC as well as LUNARC helping users with scientific software installation and general support of these users. This initiative greatly increased the level of support that could provide to this research groups.

Special training events and workshops was also developed directed at bioinformatic users. These are listed in the training section of this report.

Parallel programming expert (Joachim Hein)

Joachim has been working as an application expert for SNIC for many years and has extensive knowledge in parallel programming techniques such as message passing (MPI) and multithreaded programming using OpenMP. He has been actively helping many of the national users (LU and other universities) larger users optimize their codes both for efficient running as well as for scaling them to larger European systems.

Other experts

IT architect (Monica Lassi)

Monica joined LUNARC in 2018 as an IT architect. She holds a Ph.D. in library and information science (2014), on a dissertation exploring the socio-technical design of an e-infrastructure to facilitate resource sharing and scientific collaboration.

Since 2017, she has worked for the upper echelons of the University Management team to coordinate the development of the research data strategy at Lund University. The research data strategy is founded in the research communities' current practices and needs for support in research data management. A fundamental part of the work with the strategy for research data management lie in facilitating meetings and communication between different types of actors, nationally as well as internationally. Much of the international work is done through the networks of NeIC and EOSC.

Further, Monica has taken active part in the Swedish Research Council funded e-infrastructure Swedish National Data Services (SND) since 2016: first as Lund University's representative in formulating the application to the Swedish Research Council and in setting the consortium organisation, and since 2018 as Lund University's representative in SND's steering committee.

During fall 2019, she contributed to writing the proposal for a new government agency for Swedish e-infrastructure for research (Inriktningsförslag för organisering av svensk e-infrastruktur för forskning), submitted to the government by the Swedish Research Council and SUHF, the Association of Swedish Higher Education Institutions.

Monica is a team member of the InfraESOC project EOSC-Nordic (2019-2022), in the work package FAIR data. She is the task leader of a task focussing on researchers' views on incentives and disicentives for making data FAIR. In this capacity, she is a member of the EOSC-5-FAIR task force for synchronising the five InfraEOSC projects' work on FAIR data. She has also contributed as a panel member on the EOSC-Nordic workshop on mapping open science policies in Copenhagen in February 2020.

Senior Developer / Project manager (Roger Larsson)

Roger Larsson has been working as a senior developer at LUNARC since 2016. Working on the development of L-SHIP, which is a background file transfer service. This service has been developed to offload the file transfers from sensor equipment and enable users to continue their analysis on LUNARC resources.

Roger is also responsible for designing a future infrastructure for providing services for long-term storage and other related services.

System experts

An invaluable asset to LUNARC that makes all of our resources run almost 24/7 is our team of system experts. They consist of a very knowledgeable team with years of experience in managing clusters and storage. The team consists of:

- **Tore Sundqvist** – System expert with a background from UPPMAX. Currently the main maintainer of the Aurora system.
- **Robert Grabowski** – System expert with an extensive knowledge in running grid services for the LHCK (Large Hadron Collider at CERN). Currently maintaining the Aurora-Grid system. He is also responsible for our virtual environment.
- **Marcos Acebes** – System expert / Application expert with an extensive knowledge in authentication and authorization. Maintains LUNARCs Authentication infrastructure. Marcos is also very active in the support queue as an application expert
- **Alex Contis** – System expert with extensive knowledge in monitorin- and backup infrastructures.

Data center procurement project

The LUNARC data center co-located at LDC. The current data center is old and doesn't not comply with the latest requirements with regards to energy efficiency and environmental considerations.

To be able to host modern computer equipment and reduce energy costs and environmental footprint, the current data center needed to be upgraded or an alternative had to be found.

Late 2018-2019 a project was initiated together with LDC to plan for a new data center. This work continues in 2020. Current planning is for an upgraded/new data center in 2022/23.

LUNARC within SNIC

LUNARC is one of the 10 SNIC centers in Sweden. SNIC funds and coordinates computational and storage resources among these centers.

SNIC Resources at LUNARC

From 2017-2019 LUNARC operated 2 HPC clusters for SNIC:

- Aurora – 180 node cluster
- Aurora-Grid – 20 node cluster for WLCG operations

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

See a detailed description of resources in the Resource section of this report.

Outreach

As all Lund University users of LUNARC are SNIC users, outreach efforts for LU users are also outreach efforts for SNIC users.

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.

SNIC 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.

SNAC WG

Within the Swedish national allocation committee working group (SNAC WG) Anders represents LUNARC in the allocation of computational resources to projects. Within SNAC WG, Anders has participated in the development of the allocation policies for SMALL, MEDIUM, and LARGE compute projects and also in the development of the policies of SMALL, MEDIUM, and LARGE storage projects.

MATLAB HPC mentors coordination

The MATLAB HPC mentors group was formed in June 2014 in an effort to coordinate the expertise available within SNIC and to disseminate the knowledge on the use of MATLAB in HPC throughout the SNIC organization. Anders Sjöström is the main coordinator for

this effort on the SNIC-side in Sweden. Anders has arranged meetings and miniconferences is Sweden as well as at the Super Computing (SC) conference in the US together with Mathworks.

NeIC

LUNARC is also collaborating around training in the Nordic countries through the NeIC initiative. NeIC projects that LUNARC has collaborated in has been

- Code Refinery - Training and e-Infrastructure for Research Software Development
- NeIC Dellingr – Resource sharing in the Nordics, see separate section.

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.
- Hardware and virtualization solutions for remote desktop architectures.
- Emerging storage technologies.

EOSC-Nordic

LUNARC participates for SNIC in the EOSC-Nordic EU project in "WP3: Support to EOSC-Nordic service providers (UT/ETAIS)" and "WP4 (FAIR Data)". Anders Sjöström is involved in WP3 and Monica Lassi is WP4 leader.

NeIC Dellingr

The NeIC project Dellinger focused on resource sharing across borders in the Nordic region. Anders has taken an active part in the project, investigating the legal and economic implications of sharing computational resources across nation borders in the Nordic region. The Dellingr project evolved into the Dellingr continuation project which focused on the sustainability and manageability of providing access to multiple resources located in multiple countries to users from several different countries across Europe. The Dellingr projects have produced two spin-off projects that Anders Sjöström is involved in or lead.

NeIC - Puhuri

The Puhuri project, which has been granted funding by NeIC late 2019, was born as an idea of facilitating ease-of-use of large computational resources within the EURO HPC joint undertaking. The original idea of presenting users with a common environment for access and use regardless of which resource is accessed was hatched by Anders Sjöström and has since evolved into a system for enabling prospective users of such systems to use their familiar national portals for access and project requests. Anders Sjöström has been

involved in writing the project proposal, as well as the detailed business case. The official Puhuri project start is in June 2020 and the project will run for at least two years, with a personnel budget of 9 FTE.

MAX IV International Pilot

The MAX IV international user's pilot is a test to discern how international users of the MAX IV facility can be given access to computational resources through SNIC and LUNARC. The pilot has evolved from being a part of the Dellingr project into a selfsustained project within SNIC. Providing easy access to computational resources to international users of the MAX IV facility. Both the pilot project and the current SNIC project are lead by Anders Sjöström.

Ongoing collaborations and projects

MAX IV

Since 2013, LUNARC has been actively collaborating with MAX IV to create an architecture for user and project management, data collection, data storage, computation (on-line and off-line) and visualization. In addition to a technical architecture, LUNARC has contributed to the design and implementation of MAX IV's on-line calculation system, which based on the same architecture as Aurora. The work on then initial architecture (SDM - Scientific Data Management) was reported in 140214. LUNARC and MAX IV have continuous team meetings at both operational and management level and work together to create an efficient workflow with seamless efficient transition to LUNARC's resources for off-line processing and data visualization.

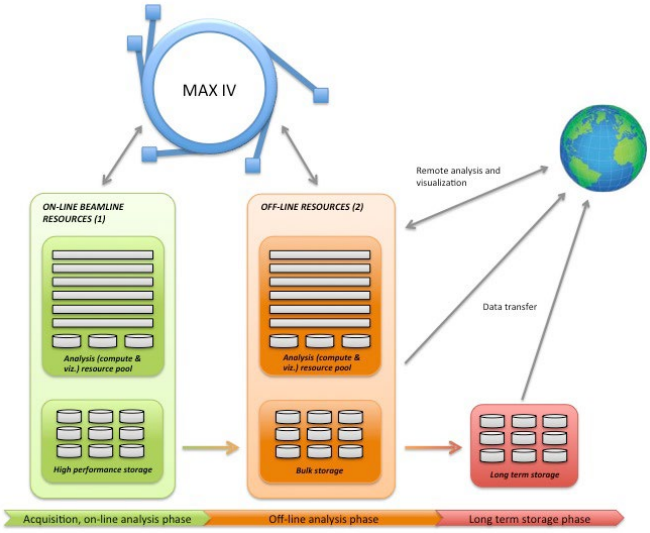

LBIC - Data Management

In 2011, LUNARC began a long-term collaboration with LBIC (Lund university BioImaging Center) regarding the construction and development of the center's infrastructure for

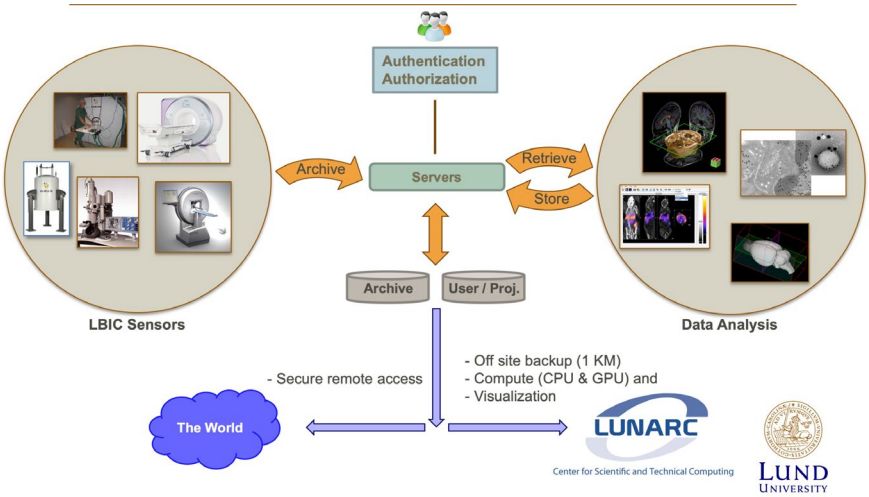

structured data collection, storage, data processing and visualization. A basic philosophy was a well-thought-out and authority-controlled structure from collected raw data to final processed results to avoid data loss and potential legal problems. All of the center's pre-clinical modalities (PET, SPECT, MRI, microscope, etc.) were included in the system. Cooperation with LBIC progressed even further. 2018 and successfully transitioned to project management of data management for 7T MR camera, a project that occupied most of the collaboration from the beginning. 2012. For more info see next item.

LBIC - 7T UHF MR Scanner

In the spring of 2011, discussions began with LBIC around a technical architecture for data management of DICOM images, raw data and associated information for sensitive (human) MRI data. The project application from VR for a 7T MR UHF camera was granted and in the spring of 2012 the practical work began on the design of a system for secure storage and handling of data from the camera that would be placed in the rebuilt and expanded MR department at SUS in Lund. Users could be affiliated with different univer-

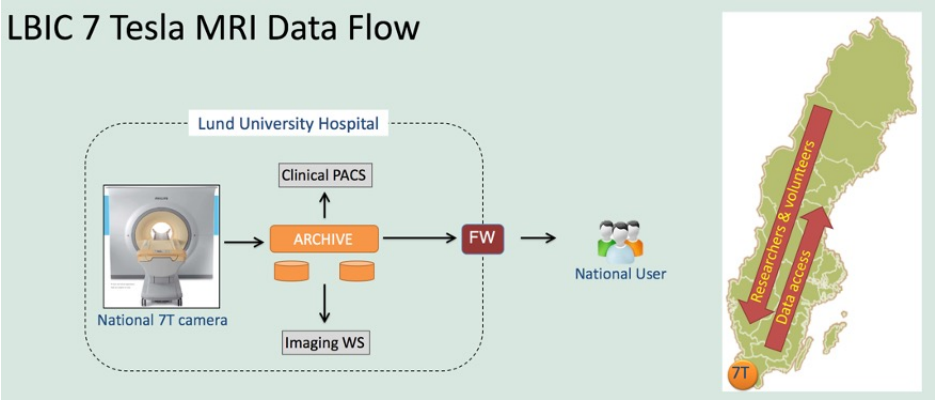

sities / university hospitals in Sweden and data could be downloaded by responsible researchers via a secure and legally approved solution. LUNARC lead the project and developed system architecture for over 4 years together with LBIC, Philips, Region Skåne, SUS and various subcontracted companies. The work was extremely complex, both organizationally and technically, as the camera was behind the hospital firewall and would handle human data. In addition to the research archive, the camera is also connected to the SUS clinical PACS and is used today for both research and to a lesser extent clinical investigation.

L-SENS

During the fall of 2016, LUNARC held initial technical discussions with a number of research groups in bioinformatics at Lund University regarding the possibilities of building a computational system approved for handling sensitive data, for example. human genome. LUNARC collaborated actively with the IT security officer for the university and a technical architecture was created and approved. The system was named L-SENS and was launched in the spring of 2017. This first incarnation of L-SENS (L-SENS generation 1) is static in that each PI and research group (assuming their

data is governed by the same legal requirements) works within their own secure calculation clusters with associated storage systems, isolated from the outside world, both at work and physically. L-SENS became popular and LUNARC currently has a total of 4 separate L-SENS generation 1 clusters used by research groups in bioinformatics. Many security audited special solutions have been implemented. One example of this is the direct connection of sequencing machines for automatic upload of sequenced through to LUNARC's secure storage system. LUNARC has a constant stream of requests from both individual researchers and groupings regarding access to the L-SENS system.

RS-SENS

RS-SENS, which has been in operation since 2019, is architecturally identical to L-SENS, the difference being that the system is logically located in the Region Skåne hospital network. RS-SENS was created because the region could not find competence in the area internally or at them outsources IT operation partners. Staff at Lund University Hospital had good experience with L-SENS from their research side and wanted the same service in the analysis of clinical data. The challenge with RS-SENS was to create a solution that could meet the region's strong security requirements when handling patient data. There are many similarities to the 7T MR project (see above) and the fact that LUNARC had very well projected the data management for the 7T camera meant that there was strong confidence in LUNARC as an organization.

LUBI-SENS

LUBI (Lund University Bioinformatic Infrastructure) was allocated infrastructure funds for the construction of a secure computational system similar to L-SENS. The system was commissioned by LUNARC in 2019. Unlike LU-SENS, unlike L-SENS, each has unique legal agreements governing the management of their data. This fact required that LUNARC had to create a new smaller monolithic system architecture that allows for higher granularity regarding secure isolation of users' data and computational processes. The architecture is based on a so-called hypervisors and virtual machines to create separate computational systems that all run within the same physical computer environment. The advantage is that one and the same cluster and storage system can be used safely by researchers with different legal requirements.

Humanities Lab (HUMLAB) and Archeology

For many years, LUNARC has had a continuous exchange of knowledge with HUMLAB in data management and data flows. The work has involved scientific workflow and structured handling of data from the laboratory's various types of advanced systems such as eye-trackers, EEG, articulography, motion capture and other sensor systems. In recent years, LUNARC has had an in-depth development project together with both HUMLAB and Archeology, which also involved adaptation of commercial program code. The project, which is still active at the time of writing, handles optimization of data flow and parallel calculation as well as visualization of photogrammetric data primarily from autonomous flying vehicles so-called UAVs or drones. The basic idea is that the flexibility and the minimum system requirement on the client computer for LUNARC HPC Desktop are used to upload collected data from the field, start massive parallel computation on

LUNARC's calculation clusters and visualize the results in real time. Since verification of collected data is important before an excavation area is left (then it is usually impossible to collect sensor data), this development project has greatly contributed to reducing lead times and enabling better and verified results in archaeological excavations. Of course, the system is also used from the office and other arbitrary locations as it relieves the local workstations and allows for processing of multiply larger data sets with minimized computation times and higher precision in the results. Both LUNARC's computational clusters Aurora and dedicated graphics clusters L-VIS with associated on-demand function have been used with great success.

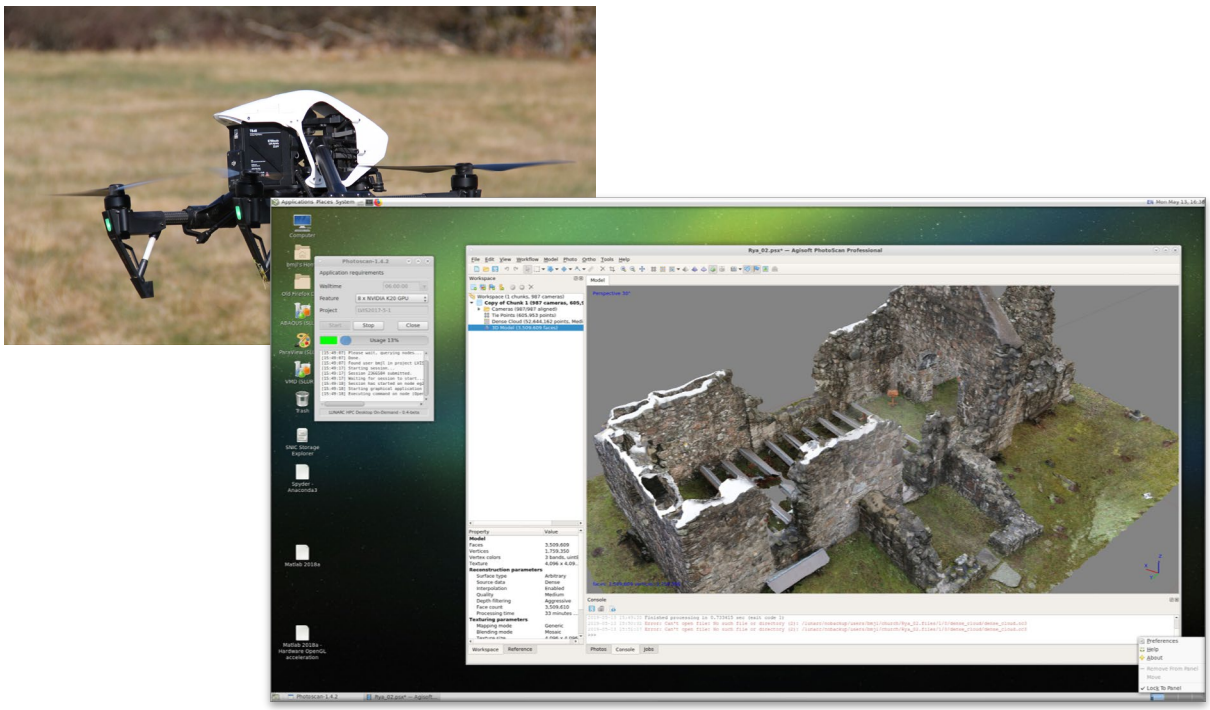

LUNARC L-SHIP

L-SHIP is a parallel data transfer system developed by LUNARC that focuses on performance, ease of use and flexibility. The idea for L-SHIP was born within the long-term collaboration with LBIC and the medical faculty at Lund University. The concept of HPC for LUNARC is not only massive parallel calculation but the whole process from data production, data storage and analysis. Focusing solely on the calculation part can be seen as a sub-optimization of a part of a larger scientific workflow. More specifically, L-SHIP was developed to increase user flow and thereby ROI (Return of Investment) for advanced and expensive laser-based microscopes. A microscope often consists of the instrument itself and an associated support computer for data collection, storage and analysis. The collection phase is typically substantially shorter than the analysis phase, but the instrument cannot be made available to the next user until ongoing data collection and analysis is completed. L-SHIP decouples the collection phase from the analysis phase and later parallels it via LUNARC HPC Desktop. As soon as the first user is done with their data collection, the collected microscopy data is sent via parallel processes efficiently to LUNARC's storage system and the analysis phase can be carried out via LUNARC HPC Desktop from

any location e.g. the researcher's workspace. The next user can immediately begin their data collection while at the same time a user analyzes their collected data remotely from the LUNARC's HPC Desktop. L-SHIP is, of course, applicable in various fields of science and faculties as well as for different sensor-based workflows.

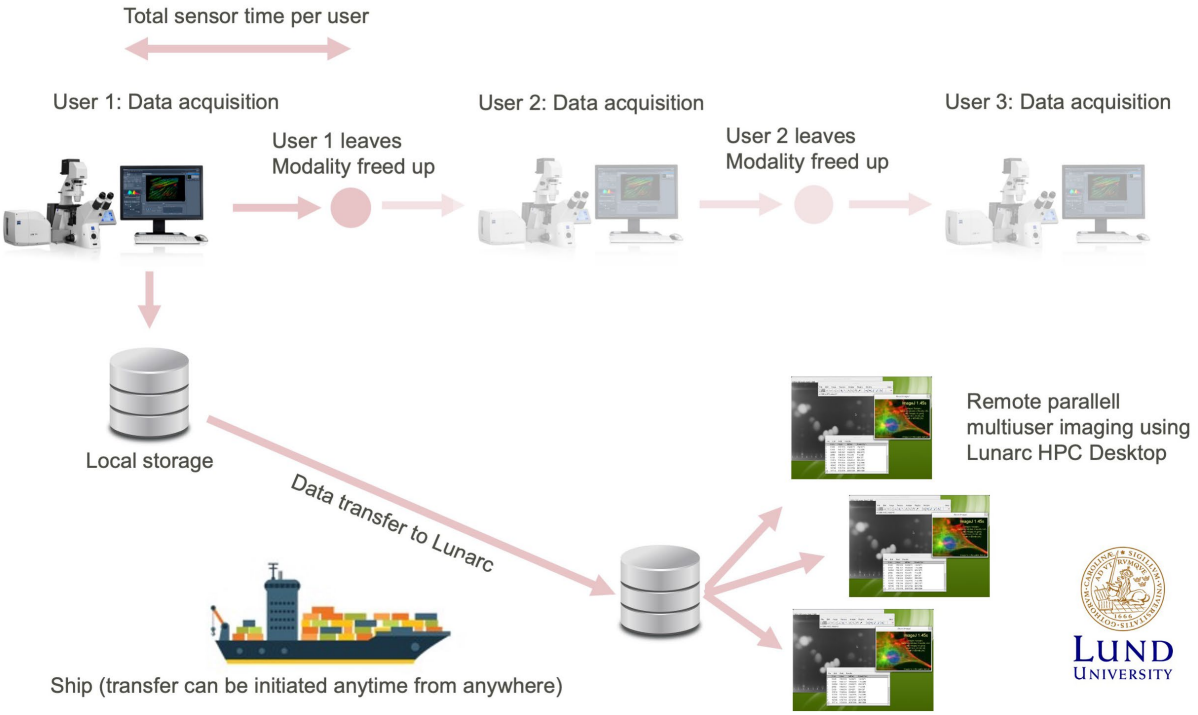

LUNARC HPC Desktop

In 2010, LUNARC began work on developing suitable techniques for a distributed visualization solution. The main principle is based on the fact that available network performance works against the purpose of an HPC center, namely to be able to handle very large data sets both in terms of synthetic and sensor-based data. The initial ideas and

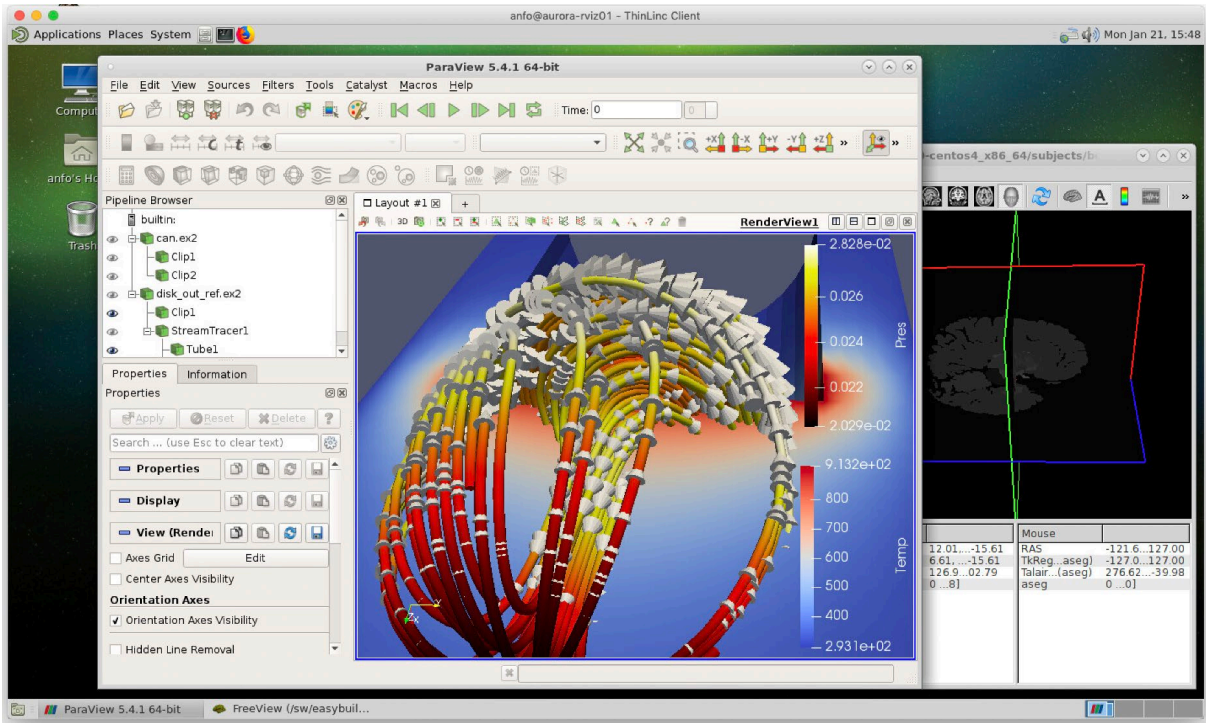

methods originated from industrial solutions developed by pioneers in computer graphics adapted to an academic setting. In the first few years, various hardware-based proprietary and high-performance point-to-point solutions were tested, which then gave way to an open source software solution. LUNARC's deep knowledge in computer graphics and visualization was used when the underlying scalable graphics system was developed and LUNARC HPC Desktop was launched to all users in 2013. The solution gave each user the opportunity to connect quickly and securely as an alternative to standard terminal login with SSH. from any location with a graphical Linux desktop. System management such as file transfer and file editing became much easier thanks to graphical tools and a whole new world of graphical applications became available. Graphic code development environments, code debuggers, MATLAB, with several programs were directly possible to run on LUNARC's HPC cluster. To meet the users' need for graphical post-processing over a large data set in a high-performance and scalable way, a technical architecture was created consisting of a back-end resource pool with graphic application servers. Demanding graphical codes such as ParaView and ABAQUS are executed on these specially configured nodes and the program windows are assembled on the desktop together with standard graphic desktops such as e.g. file manager for further distribution to the user. This entire background process is completely transparent to the user and all programs are available through the desktop menu system. LUNARC HPC Desktop is a success and has an ever-increasing user base.

LUNARC HPC Desktop On-Demand

The handling of demanding graphic applications in the original version of LUNARC HPC Desktop takes place on shared resources (see also the point above), which for some demanding user scenarios can be restrictive. Finite system resources such as primary memory, internal bandwidth, CPU and GPU cores are fixed and users are united on these resources within the graphics node. To solve this limitation, LUNARC developed HPC On-Demand which, uses the resource manager (SLURM) to allocate an entire graphics nodes

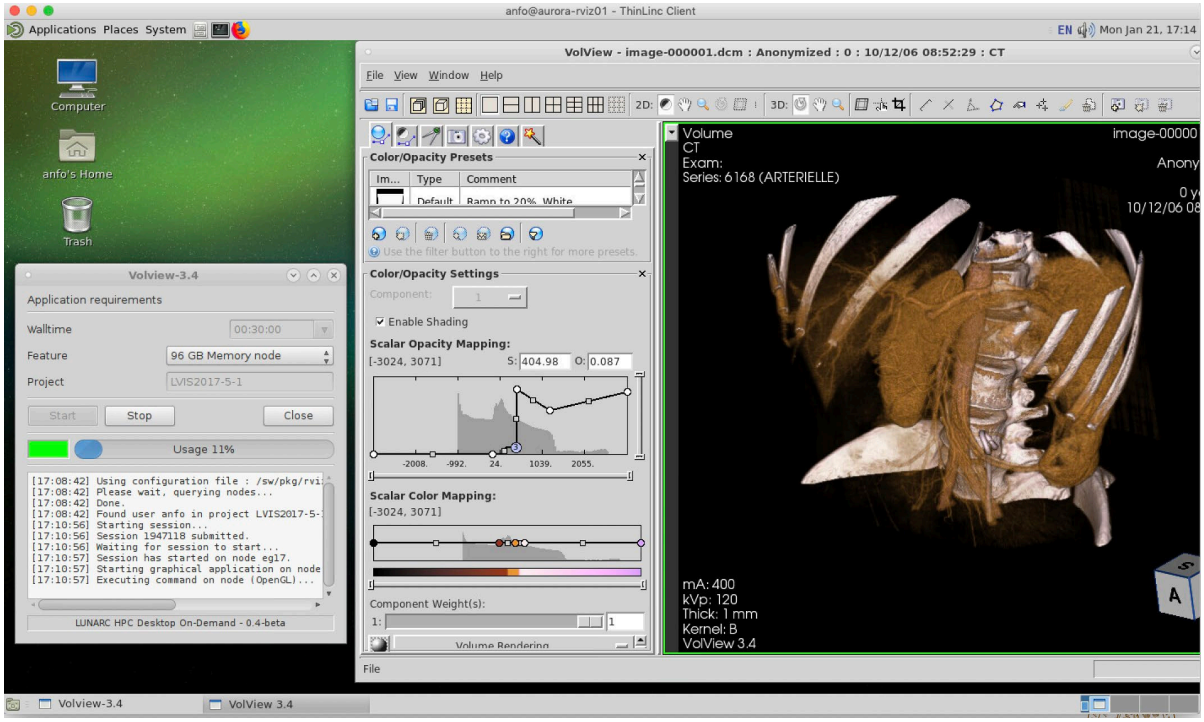

during a certain user-predetermined time for a specific application. All hardware resources in the node are available to the user throughout the allocated time. The resource is called L-VIS and consists of 24 nodes and a total of 68 graphics cards and is similar to the original version of HPC Desktop based on Linux. The system was implemented in production during 2018.

LUNARC HPC Desktop with Microsoft Windows integration

More user groups, for example in medical imaging, archeology, and material sciences, have expressed a need for other operating systems on the LUNARC HPC Desktop. To support this and create a seamless way of working, where a user can maintain their (often large) data sets on LUNARC's central storage system, LUNARC has developed an operating system independent graphical application backend. The system is part of LUNARC's existing infrastructure and benefits from the center's high internal network performance for connection to central storage and calculation clusters. Although the solution can handle any operating system, it is initially configured to provide support for Microsoft Windows 10. The hardware supports multiple users and different system configurations can be requested depending on the nature of the work. The system is equipped with a powerful graphics system which together with a powerful underlying server guarantees both 2D and 3D performance distributed to the users via LUNARC HPC Desktop. The resulted is a desktop solution where graphically demanding applications independent of operating system requirements can coexist in one and the same desktop solution.

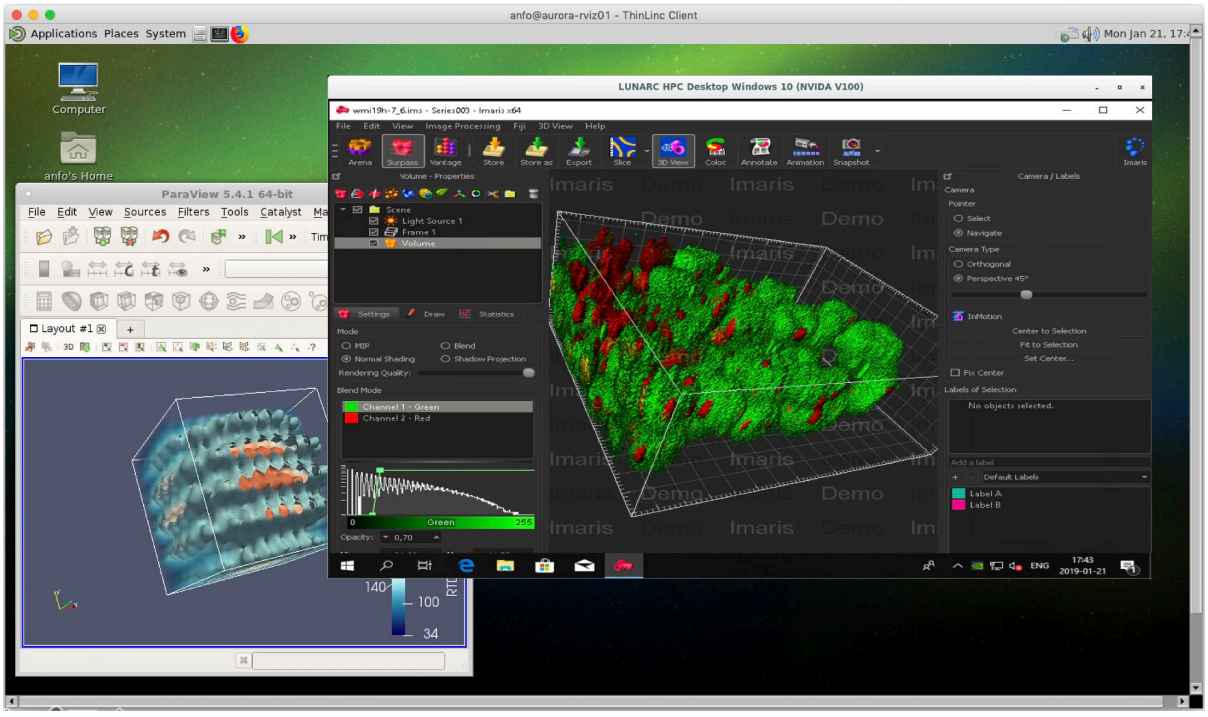

A user-friendly and secure user environment

The core service for the users of LUNARC resources is to be able to efficiently run large-scale computations over several computational servers with associated fast storage.

```
-------------------------------- /sw/Modules/modulefiles/Core --------
FEELnc/0.1.1                cellranger-atac/1.0.1       matlab/8.6
Protege/5.2.0               cellranger-atac/1.2.0 (D)   matlab/8.7
Singularity/default         cellranger/2.2.0            matlab/2017a
TIDDIT/20171003             cellranger/3.0.0            matlab/2017b
abaqus/V6R2017x             cellranger/3.0.2            matlab/2018a
allinea_forge/6.0.2         cellranger/3.1.0      (D)   matlab/2018b
allinea_forge/6.0.5         clion/2018.1.6              matlab/2019a
allinea_forge/6.1           comsol/5.2                  metashape/1.
allinea_forge/6.1.2         comsol/5.3                  ncftp/3.2.6
allinea_forge/7.0           comsol/5.4            (D)   netlogo/5.3.
allinea_forge/7.0.4   (D)   cplex/12.8                  nextflow/19.
```

When the Aurora computational resource was implemented, a lot of work was initiated to renew the tools that users frequently used on our systems. The entire software installation environment was reimplemented so that the support staff at LUNARC easily can provide optimized and well-functioning versions scientific software packages in a fraction of the time it usually takes to install these. The installed packages are then provided in a searchable system where the user can easily select the correct software package and associated compilers and libraries are selected automatically.

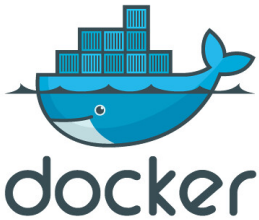

In addition to the software packages provided by LUNARC, support has been added to be able to run various software packages that are only provided as containers, which increases the opportunities for research groups to use the resources in a more flexible way.

A popular numerical environment is Jupyter notebooks. This is a web-based computing environment where you can combine report writing with Python code execution, producing well documented execution flows. To make it easy to use, this environment has been integrated into our desktop and you can easily start a notebook directly from the desktop menu.

LUNARC was also early in providing 2-factor login to the systems to make the systems more robust and less vulnerable to intrusion attempts. The system was introduced with SMS-based login, but from 2018 and 2019 login is accomplished with help of a smartphone app, which is perceived as easier by users.

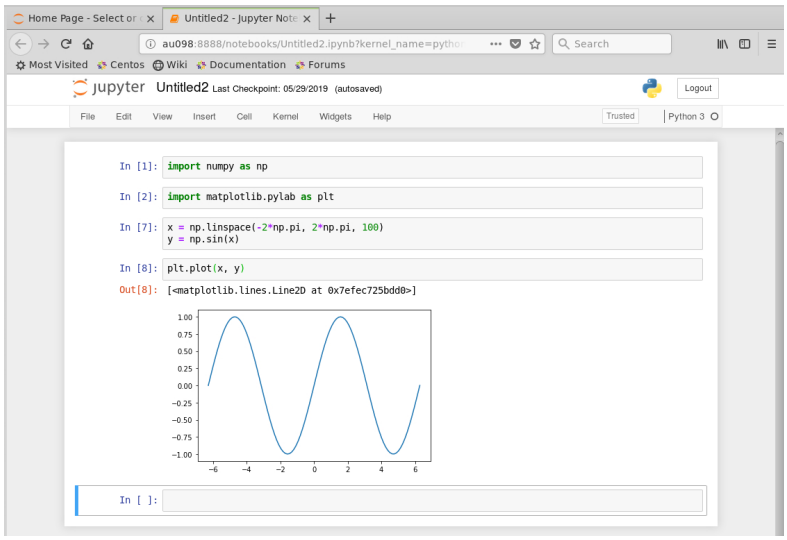

Figur 1 - Jupyter Notebook

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/projects 4 PB SpectrumScale parallel file system
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

L-VIS

L-VIS 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.

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

L-SENS

L-SENS is based on similar hardware as Aurora with separate secure storage systems.

L-SENS has been operational since 2017.

RS-SENS (Region Skåne)

RS-SENS is based on the same architecture as L-SENS. The resource has one login server, 8 compute nodes with 192 GB RAM and 144 TB storage server. The system is connected directly to the Region Skåne network.

Queueing system	SLURM
Home space	/home NFS mounted, available on all nodes
Linux distribution	CentOS 7.x x86_64 (RHEL7 compatible)
Software	Custom

LUBI-SENS (L-SENS Generation 2)

LUBI (Lund University Bioinformatic Infrastructure) was allocated infrastructure funds for the construction of a secure computational system similar to L-SENS. The system was commissioned by LUNARC in 2019. Unlike LU-SENS, unlike L-SENS, each has unique legal agreements governing the management of their data.

The system consists of 8 shared nodes connected to a secure file server and an virtual environment for providing unique and separate login nodes for each research project.

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 (SVS).

A new more efficient data center is being planned to handle needs for Lund University. The new center will be designed to be very environmentally friendly and reuse excess energy. By placing the data center in Lund, it can be placed close to high speed data links going to and from MAX IV and ESS, 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 at SVS will include not only include 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, interactive computing, visualisation and storage resources, enabling these resources for more users at LU.