Vetenskapsrådet

SURVEY OF E-INFRASTRUCTURE NEEDS FOR EIGHT LARGE INFRASTRUCTURES – REPORT FROM SNIC TO THE SWEDISH RESEARCH COUNCIL

Swedish Research Council

2015

SURVEY OF E-INFRASTRUCTURE NEEDS FOR EIGHT LARGE INFRASTRUCTURES – REPORT FROM SNIC TO THE SWEDISH RESEARCH COUNCIL

Swedish Research Council
Box 1035
SE-101 38 Stockholm, SWEDEN

ISBN 978-91-7307-288-5

SURVEY OF E-INFRASTRUCTURE NEEDS FOR EIGHT LARGE INFRASTRUCTURES – REPORT FROM SNIC TO THE SWEDISH RESEARCH COUNCIL

FOREWORD

The Swedish Research Council is a governmental agency which supports basic research of the highest scientific quality in all academic disciplines. One of the decision making bodies within the Swedish Research Council is the Council for Research Infrastructure (RFI), which has the overall responsibility to see to that Swedish scientists have access to high quality research infrastructure. RFI assesses the needs for research infrastructure in a regularly updated roadmap, launches calls for funding, monitors progress of infrastructures and participates in international collaborations.

Well-functioning e-infrastructures, such as digital communication, storage and computing capacity, together with human resources to aid in the usage of these infrastructures, are a prerequisite for most scientific disciplines today; both to support research projects and as a basis for other research infrastructures. The demand for supporting e-infrastructures is high and it is expected to increase further, both in terms of 'more of the same' and new services. This was clearly described through a broad RFI initiated effort, led by Professor Anders Ynnerman, to map existing and future scientific needs for e-infrastructures. The resulting report Science cases for e-infrastrcuture[1] in 2013, presents a diverse set of science cases and points to potential breakthroughs that can be made if sufficient supporting e-infrastructures are available.

Data generating research infrastructures have turned out to demand substantial supporting e-infrastructure, which in turn poses new challenges on the service providers. RFI invited The Swedish National Infrastructure for Computing (SNIC) to survey the expected requirements for e-infrastructure services from this particular user group. The resulting report, presented here, focuses on needs reported from eight large infrastructures. The report estimates costs for data operation and data handling and offers RFI and other stakeholders guidelines for strategic planning.

On behalf of RFI I thank SNIC and the eight large infrastructures for their dedicated work with this survey!

Stockholm 2015-06-01

Juni Palmgren
Secretary General
The Council for Research Infrastructures
The Swedish research Council

[1] Swedish science cases for e-infrastructure (2014) https://publikationer.vr.se/produkt/swedish-science-cases-for-e-infrastructure/

CONTENTS

SUMMARY AND RECOMMENDATIONS

This report includes an initial inventory of the large scale needs for compute and storage infrastructure by Swedish national research infrastructures. The report is the result of an instruction by the Council for Research Infrastructures[2] (RFI) to SNIC. This document addresses the presently known and foreseeable requirements for large scale e-Infrastructure resources by a number of research infrastructures as well as the services and resources that the e-Infrastructure community, in particular SNIC, can offer to the infrastructures at the national level.

For this inventory, SNIC invited thirteen research infrastructures to describe their needs for e-Infrastructure. Emphasis was put on research infrastructures that have expressed (to RFI or to SNIC) the need for large scale compute and storage infrastructure in the next five years. Eight research infrastructures were eventually included in this report.

SNIC welcomes feedback on this inventory from all the stakeholders, in particular RFI and the research infrastructures.

The inventory is not necessarily complete and contains resource estimates that are likely to change. SNIC therefore proposes that this inventory is updated at regular intervals, and at least once a year, with updated descriptions of requirements and roadmaps for implementation and with new research infrastructures.

The following conclusions and recommendations are made:

- Traditionally, funding for research projects and research infrastructures is often being applied for (and granted), while the e-Infrastructure requirements and their costs are determined only at a later stage. SNIC encourages that research infrastructures already in their proposal or preparatory phase are required to establish a tentative plan and roadmap for the required e-Infrastructure, including a draft budget. This plan should be updated at regular intervals. Such plan should be linked a research data management plan that addresses all stages of the data life cycle, including, collection, analysis, publication, archiving and re-use.
- SNIC encourages that the research infrastructures maintain these plans and roadmaps for the required e-Infrastructure in collaboration with SNIC and other national e-Infrastructures, e.g. SUNET. For this purpose, SNIC proposes to maintain an overview of the requirements for large scale e-Infrastructure that exist within the Swedish research infrastructures. This inventory that is the result of the instruction by RFI can form the start of this.
- A number of research infrastructures have stated that they prefer, have a clear need, or have a future need to make use of existing national e-Infrastructure. The relevant research infrastructures and SNIC should increasingly work together on piloting and prototyping new functionalities and services to help refine the definition of the e-Infrastructure requirements for the research infrastructures and corresponding roadmaps for their implementation.
- Several of the research infrastructures that are in operational phase today already make use of the national e-Infrastructure that is provided by SNIC. Some of the services are provided by SNIC on a best-effort basis, without well-defined guarantees and service levels. The research infrastructures and SNIC should explore establishing agreements in which SNIC provides access to adequate services and resources, in particular where this concerns time- or performance-critical services. Such agreements must include objectives, service levels, rights and obligations, conditions of use and cost-sharing for the provisioning and usage of such critical services.
- All stakeholders, including funding agencies, must continuously inspire the research communities and research infrastructures to optimize their use of the large national shared e-Infrastructure resources. In practice, the deployment of research infrastructures, in particular those that are in proposal or planning

---

[2] One of the decision making bodies within the Swedish Research Council.

phase, may take place rather independent of (national) e-Infrastructures that exist or are being deployed at the same time. Proper inter-reference and interoperation may be lacking, irrespective incidental co-operations. In some cases, it is assumed (or hoped) by the research infrastructure or researcher consortium that the e-Infrastructure needs will be catered for externally (e.g. by SNIC) and with separate funding. At the same time, large scale (national) e-Infrastructure is not a goal in itself, but must be an integrated part of all research infrastructures and all major research efforts that can make use of it. A key objective must be to achieve a seamless interoperation between the national e-Infrastructures and research infrastructures to provide common or harmonized services to the scientific communities, tailored to their needs where possible or required, and optimize the return on investments made by the e-Infrastructures, research infrastructures and user communities. Eventually, when such seamless interoperation is achieved, and, where appropriate, defined in agreements with well-defined guarantees and service levels, then research infrastructures and research efforts will acknowledge that national e-Infrastructure forms an integral part of their activity and, as important customers, express their support and need for the national e-Infrastructure, and provide incentives for alignment.

For 2015 and 2016, the needs for e-Infrastructure for the research infrastructures NGI and XFEL are considerable. The needs for NGI are particularly urgent to be implemented during 2014 or the beginning of 2015. RFI is encouraged to assess these needs and, should RFI decide to support these, allocate funding during 2014 so that the necessary e-Infrastructure can be established during 2015.

1 MOTIVATION FOR THIS DOCUMENT

An important activity of the Council for Research Infrastructures (RFI) from the Swedish Research Council is to produce a long-range strategic plan for how Swedish researchers within academia, the public sector and industry can get access to the most qualified research infrastructure in Sweden and in other countries. This plan is presented in The Swedish Research Council's Guide to Infrastructures. The guide serves as a roadmap for funding agencies regarding Sweden's long-term need for national research infrastructures and Sweden's participation in international research infrastructures. Earlier editions of the guide were published in 2006, 2007 and 2012. The fourth edition of the guide is planned to be completed in the second half of 2014.

The research infrastructures in the Guide to Infrastructures include major Swedish facilities and services and subsequent updates thereto, which are unique to Sweden for example due to their sheer size or cost involved in their establishment and their national-level interest. They encompass central or distributed research facilities, databases or large-scale computing, analysis and modelling resources. These resources often determine the opportunities to conduct cutting-edge research in most areas, and as they become ever more extensive and costly, it is necessary to develop infrastructures jointly in large cooperative ventures, regionally, nationally and internationally.

A critical component in all the research infrastructures in the Guide to Infrastructures consists of structured information systems for data management, enabling information, processing and communication to professionally and efficiently conduct science with or on the very facilities they are concerned with. These systems (commonly referred to as e-Infrastructure) include ICT-based infrastructures for computing, storage, communication and visualization of research data and a variety of middlewares (e.g. grids/clouds) to conduct science in a distributed setting. Nowadays, e-Infrastructure is an integrated component in almost all scientific workflows.

Today, RFI has only a partial overview of the needs for e-Infrastructure by the research infrastructures it supports (or considers to support), which complicates planning and budgeting by RFI. In addition, earlier in 2014, the RFI Council discussed the report 'Swedish sciences cases for e- infrastructure'[3]. In its examination of the report, the RFI Council noted that the needs of national and international research infrastructures for large-scale resources for computing and storage were not highlighted. As a follow-up to the report, RFI instructed SNIC to make an inventory of these needs and report to RFI which resources are required to meet the identified needs. This document is the result of that inventory. The instruction from RFI is given in Appendix A.

To some extent, the inventory is also meant to address a 'traditional' problem where research projects and research infrastructures are being applied for (and granted), while the e-Infrastructure requirements are determined only at a later stage. Moreover, the deployment of research infrastructures may take place rather independent of (national) e-Infrastructures that exist or are being deployed at the same time and proper inter-reference and interoperation is lacking, irrespective incidental co-operations. In some cases, it is assumed (or hoped) by the applicant that the e-Infrastructure needs will be catered for externally (e.g. by SNIC) and as such additional (from the Swedish Research Council or SNIC) resources and funds be allocated to the research infrastructure. In case the agreements with these parties are not established upfront, this may lead to a research infrastructure that has insufficient access to (or lack of) e-Infrastructure resources, budget problems, and eventually suffers from delays in production readiness or reduced output by and quality of the infrastructure.

The interest of SNIC in the instruction from RFI is therefore obvious. SNIC is a national e-Infrastructure provider of services for large scale computing (High Performance Computing (HPC)) and data storage and corresponding user and application support, all of which are made available through open procedures such that the best Swedish research is supported. This includes fair allocation and use of shared resources, as well as secure and distributed ways for computation and data storage. An important aim of SNIC is, where this is

[3] https://publikationer.vr.se/produkt/swedish-science-cases-for-e-infrastructure/

possible and appropriate, to tailor its services to the researchers' and research infrastructures' needs and to establish and strengthen links between SNIC and other national infrastructures to help provide the best services at the best conditions to Swedish flagship research facilities.

However, in situations where the demand for access to SNIC resources by far exceeds the available resources, the user of the SNIC infrastructure will perceive the services by SNIC as inadequate. For research infrastructures that must rely on access to SNIC resources, such situations are not acceptable. It is therefore important that SNIC and these research infrastructures define a framework for long-term collaboration that includes a high-level the roadmap for compute and storage infrastructure that is needed by the research infrastructures and expected from SNIC. During the past few years, SNIC initiated discussions and sharing information with other research infrastructures to describe their requirements for their use of e-Infrastructure, including both concrete and perceived requirements. The instruction from RFI led SNIC to invite a larger number of Swedish research infrastructures and as such provides important input to further evolve or initiate such collaborations. SNIC sees this report as a first iteration in a longer term dialog with other research infrastructures. This is further addressed in Section 6.

Finally, along with investments in new or next-generation research infrastructures, SNIC finds it important that investments in large scale e-Infrastructures resources are coordinated, thereby achieving efficiency in management and utilization of these resources and avoiding that parallel solutions are set up without connection to existing or planned investments. A key objective must be to achieve a seamless interoperation between the national e-Infrastructures and research infrastructures to provide common or harmonized services to the scientific communities, tailored to their needs where possible or required, and optimize the return on investments by the research infrastructures and user communities.

2 REQUIREMENTS GATHERING

In the past ten years, RFI has awarded grants for the preparation, implementation and operation of a range of large national infrastructures and for the Swedish participation in large international infrastructures that have, or will have, large needs for e-Infrastructure. For this inventory, SNIC contacted thirteen of these research infrastructures to describe their needs for e-Infrastructure.

Emphasis was put on research infrastructures that have expressed (to RFI or SNIC) the need for large scale compute and storage infrastructure in the coming five years. A rigid definition of what is meant by ‘large’ was not made, but can for example mean storage/data in the order of 1 Peta Byte or more, computation in the order of 1 million core hours per month or more, or significant personnel effort that is needed for deploying and maintaining the e-Infrastructure (e.g. at least 2 FTEs).

The needs for e-Infrastructure for eight research infrastructures are eventually included in this report. Some of the other research infrastructures that were invited reported that they do not currently classify as research infrastructures with large scale needs for computation and storage. This may change in the coming years, but it is for most unclear when and to what extent this will happen. These may be included in future versions of the report.

Research infrastructures exist in different forms and sizes and their needs for e-Infrastructure solutions are usually equally different. However, the need for e-Infrastructure, and in particular computation and storage, typically comes from two directions:

A. In part directly from the infrastructure operations (‘production part’). This includes for example the production and storage of primary/raw data, pre-processing, and delivery of data products.
B. In part from the infrastructure users (‘research part’). This includes for example data analysis and simulation made by researchers.

A typical scenario is that a research infrastructure needs access to computing and storage services to process and store the raw data from the experiments (e.g. sequencing or beam time) before they can be delivered to the user or made openly accessible. A research infrastructure may have the obligation to keep the experimental data in some form for a number of years, in particular when it concerns data that cannot (easily) be regenerated. The researcher subsequently also has a need for computation and storage in his/her work to analyze the data from these experiments and share the research results and corresponding (raw and derived) data sets with colleagues.

Where possible, concrete numbers for resource estimates are included in the descriptions, in particular for the coming three years (2015-2017) and where possible up to 2019.

For each research infrastructure, the following information is provided:

1) Scientific discipline(s)
2) Coordinator(s)
3) Participating institutions
4) A short description of the research infrastructure with emphasis on those aspects that are relevant for describing the needs for e-Infrastructure, e.g. distribution of facilities, data generators, data volumes and data complexity
5) References to recent contracts or other relevant documents from funding agencies by which the infrastructure is supported (or referred to).
6) A description of the e-Infrastructure requirements. Where possible this description takes into account the requirements for both the production part and for the research part. Where appropriate, the description includes the work/dataflows between the components/nodes in the infrastructure and the data sets and products that are collected and/or produced by the infrastructure. e-Infrastructure capacities and services are described, in particular:
    a) Computing services, e.g. general purpose computing (x86-based) and specific computing resources (GPU, etc.)

b) Data services, including
    - Types of storage: Long-term, persistent storage and archiving; project, temporary storage; disk, tape, etc. Where known, also describe specific storage requirements (e.g. latency, bandwidth)
    - Need for data publishing, databases, access, retrieval, etc.
    - Need for curation, annotation, metadata services
    - Other

c) Support, application expertise, training

d) Network services

e) Other

7) Roadmap for implementation, including a brief time schedule for the implementation of the e-Infrastructure (2015-2020) and a specification of the required resources, total per year, both for the production part and research part of the infrastructure. Concrete resource specifications (e.g. CPU hours) use 2014 performance characteristics and are not scaled with predicted technology performance improvements in the years 2015-2020. Where appropriate, a justification for the specified resources is provided.

8) The description also details whether it must handle sensitive personal data and/or projects that require ethical approval.

Network requirements are also included as data communication is an integral part of e-Infrastructure. Network infrastructure to interconnect research and education communities in Sweden is the responsibility of SUNET.

3 OVERALL ASSESSMENT

The e-Infrastructure requirements for the individual research infrastructures are given in Section 4. In this section, we provide some observations and an initial analysis of requirements, as well as some cost estimates for the research infrastructures that see a need for considerable investments in e-Infrastructure during 2015 and 2016.

3.1 Observations

The following general observations are made:

- Research infrastructures come in different forms and sizes, which makes it hard to present them in a unified manner. The facilities and services that constitute the infrastructure may be centralized in a single organization (e.g. MAX IV), but can also be distributed across a larger number of organizations across Sweden (e.g. BILS). The chosen operational structures for the infrastructure are usually influenced by the location(s) of the experimental facilities and by the location(s) where the experimental and derived data are made available for further analysis and modelling.
- Research infrastructures are at different phases of their deployment, varying from a preparatory/design phase (e.g. EISCAT_3D) to an operational phase (e.g. NGI, WLCG). The matureness of the specification of the e-Infrastructure requirements varies accordingly.
- The research infrastructure's role in the production part of the infrastructure is usually clear. However, its role in the research step for the scientific analysis of data as well as long-term curation is sometimes less clear or still under discussion (e.g. MAX IV and EISCAT_3D).
- In addition, the ownership of the experiment data and responsibilities for the long-term curation of the data differ between the infrastructures. Some infrastructures keep all or most of the experiment data, and must cater for this accordingly (e.g. WLCG, EISCAT_3D), while other infrastructures hand the experimental data over to the researcher (e.g. MAX IV, NGI, XFEL). In the latter case, it is not always clear that the researcher has access to the necessary e-Infrastructure for further scientific analysis of the data.
- It is often not trivial for research infrastructures to predict their e-Infrastructure requirements accurately for three or more years in the future (i.e. beyond 2017). Especially for research infrastructures that are in a preparatory or design phase, e-Infrastructure requirements and corresponding roadmaps for implementation may only be approximations that are based on a number of assumptions, e.g. regarding time schedule for implementation, available funding, number of experiments (e.g. beam time or samples), projected data volumes, data compression, etc.
- In case of participation in international infrastructure that is still in a preparatory or design phase, it is possible that e-Infrastructure requirements are only known or estimated for the total infrastructure and not for Sweden only (e.g. EISCAT_3D). For other infrastructures, the Swedish resource requirements may be derived from international specifications (e.g. WLCG).

Regarding SNIC's role, there seems consensus among the research infrastructures that SNIC is a suitable infrastructure to take care of the research part. This way, the required resources are part of a larger shared SNIC infrastructure. This will make it for example easier to cope with peak demands from a research community and avoid that parallel solutions are set up. In return, SNIC can increase its efficiency in utilization, cost, support and management of its resources. This partnership requires a close interaction to make sure that the research and production parts are interoperable by aligning the necessary tools, for example for authentication, data transfer and scientific analysis.

SNIC could also have a role in the production part of the research infrastructure in case it exhibits peak demands for resources or in case the workloads can be mapped onto the general-purpose computing and storage resources (e.g. as for WLCG) or in case the production and research parts must be closely coupled in some way. Specialized resources and system configurations that are suitable for only a single research infrastructure and that must be dedicated to take care of specific workloads may be set up separately.

3.2 Initial analysis of requirements

Some initial requirements analysis and the challenges and opportunities for e-infrastructures are outlined. This includes technical and functional aspects, but also policy aspects, such as resource allocation, governance, and cost-sharing that should be addressed to ensure long-term sustainability.

Several of the research infrastructures have stated that they prefer, or have a clear need, to make use of existing national e-infrastructure. Therefore, improving the interoperability between these structures will have an added value for all the user communities. Several areas can be identified where the e-Infrastructures and research infrastructures can work in common:[4]

- Resource access:
  - Consistent identity management is a fundamental requirement. All research infrastructures must use Authentication and Authorization Infrastructures (AAI). These are typically separately managed and not identical in technology. A unified single sign-on service can ensure that an individual's identity can be used across network, compute and data services and across infrastructures. Harmonizing policies for authentication, authorization, and potentially accounting and auditing, will simplify access to underlying e-Infrastructures. Deployed AAI systems must interoperate so that a user's identity can be established once and accepted by the e- Infrastructures and research infrastructures.
  - Access control to resources, data and applications on a community level is necessary for a subset of the research infrastructures and their user communities. Different research infrastructures may require different levels of granularity and use different semantics. Hence, a goal would be to offer consistent support for research infrastructures by harmonizing current features for access control.
- User support (i.e. procedures for user query handling and dealing with perceived performance issues) and training. To effectively use e-Infrastructure, the users of a research infrastructure should have quick responses to their queries and access to high-quality documentation. All national research infrastructures must offer specific user support services that may have to cooperate such that requests can be issued to appropriate support groups across different technologies and geographical regions. The users of research infrastructures may have a need for training, education or external expertise in their use of e-Infrastructure. In this respect, training may have most impact when tailored to the specific needs of the target user community.
- Research data management practices ensure that research infrastructures and researchers are able to meet their obligations to funders, improve the efficiency of research, and make data available for sharing, validation and re-use. To support this, it is imperative that data management is done properly through all stages of collection, analysis, publication, archiving and re-use. Data management includes a range of things, including:
  - The ability to provide long-term storage and accessibility (possibly measured in decades rather than years), identified as important by several research infrastructures. Persistent Identifiers (PIDs) and metadata are key issues. There exist services for registering, storing and resolving digital object identifiers (DOIs).

    Effective access to persistent data from a national e-Infrastructure has several implications, including for example guarantees of quality of service and access for long-term storage will be required for centers offering persistent data, seamless authorized user access to data across the infrastructures, the middlewares deployed by the e-Infrastructures must support access to persistent data using PIDs, and provenance of data allowing the origins of data to be recorded and traced and its movement between databases and research infrastructures.

---

[4] See also deliverables from e-Infrastructure Reflection Group (www.e-irg.eu) and European e-Infrastructure Forum (www.einfrastructure-forum.eu).

- Planning is also an integral part of the data management process. In environments where the responsibility for research data management lies with the researcher, research infrastructures should have comprehensive data management policies and procedures to support their researchers. The research infrastructures can cooperate on formulating and adopting guidelines to safeguard data, to ensure high quality and to guide reliable management of data for the future without requiring the implementation of new standards, regulations or high costs.[5] Such guidelines are of interest to organizations that produce data, organizations that archive data, and to the consumers of data.

- Sensitive data. A number of research infrastructures expressed the importance of e- Infrastructure to handle sensitive personal data and/or projects that require ethical approval. These require secure analysis and storage with clear rules and routines for access to data that follow legislation. See Section 3.2.
- Security incident handling. All the e-Infrastructures and research infrastructures must have dedicated security structures, procedures and measures to ensure the secure operation of the infrastructures. The security incidence response groups in the various infrastructures must cooperate. Such cooperation is already in place between national e-Infrastructure partners and international e-Infrastructure initiatives (e.g. through the European Grid Initiative (EGI)), but must be generalized to ensure an effective and timely response to security threats that can exist across the whole research infrastructure ecosystem.
- Workflow support. The different research infrastructures use many different workflow tools or frameworks employed by the user communities and this diversity is certain to remain. This requires that the underlying e-Infrastructure supports these workflow tools and environments. Cross infrastructure workflows requires that the AAI, access control and various data management interoperation aspects must be in place.
- Integration. Research infrastructures are typically not concerned about how and where computing and storage resources are provided, but are primarily interested in easy-to-use, powerful and secure facilities and services. Externally operated resources can be a good solution for a user who needs additional resources on demand, e.g. commercial clouds. However, the tasks of research infrastructures' users may require computation that is for example only possible on sophisticated high performance computational resources that are not provided through clouds. Similarly, there are still important policy questions to be addressed concerning large-scale data management and archiving on commercial services. Each computing paradigm has its particular use, advantages and drawbacks, and eventually, a custom fit solution for each user community is preferred. Proper use of standards must make it easier to bring cloud, grid and supercomputer services together and define interfaces designed to simplify and promote their interoperability.

Swedish research infrastructures need to collaborate with parties beyond national borders. The national e-Infrastructures must increasingly leverage their existing international contacts for the benefit of the research infrastructures. e-Infrastructures from different countries must collaborate to provide pan-European or global e-Infrastructure that delivers uniform serves to an international user community. As such, the participation of SNIC in international e-Infrastructure initiatives must be evaluated on the value it has to Swedish research infrastructures and researchers. Where possible, such value must already be assessed at the proposal or planning stage of international collaborations.

3.3 Sensitive personal data

The research infrastructures MAX IV, NGI, BILS and Swedish Bioimaging expressed the importance of e-Infrastructure to handle sensitive personal data and/or projects that require ethical approval.

Legislations that govern the handling and access of personal data are the Personal Data Act[6] (PUL) and the Ethical Review Act[7]. Also to be considered is the Public Access to Information and Secrecy Act[8]. The Swedish

---

[5] E.g. Data Seal of Approval, http://datasealofapproval.org/

[6] SFS nr 1998:204, Personuppgiftslagen - PUL

Data Inspection (Datainspektionen) has produced together with the National Board of Health and Welfare (Socialstyrelsen), Statistics Sweden (SCB) and the Central Ethical Review Board (Centrala etikprövningsnämnden) a booklet on the rules and legislations that govern the use of personal data in research.[9] Personal data that contains information about for example health is to be considered sensitive, and thus human genome data falls in this category. It should be noted that coded and encrypted data is still personal data as long as there exist keys, even if the keys are not available to the researcher.

To handle personal data according to legislation puts certain requirements on the IT-systems and the operating procedures. Universities and other government agencies have their own written guidelines for information security and IT security that should be followed. SNIC and its partner centers do not currently have such resources or routines. Up to now, analysis with sensitive personal data has been handled by local (small) IT-systems at the researchers departments or home institutions. However, with increasing amounts of data (as generated by next generation sequencing), databases and software tools, the need for a national e-Infrastructure for sensitive personal data has increased. Therefore it makes sense that one or more of the SNIC centers would provide a secure environment for computation and storage of sensitive data. In this respect, the SNIC center UPPMAX in Uppsala has established a pilot project for sensitive data in the life sciences together with BILS.

To have users from several universities and research institutes working in different research groups analyzing data on a national e-Infrastructure requires routines and procedures on a wide scale that currently does not exist. Probably, users also need to be educated to avoid "leaking" of personal and genomic data or to mitigate the consequences of such leak.[10]

Early 2014, The Science for Life laboratory (SciLifeLab) initiated the national Swedish Genome Program. The program has two parts: (i) whole genome sequencing to identify the genetic causes of diseases of high health relevance and (ii) the establishment of a reference database of genetic variation in the Swedish population based on whole genome sequencing. Both parts will generate sensitive personal data.

Hence, requests for compute and storage resources can be expected already during the autumn of 2014 and early 2015. Therefore, there is an urgent need to supply infrastructure for handling sensitive data as soon as possible.

3.4 Cost estimates

This section includes some cost estimates for the research infrastructures that are described in Section 4. Since it is non-trivial for research infrastructure to predict their e-Infrastructure requirements for three or more years into the future, we restrict the discussion to those infrastructures that see a need for considerable investments during 2015 and 2016.

We make the following observations for 2015-2016:

- The needs for e-Infrastructure for the research infrastructures NGI and XFEL are considerable. The needs for NGI are particularly urgent to be implemented during 2015.
- EISCAT_3D has no hardware requirements for 2015-2016. The required personnel during 2015-2016 is 1-2 FTE to further design and prototype the e-Infrastructure. This effort is shared between the participating Nordic countries.
- The e-Infrastructure needs reported by BILS are essentially those for NGI for 2015-2016.
- The hardware requirements for MAX IV, Swedish Bioimaging and Onsala Space Observatory for 2015-2016 are modest. The requirements for support personnel can however be considerable, for example to prototype, build and operate the e-Infrastructure.

---

[7] SFS nr 2003:460

[8] SFS nr 2009:400

[9] "Personuppgifter i forskningen – vilka regler gäller?" March 2013. www.epn.se, www.datainspektionen.se, www.scb.se, www.socialstyrelsen.se

[10] Steve E. Brenner. "Be prepared for the big genome leak." Nature 498, 139. June 2013. doi:10.1038/498139a

- The e-Infrastructure needs for WLCG will be handled separately. SNIC will propose a detailed budget for 2015-2016 to RFI in September 2014.

Cost estimates for 2017 and later years can be made (and refined) in the coming years and funding to implement the requirements can be secured during 2015 and 2016.

To get an idea of the level of funding that is required to implement the e-Infrastructure needs for NGI and XFEL, we use unit costs for the various resource types (in 1000 SEK). These are given in the table below. By multiplying the resource specifications of the research infrastructures (given in Section 4) with these unit costs, one gets an estimate for the total funding that is required. It is emphasized that these unit costs are rough estimates, especially for 2017-2019.

Unit	**2015**	**2016**	**2017**	**2018**	**2019**
CPU: one million core hours (1)	51.0	35.8	25.0	17.5	12.3
GPU: GK110 equivalent (1)	10.0	7.0	4.9	3.4	2.4
Storage: one Petabyte / year (2)	250	200	160	128	102.5
Network: 10 Gbit/s link (3)	300	300	200	200	200
Network: 100 Gbit/s link (3)			600	600	600
Support: one FTE (4)	1 000	1 030	1 061	1 093	1 126

(1) Based on cost for hardware that is operated for a 4-year period. The cost decreases 30% per year. Operational costs (floor space, cooling) are not included;

(2) Cost for storage hardware that is operated for a 4- year period. The cost decreases 20% per year. Operational costs (floor space, cooling) are not included;

(3) SUNET service, annual fee;

(4) Cost per FTE, 3% increase per year.

3.4.1 NGI

The total requirements for e-Infrastructure for NGI are given in Section 4.3. Using the unit costs from the above table, we get the following cost estimates for the production and research parts of the e- Infrastructure for NGI for 2015-2019 (in 1000 SEK):

	2015	**2016**	**2017**	**2018**	**2019**	**Sum**
Production*	5 125	7 335	7 526	7 913	7 870	35 769
Research	13 950	15 963	14 644	13 361	12 406	70 325
Required funding	15 450	17 508	16 236	15 001	14 096	78 291

* excluding the 19 FTE support staff for the bioinformatics platforms.

The cost estimates include hardware (computing, storage, networking) and personnel, but not local infrastructure expenses for example for floor space, electricity and cooling.

The Knut and Alice Wallenberg Foundation is considering a substantial increase in its funding in the field of human genome sequencing – reagents, instruments and production e-Infrastructure (hardware only). This increase is already reflected in the estimates provided in Section 4.3. This funding is likely to be decided during 2014. In such case, the establishment of the research part of the e-Infrastructure will be urgent. For the above table, the resource requirements have been included for both the production and research parts. The estimated required funding includes only 1.5 FTE for supporting the production e-Infrastructure, and all the resources for the research part of the e-Infrastructure.

Both the production and research part of the e-Infrastructure must be able to handle sensitive data and therefore the workloads cannot be mapped onto existing SNIC production resources.

3.4.2 XFEL

The total requirements for e-Infrastructure for XFEL are given in Section 4.2. The following cost estimates are given for the (combined) production and research e-Infrastructure for XFEL for 2015- 2019 (in 1000 SEK):

	2015	2016	2017	2018	2019	Sum
Production & Research	9 247	13 059	16 468	7 342	7 750	53 866
Required funding	9 247	13 059	16 468	7 342	7 750	53 866

The cost estimates include hardware (computing, storage, networking) and personnel, but not local infrastructure expenses for example for floor space, electricity and cooling.

4 RESEARCH INFRASTRUCTURES AND THEIR REQUIREMENTS

This section describes the e-Infrastructure requirements for eight Swedish research infrastructures. The descriptions use the structure that is outlined in Section 2 and, where appropriate, distinguish between the production part and research part of the required e-Infrastructure.

4.1 MAX IV Laboratory

The MAX IV Laboratory is a national facility hosted by Lund University that operates accelerators producing X-rays of very high intensity and quality. The MAX IV Laboratory is the successor of the MAX-lab and includes both the operation of the present MAX I, II, III facilities (MAX-lab) and the MAX IV project aiming at constructing the new MAX IV facility at Brunnshög in the North Eastern part of Lund. The MAX IV source will be the most brilliant synchrotron light source in the world and will by far exceed the performance of other third generation synchrotron radiation facilities.

1) Scientific disciplines.

Physics, Engineering, Materials Science, Chemistry, Biology, Medical Sciences, Environmental Science, Cultural Heritage and Archaeology

2) Coordinators.

Christoph Quitmann (Director), Tomas Lundqvist (Life Science Director), Krister Larsson (IT-strategy)

3) Participating institutions.

Lund University (host), Uppsala University, Linköping University, Gothenburg University, Umeå University, Luleå University, Karlstad University, Stockholm University, KTH Royal Institute of Technology, Karolinska Institutet, Chalmers University of Technology, Swedish University of Agricultural Sciences

Vinnova, Region Skåne, Vetenskapsrådet, Region Skåne, Finnish Academy of Sciences, Tartu University (Estonia), Danish Technical University, Copenhagen University and Århus University.

4) Short description of the Research Infrastructure.

The data generators at the MAX IV facility are located at the experimental stations, called beam lines. The new facility will accommodate 26 beam lines when fully developed (2026). Presently, 13 beam lines are funded and will become operational in 2016 and 2017.

Each beam line will have its own setup, adopted for a particular experimental technique, which includes instrumentation for optics and sample positioning and environment control. Each beam line also includes detectors for the data acquisition, which have different data formats, peak rates and volumes as required from the relevant experimental technique. It has been estimated that MAX IV will have 2000+ visiting users collecting data per year and it is a top priority to make sure that the data is of the best possible quality and in a format allowing non-expert scientists to extract meaningful information. Data will be stored at MAX IV Laboratory for a limited period and then transferred off site either to the user's home institute or to central archiving/storage facilities.

Getting data transferred off-site requires organizing the data storage for all experiments and knowing *who* collected *what* and *when* in order to keep track of it. Access restriction needs to be applied allowing access only to the principal investigator and scientific collaborators identified by him/her. Transfer of data should include a user-friendly download service (i.e. web browser based, one-click) from MAX IV storage services as well as an automated transfer service of data, retaining the ownership and location information. The latter will be of particular importance for scientific communities that collaborate and have developed their own analysis platforms.

Although the role for MAX IV in the long-term curation of scientific data is still unclear, it is important to make provision for having control of what data should be transferred to an archive or deleted and this requires an organized approach.

Some of the beam lines will need to pre-process and reduce data during the experiment, before further analysis can take place. The reasons for this include: compression, evaluation that the collected data is suitable and valid for the scientific investigation, correcting and calibration of raw data, combining multiple data files into one data set and file format conversion. Some beam line data collection techniques, such as imaging, will require advanced and complicated analysis to be performed at MAX IV, which must therefore provide resources in terms of compute power, software and staff competent in using the complex software packages involved.

The recent developments in detector technology have increased the data volumes, with peak rates of ≈1 GB/sec and 40 TB/day, to a point where every component in the data flow from detector to storage has to be fine-tuned. This has made it increasingly difficult and costly to transfer data across the network for online analysis at a different computing facility. The data usually benefits from being pre-processed where it is initially stored, with a high-speed connection (i.e. Infiniband) between file storage and the computing nodes.

It is desirable that MAX IV provides a more complete analysis service, with tools for the final analysis of pre-processed data, however this is unclear since it depends on local resources and knowledge being available and also the working practices of the visiting users. Although many types of experiments will have a well-developed data reduction step, some experiments will still have large data volumes after the pre-processing making the data download step time consuming and subsequent storage expensive.

The above needs require a well working e-Infrastructure where data can be transferred, stored and analyzed. MAX IV is currently investigating how this should be accomplished offering suitable and cost-effective services.

5) References to recent contracts or other relevant documents from funding agencies.

- MAX IV original funding agreement (Avsiktsförklaring avseende etablering av MAX IV, Lund Universitet 2009-04-27, Dnr LS 2009/431)
- MAX IV Strategy Report to the Swedish Research Council 2012
- Operation funding decision Dec 19, 2013 (Swedish Research Council Dnr 2013-2235)
- The Swedish Research Council´s guide to infrastructures (2012)
- Swedish Science Cases for e-Infrastructure (2014)

6) Description of e-Infrastructure requirements.

A. Production requirements

Data from the detectors will be stored on the central storage system at the MAX IV Laboratory, where the goal is to save data without local storage on each beam line. In the cases where this is not possible, i.e. due to detector setups, a local disk will be used on the beam line for caching including pre-processing where required. Data will then be transferred to a common disk array from the cache. This central storage system will consist of fast disks using a high performance parallel file system for beam lines with high rate detectors while beam lines with more modest demands might use a slower, less expensive storage solution.

The central storage system will be directly connected to compute resources and these will be accessed using a remote graphical desktop or terminal interface. The computing resources will mainly be of the general-purpose type, but some applications will require GPU nodes.

As mentioned before, data needs to be processed at the MAX IV Laboratory and transformed into a format usable to the scientist. Depending on the experiment, this transformation involves a few simple steps calibrating the measured data or, at the other extreme, a fully automated workflow with parallel steps. The latter will use a lot of computing power, and require a more complex workflow and analysis tools.

Data will be kept on disks for approximately two months after data collection in order to be available for easy data transfer off site.

The needs to organize the data will require a metadata catalogue and a presentation tool, e.g. a web portal. The data catalogue will contain information about the conditions at the experiment, who has what access rights, location and links to the raw data files and the possibility to enter some analysis results and annotations for whole experiments and individual data sets. The data in the metadata catalogue will be saved indefinitely, even if the raw data is deleted. For highly automated beam lines, there is a need for added features to the metadata catalogue and web portal, allowing sample tracking and connection to the data acquisition service.

Fast detectors will produce data streams requiring the whole bandwidth of a 10 Gbit/s network connection and tuned file systems to manage the stream. In these cases the data storage and necessary compute power needs to be located close to the source. Although it could be technically possible to add enough bandwidth to overcome this, it would be very costly. Overall, adding more specific components to the infrastructure also increases the risk of a service failure in a data workflow which has to operate 24/7 during an experiment. To provide the best support to ensure the continuous operation, the e-Infrastructure for data production has to be located at MAX IV.

A user handling service for managing the process of all users' visits is needed at the MAX IV Laboratory. This service includes experimental proposal submission and their review process, including experimental safety handling. Beam line access scheduling and reporting would also be a part of this service.

In terms of required competences, there is a need to have expertise available at MAX IV for all parts of the e-infrastructure. Besides staff to support the e-infrastructure, data acquisition and implementation of online data analysis workflows described above, MAX IV will require support for a core set of analysis applications. This would include installation, license management and general application support as well as expertise in a common framework and workflows enabling scientist to perform their research.

B. Research requirements

The role of MAX IV Laboratory in the research step is still under discussion for the scientific analysis of data as well as the long-term curation.

Presently, there is no provision for a scientific analysis service offered by the MAX IV Laboratory. Instead, the data would have to be transferred to the scientist home institute for analysis after the pre-processing has taken place at MAX IV. Due to the file size produced by new detector technologies, this is an increasingly difficult task, where a data set would require days to download (equal or longer than needed for acquisition). Some scientific communities, such as the structural biologists, have started their own project in collaboration with the SNIC center NSC in Linköping, in order to solve this. Recent initiatives in Denmark have resulted in an application to the Nordic e-infrastructure Collaboration (NeIC), that if successful will provide a Scandinavian solution for Imaging of materials. MAX IV has identified a data transfer service as a key service component to enable scientific analysis.

Such a data transfer service would transfer data files automatically or by request to other computing centers or other institutes while retaining ownership of the data and keeping track of the location of files. To ensure that data can be acquired, stored, pre-processed, transferred and analyzed within one workflow, MAX IV has built a data management prototype with the SNIC center Lunarc in Lund during spring 2014. The prototype also highlights the need of a common authentication system and the required resources for running the services.

It is important to note that, although the community efforts mentioned would solve an analysis problem, they are only targeted at Swedish or Nordic users. MAX IV has users from all over the world, and if an analysis service would be included in the approval of a proposal, in it has to be agnostic in terms of nationality.

Open Access, as defined by the EU, would apply to data collected at MAX IV and at other research Infrastructures in Europe. Many facilities, who MAX IV are collaborating with, are already preparing for this and their services are adopting to accommodate the needs. For example, the PaNdata Open Data Initiative project, within FP7, has resulted in a metadata catalogue providing means to provide open access to data and as an added bonus, the possibility to mint DOIs (Digital Object Identifiers) for citing experimental data.

In terms of Open Access and long-term curation of data, the MAX IV services would be capable of handling the task in terms of managing the data. The data store could be located elsewhere, e.g. SweStore. However, MAX IV will not provide this without input and directive from the funding partners.

7) Roadmap for implementation.

- 2015
  - Develop basic infrastructure in terms of services for acquisition, storage and analysis of data
  - Provide first version of a scientific data management service
- 2016
  - Deploy scalable high-speed file systems and compute nodes according to the needs of the beam line projects
  - Develop new services for automated workflows and data treatment
  - Deploy metadata catalogue
  - Deploy common analysis platform for data treatment
  - Develop a remote access service for data collection
  - Deploy user management service, DUO (Digital User Office)
- 2017
  - Increase storage system to meet beam line demands.
  - Firewall upgrade likely to facilitate off site data transport
- 2018
  - Increase storage system to meet beam line demands.
  - Deploy workflows and analysis services for high volume beam lines.
- 2019
  - Increase storage system to meet beam line demands
  - Firewall and external network upgrade envisioned

A. Production requirements

	2015	2016	2017	2018	2019	Unit
CPU	0.2	0.9	1.5	2.0	4.0	million core hours
Storage	50	280	700	800	1250	Terabyte
Support	7	9	10	10	10	FTE
Network*	2x1	2x1	2x10	2x10	2x40	Gigabit/s

* Firewall throughput to external network.

The computing needs for most beam lines are not very high, many will not require much more than offered by a high-end desktop computer. Instead, it is the combination with fast disk access that puts a high demand on the e-Infrastructure. The access to fast computing nodes based on GPUs is projected but not currently known. There are quite a few uncertainties in the projected data volumes. Factors which influence the total volume include (but are not limited to):

- Compression ratio of data files
- Detector type. No detectors have been procured at the writing of this document
- Sample characteristics defining data acquisition speed
- Degree of automation available at beam line
- Sample preparation time
- Sample alignment

The final data volumes could be higher as well as lower than these estimates.

The bandwidth of the external network connection will depend on what analysis tools are offered in the end. These estimates are based in the assumption that data reduction takes place at MAX IV Laboratory and the network is used to transfer data of site after this step. Collaborations, where raw data is analyzed off site, will increase the need.

B. Research requirements, long term storage (e.g. no data deletion)
Since there is no initial requirement to be able to do post analysis on processed data at MAX IV Laboratory, no numbers for e.g. compute capacity have been added. The only numbers given are the accumulated data volumes for the time period, based on the numbers in the previous table.

	2015	2016	2017	2018	2019	Unit
Storage	250	1650	4900	7500	10250	Terabyte

8) Sensitive data.
The infrastructure will deal with sensitive personal data and/or projects that require ethical approval. This must be taken into account when designing the supporting IT systems. However, it will not affect the studied time period in a major way as the need will mainly be driven by experiments at the medical imaging beam line MedMAX that is projected to come on line after 2019. The anticipated user community will use MedMAX for in vivo imaging of small animals in applications for physiological and morphological characterization of organs, as well as for micro-localization of toxic elements and tumour-targeting molecules in tissue samples.

4.2 XFEL - The European X-ray Free Electron Laser

The European XFEL facility is more than three kilometres long and stretches between the DESY site in Hamburg-Bahrenfeld and the town of Schenefeld in Schleswig-Holstein. The European XFEL will produce extremely brilliant, ultra-short pulses of spatially coherent X-rays with wavelengths down to 0.1 nm and below, and this radiation is accessible through ten experimental stations. Operated as a user facility, the XFEL is expected to provide results of fundamental importance in material sciences, plasma physics, planetary sciences, astro-physics, chemistry, structural biology and biochemistry, with significant effect on applied and industrial research. The facility will be commissioned in 2016. User operation will start in 2017 with one beam line and two experiment stations.

1) Scientific disciplines.
Physics, Chemistry, Biology, Mathematics and Computing.

2) Coordinators.
Filipe R.N.C. Maia and Janos Hajdu, Uppsala University.

3) Participating institutions.
In Sweden: Uppsala University, Stockholm University, KTH Royal Institute of Technology, Lund University, University of Gothenburg, Umeå University, Chalmers University of Technology.

Vetenskapsrådet (Sweden), DASTI (Denmark), CEA (France), DESY (Germany), NIH (Hungary), MIUR (Italy), NCBJ (Poland), OJSC RUSNANO (Russia), Ministry of Education (Slovakia), MINECO (Spain), SBFI (Switzerland).

4) Short description of the Research Infrastructure.
X-ray free-electron lasers (XFELs) deliver X-ray radiation with a peak brilliance more than ten billion times greater than what was available before. Such a jump in a physical parameter is both remarkable and rare and can lead to revolutionary new advances in science.

The European XFEL will be capable of producing billions of shots per day when it comes online in 2016. This represents more than a hundred fold increase in capacity and data rate as compared to existing XFELs.

This abundance of data promises to lead to a revolution in structural and materials sciences but has to be combined with effective and efficient algorithms for data handling and data analysis, as well as adequate computing infrastructure.

Data banks with experimental data are crucial for education and research, aiding the development and validation of new theories and techniques. The Protein Data Bank is a remarkably successful example of such a database. The Coherent X-ray Imaging Data Bank (CXIDB, www.cxidb.org) was set up by Filipe Maia and is dedicated to the archival and sharing of data from free-electron lasers. Such data are currently available only to an extremely limited number of people. CXIDB enables anyone to upload experimental data and browse data deposited by others. This data bank is one of the "Approved Data Banks" of Nature's new publication called Scientific Data, and is currently maintained on a temporary basis at the Lawrence Berkeley National Laboratory. CXIDB will play a central role when XFEL begins user operations and needs a permanent home.

5) References to recent contracts or other relevant documents from funding agencies.

The Swedish participation in XFEL has been supported by 38 grants so far. Recipients: Janos Hajdu, Filipe Maia, Marvin Seibert, Bianca Iwan, Inger Andersson, Jakob Andreasson, Gergana Angelova, Jan Isberg, Jan-Erik Rubensson, Nicusor Timneanu, Volker Ziemann (UU), Christian Bohm, Anders Hedqvist, Mats Larsson, Reinhold Schuch (SU), Raimund Feifel, Richard Neutze (GU), Ulrich Vogt (KTH). Granting agencies: Vettenskapsrådet, KAW, SSF and ERC.

6) Description of e-Infrastructure requirements.

A. Production requirements

The European XFEL, when fully operational, is expected to produce datasets of up to 2 PB per day. Most of the infrastructure needs for production will be addressed directly by the European XFEL. Here, we concentrate on the necessary extra infrastructure to support Swedish researches within Sweden to process and store their data, after their experiments at XFELs.

The estimated Swedish share of data that can be expected from the European XFEL is based on the initial operations parameters of the European XFEL and the current usage of X-ray free-electron lasers by Swedish scientists. In particular, we assume that most of the data will come from area pixel detectors. The initial detectors will only be able to record 3520 images per second, not the full 27000 that the machine produces. This latter value will be reached later. The specifications of XFEL are likely to be upgraded throughout the life of the facility, easily increasing the data production by an order of magnitude or more.

Comparing with the Linac Coherent Light Source (LCLS), the leading X-ray free-electron laser today, which records data at 120 images per second, the European XFEL's initial output corresponds to roughly 30 times more data. The LCLS currently produces a total of 3.5 PB of data per year (including various shutdowns and maintenance). Scaling this number to the European XFEL and assuming similar operational schedules, gives a yearly data production of 100 PB.

Sweden is receiving about 280 TB of data per year from LCLS. Assuming that the share of applicants per country for beam time at the LCLS (see Figure 1) roughly corresponds to the share of the data that each country would receive from the European XFEL (which might be an underestimate given the geographical proximity of the XFEL facility) gives a Swedish yearly data production of 8.4 PB. This is expected to grow.

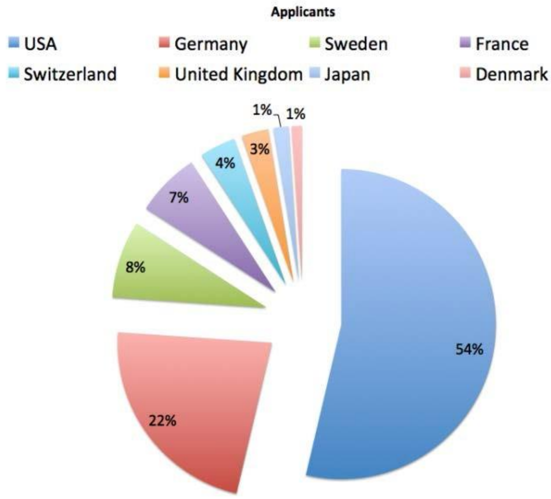

Figure1. Share, by country, of applicants to the LCLS.

While some preliminary filtering of the data will be done at the facility, experience at other facilities (FLASH at DESY and the LCLS at SLAC) has shown that a significant fraction of the data will need to be brought home. The reason for this is that the facilities do not have long-term storage capacity (data are erased after a fixed period of time, typically after 3-6 months), and the facilities do not have an environment optimized for data analysis. In addition, the facilities lack a focused effort to make all data available to the public.

Thus the first challenge is having enough network capacity to transfer the data to Sweden. We expect to require about 100 Gbit/s of network bandwidth during a few weeks following each experiment to bring the data to Sweden within a manageable time frame. Storage capacity is needed for this data, with high sequential read bandwidth because a significant part of the initial analysis requires going through the large datasets trying to filter out images. High-availability, high IOPS and low latency are not a requirement.

B. Research requirements

Due to the large diversity of experiments that can be done at an X-ray free-electron laser and the inherent unpredictability of emerging techniques, it is difficult to give an accurate description of the infrastructure needs for research.

A two-pronged strategy is envisaged:

- Using CPUs to form the backbone of a flexible analysis infrastructure for a diversified range of low computational cost data analysis problems. For example, simple filtering and classification schemes, which are likely to be I/O bound, would fall here.
- Using GPUs for more well-defined and computationally demanding problems. The prime example is the 3D assembly of a Fourier diffraction volume from a large number of, low signal-to-noise, 2D Ewald sphere sections. GPU implementations of the EMC algorithm for such a problem already exist. Another example is 2D and 3D phasing of diffraction volumes, which can be accomplished with Hawk, an image reconstruction package with GPU support.[11]

LCLS data is currently analyzed using a cluster of 32 nodes, with 4 GPUs each. Scaling current computing resources with data production ratio gives an estimate of 1000 nodes, each with 2 CPU (currently 6-core Sandy

[11] http://xray.bmc.uu.se/hawk/

Bridge Xeon), 64 GB RAM, 4 GPUs. QDR Infiniband fabric seems adequate with the most communications intensive application, the 3D Fourier Assembly, which would also benefit from large GPU memories. Since all current algorithms are still some way from being fully standardized, estimates are based on preliminary scaling tests. These tests confirm that current EMC implementations scale linearly to 32 nodes in low-resolution experiments, and should scale much further with higher resolution (increasing the amount of work per node).

Besides providing the necessary support for data analysis, another important element to maximize the output from expensive X-ray free-electron laser experiments is to make the data as widely available as possible. The Coherent X-ray Imaging Databank (CXIDB) was built with that purpose in mind.[12] It is currently hosted by NERSC at the Lawrence Berkeley National Laboratory through an agreement, which has been extended annually, and no future guarantees exist. Creating a mirror of the database, with long-term assurances, would greatly help data dissemination at a moderate cost. The data storage needs are estimated to be about 500 TB per year. Most of these can be stored in high latency devices such as slow-spinning disks or write once optical media. Data, once deposited, are not expected to change. The network requirements are also modest as users typically only download sections of datasets.

Finally, there is a need for more software for visualizing the large datasets that will be produced. The biggest constraint here is simply lack of manpower. An open source python based visualization tool for LCLS datasets, Owl[13], is already being developed. This uses h5py and OpenGL to quickly visualize Terabyte sized datasets using modest memory resources. Such efforts should be expanded and adapted to cope with the increased data volumes of the European XFEL.

7) Roadmap for implementation.

The roadmap for developing and implementing the large-scale Swedish research infrastructure for handling data from XFEL consists of the following elements:

A **survey of the field** and an **analysis of needs** has been completed and a concept for the research infrastructure has been developed with the aims of (1) receiving data of Swedish users from XFEL, (2) providing software to sort and analyze the data, (3) archiving the data for Swedish users, and (4) creating the possibility of depositing useful data with CXIDB.

A **test facility** has been set up to refine requirements and to develop and test software, hardware and networks and communications. This **pilot project** began recently. It is based on a collaboration between SNIC-UPPMAX and the Laboratory of Molecular Biophysics (LMB). LMB made their new GPU cluster (the largest in Sweden today) available to UPPMAX and SNIC. UPPMAX provides space for the cluster and storage media, and runs the system as an "UPPMAX system". LMB is part of the DataXpress User Consortium at XFEL and works with XFEL scientists on data handling. The pilot project in Uppsala strengthens **collaborations** between UPPMAX, LMB, SNIC and XFEL.

The large-scale research infrastructure project is approaching the level of readiness for realization.

Resources presently available for the pilot project are (1) GPU cluster and storage media placed at UPPMAX; (2) limited personnel from the existing staff of the partner labs.

Future needs: Resources in the range of 4000 GPUs (Kepler GK110 equivalents) for timely 3D reconstruction of the expected high-resolution data from the European XFEL are needed. This need will be intermittent in nature, but a smaller subset can be used for methods development.

Storage for at least 8.4 PB new Swedish primary data each year is necessary. Data retention rules will need to be developed to determine what to save for how long, but at the very least ten years of this should be expected. Add approximately 20 % for processed data in different forms, totaling 100 PB over ten years. This is probably a very conservative figure. Sequential access to primary raw data needs to be high-bandwidth, but can

[12] http://www.cxidb.org/

[13] https://github.com/FilipeMaia/owl

be rather high in latency for random access. Additional needs for e.g. CXIDB are small but not negligible compared to this.

Data analysis needs high-bandwidth access to individual TB-PB size datasets. Significant bandwidth (100 Gbit/s) is also needed for timely transfer from the European XFEL to local resources.

One Full-Time Equivalent (FTE) could be sufficient to manage the cluster and CXIDB, if resources are pooled with UPPMAX so that more manpower is available for specific maintenance operations. Two dedicated software engineers would be needed for making existing algorithms "production-grade" and working on user-friendly and dependable visualization and data management tools.

Production requirements and Research requirements

The table below lists the resources that are required for production and research. The two are merged as there is a tight dependency between the two.

	2015	**2016**	**2017**	**2018**	**2019**	**unit**
CPU	1.5	3.0	10	15	15	million core hours
Storage	3 000	8 000	15 000	25 000	35 000	TeraByte
Support	3.0	3.0	3.0	3.0	3.0	FTE
Network	10	20	100	100	100	Gigabit/s
GPU	512	2 048	4 096	4 096	4 096	GK110 Equivalent

8) Sensitive data.

The XFEL infrastructure does not deal with sensitive personal data.

4.3 NGI – National Genomics Infrastructure

The National Genomics Infrastructure (NGI) is the largest technical platform at SciLifeLab and provides access to technology for massively parallel/next generation DNA sequencing, genotyping and associated bioinformatics support. The next-generation DNA sequencing techniques can be used for a variety of studies. Modern genome analyses critically depend on expertise in computational biology (e.g. bioinformatics, biostatistics, and theoretical systems biology). Such expertise is closely integrated with NGI in order to optimize throughput, data handling, and basic analysis. The technical development in the genomics area has been overwhelming in recent years and the next generation sequencing instruments now allow large-scale genomics on a previously unattainable scale. The platform comprises two nodes: NGI Stockholm (Genomics Applications and Genomics Production) and NGI Uppsala (SNP&SEQ Technology Platform and Uppsala Genome Center). NGI was launched January 2013, originating from the RFI-funded infrastructure SNISS.

1) Scientific disciplines.

Genomics, Massively parallel sequencing (Next Generation Sequencing), SNP-genotyping, Bioinformatics.

2) Coordinator.

Joakim Lundeberg, KTH Royal Institute of Technology.

The facilities are led by Joakim Lundeberg (KTH Royal Institute of Technology), Ulf Gyllensten (UU, UGC) and Ann-Christine Syvänen (UU, SNP&SEQ).

3) Participating institutions.

KTH Royal Institute of Technology, Uppsala University

NGI is supported as a national facility by RFI and by SciLifeLab.

4) Short description of the Research Infrastructure.

NGI offers all applications of massively parallel sequencing as a service to Swedish researchers, including consultations on project design prior to project start and primary bioinformatics support to its users.

The Stockholm node and the SNP&SEQ Technology Platform in Uppsala produce sequencing data using HiSeq and MiSeq technology (Illumina). The Uppsala Genome Center produces sequencing data using the Ion Torrent and Proton equipment (Life Technologies) and by single molecule sequencing using the PacBio system (Pacific Biosciences).

In 2013, NGI completed 550 sequencing projects encompassing 14 000 DNA or RNA samples, and 70 SNP genotyping projects encompassing 30 000 DNA samples. In 2013, about 40% of the users of NGI were from other Swedish universities than the host universities of NGI. The volume of the sequencing services is expected to increase by at least 50% in 2014. The capacity and data out-put of the sequencing systems is expected to increase exponentially as shown below.

In total, NGI has a staff of about 70 FTEs, including laboratory engineers and technicians, bioinformaticians, facility managers and project coordinators.

5) References to recent contracts or other relevant documents from funding agencies.

Swedish Research Coucilcontracts:

- SNP&SEQ Technology Platform, running grant, contract no. D0203101 (2103-11-07)
- NGI (former SNISS), running grant, contract no 829-2009-6239

6) Description of e-Infrastructure requirements.

Sequence data is produced by NGI at facilities in Stockholm and Uppsala hosted by SciLifeLab. After production, data is transferred to SNIC-UPPMAX where it is pre-processed. Raw data is archived to SweStore. Pre-processed data is delivered to research projects on the SNIC-UPPMAX file system for further analysis by scientists. A best practices analysis pipeline is executed upon request from projects prior to delivery. SNIC-UPPMAX has a separate production cluster at Uppsala Biomedical Centre, while the research cluster is located at the main computer hall at Ångström Laboratory.

A. Production requirements

The infrastructure needs for production can be summarized as follows:

- Temporary disk-based storage and general purpose computing resources for data pre- processing and best practice analysis. Storage bandwidth should be >10 GB/s
- Disaster backup (tape), currently works well via SNIC – required bandwidth for 2-4 TB per day
- Network 10 Gbit/s to SUNET backbone
- Encrypted backup at an expected level of 1-2 TB/day (incremental)
- Archiving of raw data to SweStore (250 GB per whole genome)
- Redundant production resources to ensure high availability, must be located in separate computer hall and connected by network of preferably 20-40 Gbit/s bandwidth
- Production cluster should preferably have emergency electricity, emergency cooling, reserve electricity and reserve cooling
- Production infrastructure should conform to PUL regulations
- Network between production and research cluster should preferably be 20-40 Gbit/s
- A wide range of continuously maintained bioinformatics software and reference data
- System administration for resources
- Application experts for maintaining software and reference data

B. Research requirements

The infrastructure needs for research (by NGI users) can be summarized as follows:

- General purpose computing resources
- Disk-based storage for projects with bandwidth >10 GB/s
- Disk-based scratch storage with bandwidth >10 GB/s
- Long-term storage (archiving) for project-specific data and results
- Backup expected to be 5-10 TB/day (incremental)
- For substantial part of projects, resource should conform to PUL regulations
- A wide range of continuously maintained bioinformatics software and reference data
- System administration for resources
- Application experts for maintaining software and reference data
- Dedicated support-function for end users (currently 0.75 FTE is dedicated to this)
- Databases for collecting sequenced data in accessible and interoperable formats and interfaces, and corresponding data access committees (DAC), as well as services for publishing results to public repositories, e.g. EGA (www.ebi.ac.uk/ega/)
- Extensive education and training of user community inexperienced with HPC

7) Roadmap for implementation.

The infrastructure needs for production and research for 2015-2019 are given in the two tables below. The numbers are based on the following estimates. Current sequencing capacity and in addition a large number of whole genomes: 15000 whole genomes in 2015, 30000 in 2016, 35000 in 2017, 40000 in 2018 and 45000 in 2019. A substantial investment in sequencing capacity is made in 2015 and 2016 to a level of 30-50000 genomes per year where the level is expected to be for a longer term. A reference population of approximately 5000 Swedish individuals is also expected to be sequenced during 2015-2016. The numbers have been estimated by NGI together with SciLifeLab management. Note that this represents almost a 10-fold increase from sequencing capacity of 2014. We have calculated on 30x genome coverage (which is standard today) which might very likely increase in the future e.g. to 100x – this would increase storage demands even more.

A. Production requirements

	2015	**2016**	**2017**	**2018**	**2019**	**unit**
CPU	25	50	55	65	70	million core hours
Storage	3 000	5 000	5 000	6 000	6 000	TeraByte
SweStore	4 000	12 000	21 000	31 000	42 000	TeraByte
Bioinformatics Platforms*	19	19	19	19	19	FTE
Support UPPMAX	1.5	1.5	1.5	1.5	1.5	FTE
Network	2x 10	2x 10	2x 10	2x 10	2x 10	Gigabit/s

*includes data delivery, systems/software development and R&D at SNP&SEQ.

B. Research requirements (by NGI users)

	2015	2016	2017	2018	2019	unit
CPU	100	160	190	210	240	million core hours
Storage	11 000	16 000	19 000	22 000	25 000	TeraByte
SweStore	4 000	6 000	8 000	9 000	10 000	TeraByte
Support UPPMAX	4.5	4.5	4.5	4.5	4.5	FTE
Network	2x 10	4x 10	4x 10	4x 10	4x 10	Gigabit/s

Another assumption is that data (raw data, mapped data, and variants) from whole genome sequencing are archived to SweStore (250 GB compressed data per sample), and stored for a maximum of 4 years. The adoption of new archiving formats like CRAM might reduce the size on disk. The regulations that all scientific data should be archived for 10 years seems out of the question from a practical point of view.

The requirements for production have a high degree of certainty as they are calculated from the number of sequenced samples. The computational requirements for research are based on the assumption that an acceptable research level is 3 times as much as production (today the ratio is 5:1) and the estimate hence has a higher degree of uncertainty. Research storage is based on the assumption that projects are two years long on average and make better use of resources than today.

Support is divided into

- Production: System administration 1.5 FTE (0.75 FTE storage + 0.75 FTE computing), platform bioinformatics support 19 FTE (4 FTE SNP&SEQ, 11 FTE KTH-Stockholm, 4 FTE UGC)
- Research: System administration and technical support 2 FTE (1 FTE storage, 1 FTE computing), software and data maintenance 0.5 FTE, dedicated technical user support 1 FTE, education 1 FTE.

8) Sensitive data.

NGI serves a wide range of research projects, some which have samples of human origin and hence a substantial proportion of NGI data is to be considered as sensitive, the fraction is estimated to 80%.

4.4 BILS - Bioinformatics Infrastructure for Life Sciences

BILS (Bioinformatics Infrastructure for Life Sciences) is a distributed national research infrastructure. The aim of BILS is to provide bioinformatics infrastructure and support for life science researchers in Sweden, on both local and national levels. The organizational structure should allow for changes in services over time as new techniques are developed and utilized. BILS is the Swedish ELIXIR node and coordinates the Swedish contributions to the ELIXIR infrastructure. BILS is hosted by Linköping University. BILS is supported as a national infrastructure/facility by RFI and by SciLifeLab.

1) Scientific disciplines.

Life sciences, Biology, Medicine, Agricultural, Pharmacology, (Chemistry).

2) Coordinator.

Bengt Persson, Uppsala University.

3) Participating institutions.

Linköping University (host), Umeå University, Uppsala University, Swedish University of Agricultural Sciences, KTH Royal Institute of Technology, Stockholm University, Karolinska Institutet, Naturhistoriska Riksmuseet, University of Gothenburg, Chalmers University of Technology, Lund University

4) Short description of the Research Infrastructure.

BILS is organized as a distributed national research infrastructure, with a number of nodes. The nodes provide support for bioinformatics issues, e.g. consultancy in experiment planning, analysis and interpretation of results, bioinformatics calculations, and training. Each node is responsible for one or several areas of expertise. Depending on the user needs, some areas are represented at more than one node.

In life sciences, the major data generation is currently from genome sequencing. The largest sequencing facilities are at SciLifeLab (including NGI), but there are also sequencing facilities at GU and LU, and additional facilities are under construction (LiU and others). Furthermore, some research groups have their own sequencing instruments. In addition, many Swedish research groups perform sequencing abroad, but the subsequent computational analysis of data is performed in Sweden.

Thus, the total computational needs are larger than what is generated by Swedish genomics facilities only. Nearly all computational analyses on the genomic side take place at SNIC-UPPMAX, and the estimates from UPPMAX are therefore of large value in predicting the development.

Other types of data in life sciences are structural data (X-ray crystallography and NMR), mass spectrometry data, metabolomics data, and imaging data. These amounts are still (2014) minor in comparison to genomics data, but are expected to increase during the upcoming five-year period. Especially imaging is requiring much disk space and has to be taken into consideration when estimating the total needs within life sciences.

5) References to recent contracts or other relevant documents from funding agencies.

- Hantering av utlysning inom forskningsinfrastruktur 2014
- Swedish Science Cases for e-infrastructure – 2014
- The Swedish Research Council’s guide to Infrastructures (2012)
- För svensk framgång inom forskning och innovation 2013–2016 (2012)
- Vetenskapsrådets:s årsredovisningar 2010 och framåt
- “Rekommendationer om avsättning av medel till nationell och internationell infrastruktur inom ramen för utlysningen av de strategiska forskningsområdena” (2009)
- “Utlysning strategiska forskningsområden” (2009)

6) Description of e-Infrastructure requirements.

BILS users’ needs are predominantly:

- Computational resources for analysis of data, e.g. large-scale genome sequence data. These needs are estimated by NGI since the resources needed follow available sequencing capacity.
- Temporary storage during analysis. Follows sequencing capacity.
- Long-term storage and public availability of results. BILS provides this in collaboration with SNIC centers. BILS follows open source and international publication principles. For sensitive data, special considerations have to be taken (see below), but there are mechanisms for this, e.g. the Finnish REMS system providing access to groups that have ethical permits to access data will be implemented within the ELIXIR framework. The disk space needed for long term data storage will increase considerably during the next years because of achievements on the sequencing side. Even though data compression and mapping towards reference genomes will decrease the amounts needed, there will still be a huge increase. For 2015 and 2016, the estimates are quite reliable since the roadmaps for sequencing are quite clear, but for the years 2017 – 2019, the needs might turn out even bigger than presently estimated, so those numbers have still to be considered as early approximations. BILS will at least annually coordinate with SNIC regarding the needs by the life sciences community in order to enable a good planning of acquiring data resources. If possible, storage of large data sets should be associated with a fee to cover the costs, which can be used to finance the data storage at SNIC.

As mentioned above, the largest data generators are the sequencing facilities, especially NGI. The BILS requirements therefore rely on the numbers provided by NGI and SNIC-UPPMAX/UPPNEX (Section 4.3).

Computing services. Most life science computing is x86-based. Some applications (e.g. GROMACS) can preferably use GPU for increased computational performance. However, life sciences predominantly need “conventional” clusters and since the number of samples to be analyzed in each project are in the range of thousands – millions, the computational requirements are “embarrassingly parallel” and can easily be met by “conventional” clusters. In the future, BILS expects more users to require cloud based solutions for easy access. In order to meet user demands for easy access, BILS is currently developing user-friendly interfaces to HPC resources, both in collaboration with UPPNEX and in collaboration with NSC (Fido).

Data services. BILS helps users with data publishing and to push the publishing internationally as high as possible, e.g. for sequence data, the complete sequences should be put in international databases such as EMBL, Ensembl, Uniprot, while the underlying raw data are still stored nationally but publicly available. The disk space is already considered in the estimates by NGI and SNIC-UPPNEX; it is only additional access routes that are provided by BILS for the data publishing part. BILS will assist the users with annotation, for which BILS is currently hiring a data manager.

Support, application expertise, training. This is the main aim for BILS and most of the BILS staff works with support, application expertise and training. This is funded within the BILS budget. BILS has close collaboration with the SNIC centers, especially UPPMAX/UPPNEX, where systems development and training activities are performed together.

Network services. Today’s level (10 Gbit/s) is still sufficient. BILS expects in the near future (2015/2016) an increase in network capacity needed, as more and more data are transferred between EBI and Sweden, and between ELIXIR nodes. This can be solved either by dedicated connections or, preferably, by expansion of the current system with additional wavelengths (=additional 10 Gbit/s connections) as the needs increase. SUNET can easily expand the bandwidth according to the needs of large user communities.

7) Roadmap for implementation.

BILS relies on SNIC for providing most computational and storage resources needed by BILS and its users, i.e. most of the life science community in Sweden. Currently, the SNIC centre UPPMAX is providing most resources (clusters Milou and Mosler) but also the SNIC center NSC is providing cluster resources for the Fido system. Funding for these systems are predominantly external (KAW), but there are also financial contributions by BILS and SNIC to these clusters. BILS finds this collaboration with SNIC very successful, where SNIC is providing the hardware and BILS is providing the content. Only systems needed for the web site, e-mail, project handling and similar small-scale matters are provided by BILS ourselves. Also specialized systems, where root access is needed, are set-up within BILS, e.g. the genome annotation cluster (7 nodes).

In the future, it would be convenient if SNIC could provide hardware resources to national research infrastructures in the same manner as currently commercial companies are providing server solutions for private persons and companies (e.g. ManuFrog). This would enable BILS to have a virtual server for a dedicated assignment, e.g. web site, local database, easily available via SNIC. It is important that such a service is flexible enough to meet the infrastructures’ needs and also being able to provide root access for the infrastructure.

There are considerable advantages with having life science calculations and storage nationally centered at UPPMAX since this facilitates support, maintenance and coordination with BILS. For provision of results digitally (public sequence data and similar), BILS together with SNIC has set up an iRODS system with one node in UPPMAX and a back-up node at NSC . This node will be linked to the European ELIXIR nodes in the future.

Production requirements and Research requirements

	2015	2016	2017	2018	2019	Unit
CPU	125	210	300	400	500	million core hours
Long-term storage*	8 000	18 000	30 000	45 000	60 000	TeraByte
Work storage	14 000	21 000	28 000	37 000	46 000	TeraByte
Support	6.5	6.5	6.5	6.5	6.5	FTE
Network	2x 10	4x 10	4x 10	4x 10	10x 10	Gigabit/s

* includes data publication.

The numbers for 2015 and 2016 are taken from the estimates for NGI (Section 4.3). For 2017 and onwards, BILS expects the needs to increase faster than their estimates due to data coming from other sequencing facilities and from related fields (imaging, mass spectrometry). These estimates are hard to make at present since the life science field is developing so fast. Thus, annual re-estimates have to be done in order to provide accurate predictions, but the numbers above are the current best estimate.

Looking back from 2010 at UPPMAX/UPPNEX and 2011 at BILS, the numbers of projects and users have increased linearly up to 2014. BILS expects this linear increase to continue also the next years – since the number of users increase as genomics is entering more and more research areas and since the data sets are getting increasingly larger as sequencing costs decrease.

8) Sensitive data.

About 80% of the data in the previous section will be sensitive.

Many users of BILS will use human sequence data, which per definition is sensitive data which need to be treated in a special way. Therefore starting autumn 2013, BILS has together with SNIC/UPPMAX initiated the project Mosler aiming at creating a secure environment for computation and storage of sensitive data. The system consists of a cluster with storage in a secure environment using two-factor authentication. The Mosler system will be pilot used for the Swedish Genomes projects during autumn 2014. Current capacity is 24 nodes and 273 TiB storage. After the pilot phase, the system needs to be extended during 2015 to match the needs from the users. Estimates from the NGI indicate that about 80% of all sequence data from 2015 and onwards will be of human origin, and most of these data will need to be analyzed and stored in the secure Mosler system.

Future funding of Mosler could either be from BILS or from SNIC. In any case, the hardware will be maintained by SNIC, while the user contacts will be handled via BILS. For large users, user fees will be discussed as the major route to obtain scalable funding.

Hardware-wise, the secure Mosler system uses the same type of HPC cluster and storage as the UPPNEX Milou cluster. Thus, capacity- and cost-wise, the numbers given in the section will not be changed, but merely the distribution of the resources between conventional and secure systems, depending on the type of data. The allocation of resources is facilitated by that both systems are hosted at the SNIC-center UPPMAX, which also gives large-scale advantages in that the software and user support is very similar for both systems.

The Mosler system currently set-up in Sweden is based upon the Norwegian TSD2.0 solution.[14] It seems likely that these secure compute and storage environments will be future Nordic contributions in the European infrastructure for biological information ELIXIR, since the Nordic countries are at the international forefront for handling personal data. Future ELIXIR contributions will not be any economic burden for Sweden, but on the contrary will provide possibilities to share development costs on the European level.

4.5 Swedish Bioimaging

Swedish Bioimaging (Svenska nätverket för biomedicinsk avbildning) is a network aiming at facilitating collaboration between researchers in Sweden working in bioimaging and related fields. Bioimaging methods are today an integral part in both biomedical research and clinical practice, and spans the full spectrum from molecule to man.

1) Scientific disciplines.

Radiology, Nuclear medicine, Radiation physics, Biophysics, Molecular microscopy, Bioengineering, Neuroscience, Microbiology, Cellular imaging, Biochemistry, Radiopharmaceutical chemistry, Image processing and Scientific visualization.

2) Coordinator.

Örjan Smedby (Director), Linköping University.

3) Participating institutions.

Linköping University, Lund University, University of Gothenburg, Karolinska Institutet, KTH Royal Institute of Technology, Stockholm University, Uppsala University, Umeå University.

4) Short description of the Research Infrastructure.

Bioimaging comprises medical imaging methods that may be used in clinical diagnosis as well as biological imaging methods based on light microscopy. National initiatives for Bioimaging currently exist in 18 European countries. The first of these to be formed was Swedish Bioimaging (bioimaging.se), which was constituted as a network of biomedical imaging researchers in 2009. The network now has around 150 members representing 11 Swedish universities and virtually all bioimaging facilities that are nationally accessible at all Swedish universities.

Currently, Swedish Bioimaging coordinates the following facilities, all RFI funded and accessible to researchers at all Swedish universities:

- Advanced Light Microscopy at SciLifeLab, KTH (Hjalmar Brismar)
- Centre for Cellular Imaging, GU (Göran Larson)
- 7T MRI facility, LU (Freddy Ståhlberg); installation planned for autumn of 2014
- Integrated whole-body PET/MRI, UU (Håkan Ahlström); installation April 2014
- Magnetoencephalography (MEG) (jointly funded by RFI and KAW); KI (Martin Ingvar)
- MRI and nanoPET-CT for in-vivo imaging, UmU (Helena Edlund); installation planned for 2014

5) References to recent contracts or other relevant documents from funding agencies.

The network has been awarded an RFI infrastructure grant for 2010-2014, and has also been awarded a RFI grant for large databases for 2012 and 2013.

[14] Tjenester for Sensitive Data, http://www.usit.uio.no/prosjekter/tsd20/

6) Description of e-Infrastructure requirements.
For the facilities mentioned above, Swedish Bioimaging is the tool for making them accessible to researchers at all Swedish universities. In order to fulfil this task, Swedish Bioimaging needs to store and transfer image data (which, like bioinformatics data, typically require much more storage space and bandwidth than other types of research data) from the national imaging facilities to researcher groups at all involved universities.

Therefore, e-infrastructure is needed by Swedish Bioimaging to:
- store image data from national bioimaging facilities in a central facility
- retrieve image data to other locations in a safe, easy and reliable manner
- share image data between Swedish and possibly international research groups

Archiving needs differ somewhat between the light microscopy and medical imaging fields. In microscopy, there is a stronger tradition of sharing data between research groups, and issues of patient integrity are less common. On the other hand, the image data come in a variety of image formats, both proprietary and open. This makes it imperative to offer strong support functions for handling microscopic image data, including for example image format conversion facilities. Some of the tools that are needed will probably have to be developed specifically for this purpose.

For image data produced by medical imaging scanners in the format of clinical examinations, the advantages of adhering to the DICOM - Digital Imaging and Communications in Medicine - standard are numerous. DICOM is the standard for the communication and management of medical imaging information and related data. An important issue to address is how to handle information about patient identity, in order to fulfil both requirements related to personal integrity and good research practice with respect to traceability and verifiability. Particular attention must be devoted to constructing a system where the access by users is controlled as appropriate in each project using suitable coding schemes. The possibility of joining data from several populations into larger studies will also require careful consideration of integrity issues, probably most often solved by handling anonymized data.

To guarantee easy and versatile handling of medical image data in DICOM format, Swedish Bioimaging wants to include and adopt an existing Open Source Picture Archiving and Communications System (PACS) system to use the SNIC storage infrastructure. This means that data stored in DICOM format will be available both via the standard interface of the general database and via a PACS-style interface. The information security aspects regarding the data stored must be handled on a "case-by-case" basis for each modality. The data owner (PI of each study) is responsible for this procedure, which has to be solved before data are transferred to SNIC.

For e-Infrastructure support to users of the bioimaging facilities, Swedish Bioimaging foresees a need of 2 FTE staff.

7) Roadmap for implementation.
A storage solution for the Centre for Cellular Imaging at GU has been implemented by SNIC staff (Anders Follin, LU) as a pilot project in 2013-14. Similar solutions will have to be implemented as additional Swedish Bioimaging facilities are taken into use in 2014-16.

Production requirements
The quantitative estimates in the table below are very rough. Annual re-estimates have to be done in order to provide accurate predictions, but these numbers are the current best estimate.

	2015	2016	2017	2018	2019	unit
CPU	-	-	-	-	-	
Storage	300	500	700	800	1 000	TeraByte
Support	2.0	2.0	2.0	2.0	2.0	FTE
Network	2x 10	2x 10	2x 10	2x 10	2x 10	Gigabit/s

8) Sensitive data.
Confidentiality issues will probably be relevant for all the medical imaging data produced by Swedish Bioimaging facilities (cf. above). This may roughly be estimated to constitute 50% of the amounts given above.

4.6 WLCG - Worldwide LHC Computing Grid

The WLCG e-Infrastructure is a part of the Large Hadron Collider (LHC) infrastructure. LHC is the world's largest accelerator complex, dedicated to discoveries of new elementary particles and studies of physics and cosmology phenomena. The LHC has 4 major experimental stations, of which two (ALICE and ATLAS) involve Swedish physicists. Specifics of the LHC are that its storage and computing infrastructure (WLCG) is largely outside CERN. The WLCG e-Infrastructure is in operation since 2006 and is expected to operate until 2035.

1) Scientific disciplines.
High Energy Physics

2) Coordinators.
Ian Bird (CERN), Richard Brenner and Oxana Smirnova (Sweden)

3) Participating institutions.
Over 100 CERN-related institutions worldwide. In Sweden: Linköping University, Lund University, Stockholm University, Umeå University, Uppsala University and KTH Royal Institute of Technology.

4) Short description of the Research Infrastructure.
The WLCG e-Infrastructure consists of more than a hundred computing centers all over the world, and is responsible for LHC data storage and analysis, as well as simulation, specifically:

- Transfer experimental and simulated data between laboratories and researchers
- Store experimental and simulated data (disk storage) and archive experimental data (hierarchical mass storage)
- Process and analyze the data (mostly serial algorithms)
- Perform simulations of physics processes and experimental facilities (also mostly serial algorithms)

The key purpose is provision of High Throughput Computing focused on fastest possible processing of large amounts of data.

The LHC experiments at CERN record data at rates measured in Gigabytes per second. These volumes, together with derived and simulated data, are distributed worldwide at aggregate sustained rates reaching 2 GB/s. Transfer, archival and analysis of these data is done by the WLCG, which is built by all the countries that take part in the LHC experiments, including Sweden. All the data of LHC experiments are stored and processed on the WLCG; all the simulations are performed on the WLCG as well.

Contributing to this international e-Infrastructure is a necessary condition for getting unlimited access to the LHC experimental data and adequate computing power. Technically, access and utilization of this e-Infrastructure requires specialized software, adequate local computing and storage facilities, and adequate network connectivity.

5) References to recent contracts or other relevant documents from funding agencies.

- Swedish Science Cases for e-infrastructure (2014)
- "Operation and maintenance of the Nordic software platform for e-science infrastructures", Swedish Research Council contract 2012-4793
- "Operation of computer equipment for research at LHC", Swedish Research Council contract 2009-6243

6) Description of e-Infrastructure requirements.

In the WLCG context, simulation is performed in production mode, similarly to raw data taking. In this context, it is convenient to map the production part onto Tier-1 requirements, while the research part is mapped onto Tier-2 and Tier-3 requirements. It is also convenient to group CPU and storage requirements per Tier-1/Tier-2/Tier-3, since these levels are governed by different agreements with CERN.

WLCG implements a hierarchical infrastructure based on Grid technologies:

- Tier0 at CERN stores raw data from the LHC detectors, and distributes copies to 12 Tier-1 centers around the world, one of them being Nordic (NDGF-T1). At Tier-1, the raw data is reconstructed into data meaningful for further physics analysis.
- Tier-1 centers archive data on tapes and on disks, produce derived data, and contribute to simulations and data analysis. The Tier-1 clusters perform reprocessing of data when new algorithms become available. Tier-1 centers (or their funding agencies) sign a Memorandum of Understanding (MoU) with CERN and are bound by Service Level Agreements, serving international users. Sweden currently contributes to the NDGF-T1 ca. 1100 equivalent CPU-cores (~10 million CPU-hours per year), ca. 1 PB of disks and ca. 1.3 PB of tape storage. SUNET is connected to the NDGF- T1 via a 10 Gbit/s channel, which operates near saturation.
- Tier-2 centers store data on disks to carry out most of data analysis and simulations; simulations results are archived on Tier-1. Like Tier-1, Tier-2 centers sign the MoU with CERN. Swedish Tier-2 currently contributes ca. 700 CPU-cores (~6 million CPU-hours per year) and ca. 1 PB of disk storage.
- Tier-3 centers are not bound by the MoU and are used by local researchers for their daily development of analysis code, data analysis and smaller scale simulation. The Tier-3 is the main workhorse for the physicist developing analysis strategy. They receive data from Tier-2, or possibly from a Tier-1. There are three Tier-3 sites in Sweden: in Stockholm, Uppsala and an upcoming one in Lund, expected to provide ca. 1.5 million CPU-hours per year each, and between 20 and 100 TB of storage each.

CPU-cores described above are of general-purpose architectures. At the moment and in the foreseen future, WLCG does not require advanced processors.

Data storage on Tier-1 resources is long-term. A Tier-1 is also required to have mass storage systems, as indicated above. Storage on Tier-2 resources is of less long-living nature. However, in both Tier-1 and Tier-2 cases, storage management decisions are made at CERN. Tier-3 storage depends on the needs of local research groups in Sweden, and is fairly volatile but essential for the analysis work. In addition to Tier-1 and Tier-2 storage, each computing resource needs locally mounted low-latency high-performance disk cache in order to support simultaneous read and write by very large number of processes. The size of such local cache depends on the size of the computing facility. Apart from storage, a number of other servers is needed in accordance with the WLCG architecture, such as a file transfer server, cache indexing server, etc.

Data cataloguing, annotation, curation, access permissions, publishing etc. are handled from CERN and do not require additional resources beyond deployment of the necessary WLCG services.

In order to connect Nordic and Swedish resources and researchers to the WLCG e-Infrastructure, the ARC software is used, which was developed and supported by Nordic High Energy Physics groups as their contribution to the WLCG. Data storage, analysis and simulations rely on other specialized external software. The infrastructure is operated by expert system administrators on a 24/7 basis, as required by the MoU. Each computing center contributing to the WLCG has at least two qualified experts; those supporting NDGF-T1 resources take regular Operator-On-Duty shifts for the entire Tier-1.

User training and support has two aspects. Basic technology training and local support for Tier-3 must be done in Sweden, where researchers are; Experiment-specific training and support is done at CERN. Training and local support in Sweden so far had an ad hoc nature due to the lack of dedicated personnel.

WLCG requires excellent networks and fairly complex routing, including optical private network and dedicated channels to CERN and other WLCG sites elsewhere. Currently available 10 Gbit/s is routinely saturated. In general, it is difficult to assess how large the demand will be in the future.

Clearly, network must not be a limitation, neither from the technical nor from the financial perspective: charging per traffic is not feasible.

Since the global WLCG resources are heavily loaded, it is very important for Swedish researchers to have access to regional and national e-Infrastructures (Tier-3), which have lower latency. Dedicated Nordic resources can offer such an opportunity. Local facilities that are not shared with other WLCG users are an even faster option, when available. Groups which have guaranteed prompt access to the necessary computing power have an important edge in the research world. Tier-3 processing is an activity that four groups (KTH, LU, SU and UU) are currently setting up. LU is presently installing a local Tier-3 service on a dedicated resource. For KTH, SU and UU, the Tier-3 processing will not be done on local hardware, but most likely on SNIC production resources.

7) Roadmap for implementation.

The WLCG e-Infrastructure is in operation since 2006 and is expected to operate until 2035. LHC data taking proceeds in periods: the so-called Run-2 period starts in 2015 and lasts until 2017, Run-3 will start after a break in 2019. Each such period will produce exponentially increasing amounts of data. During the break in 2017-2019, no new raw data will be collected, but simulation and analysis, as well as reprocessing of derived data, will proceed anyway. Exact requirements for 2016-2019 are not known yet and are subject to changes. The tables below show only Swedish shares of WLCG, calculated according to the agreed NDGF-T1 sharing key for Tier-1, and anticipated needs for the Tier-2 and Tier-3.

A. Production requirements (Tier-1)

	2015	**2016**	**2017**	**2018**	**2019**	**unit**
CPU	12.8	24.3	30.6	37.0	50.0	million core hours
Storage	1 220	2 510	2 970	3 300	6 200	TeraByte Disk
Storage	2 050	4 160	5 380	6 500	11 000	TeraByte Tape
Support	3.0	3.0	3.0	3.0	3.0	FTE
Network	2x 10	2x 10	2x 10	2x 10	2x 10	Gigabit/s

B. Research requirements (Tier-2 and Tier-3)

Tier-2:

	2015	**2016**	**2017**	**2018**	**2019**	**unit**
CPU	7.0	14.0	18.0	22.0	40.0	million core hours
Storage	920	1 800	2 700	3 100	6 000	TeraByte Disk
Support	1.0	1.0	1.0	1.0	1.0	FTE
Network	10	10	10	10	10	Gigabit/s

Tier-3:

	2015	**2016**	**2017**	**2018**	**2019**	**unit**
CPU	1.5	1.5	1.5	1.5	1.5	million core hours
Storage	100	100	100	100	100	TeraByte Disk
Support	1.0	1.0	1.0	1.0	1.0	FTE
Network	10	10	10	10	10	Gigabit/s

8) Sensitive data.
The data stored on WLCG resources carry no personal information and require no ethical approval and no encryption.

4.7 EISCAT_3D – The Next Generation European Incoherent Scatter Radar System

EISCAT_3D is a multi-static phased array radar system dedicated to observations of the Earth's polar atmosphere in order to study how the atmosphere is coupled to space. EISCAT_3D will give improvements of more than an order of magnitude in the temporal and spatial resolution of ionospheric observations compared to present systems that use the same incoherent scatter technique. In addition, it will be the first system of its kind that offers 3D vector imaging capability from the upper atmosphere over a wide range of altitudes in the ionosphere. The project is run by the EISCAT Scientific Association, an existing international research infrastructure at a unique location for research into the polar atmosphere in the northernmost region of Europe. The present EISCAT radars are located on Svalbard and on the northern Scandinavian Peninsula. The new EISCAT_3D system will be located at the latter location and will replace the existing system there.

1) Scientific disciplines.
Atmospheric and Geospace Sciences, Plasma Physics, Space Physics

2) Coordinators.
Craig Heinselman (Director) and Ingemar Häggström (Senior scientist), EISCAT Scientific Association

3) Participating institutions.
The Swedish Institute of Space Physics (IRF), KTH Royal Institute of Technology, Stockholm University, and Chalmers University of Technology (through OSO) are co-applicants of the EISCAT_3D proposal to the Swedish Research Council. A larger group of organizations is involved, including Luleå University of Technology, Umeå University, Uppsala University and the Swedish Defense Research Agency (FOI).

The EISCAT Scientific Association is currently funded and operated by research councils and funding organizations in Norway, Sweden, Finland, Japan, China and the United Kingdom and has its headquarters in Kiruna, Sweden.

4) Short description of the Research Infrastructure.
The EISCAT_3D instrument will consist of one core site with a transmitter and four remote receive sites. Each of these five sites will generate data at a high rate. The current plan for the sites is:

- Norway: Skibotn, the core transmit/receive site, and Andøya, a remote receive site
- Sweden: Bergfors and Jokkmokk, both remote receive sites
- Finland: Karesuvanto, a remote receive site.

The first stage of EISCAT_3D includes the sites located in Skibotn, Karesuvanto and Bergfors, with the transmit power at Skibotn limited to 5 MW. Later stages of EISCAT_3D will include higher transmit power and more sites.

In addition to the antenna sites, the EISCAT_3D Operations Centre and the EISCAT_3D Data Centre play important parts for the overall use of the instrument. The Operations Centre controls the measurements, monitors the production of the standard data products from the different sites, generates non-standard products as well as the products that result from combining the measurements from the different sites (multi-static data products) and transfers them to the Data Centre. The Data Centre provides users with access to data, some computing capacity for analysis and the EISCAT_3D data archive.

5) References to recent contracts or other relevant documents from funding agencies.

EISCAT_3D is discussed in "The Swedish Research Council's guide to Infrastructures" from 2012. The process of securing funding for EISCAT_3D is on-going. Proposals for funding for EISCAT_3D have been submitted in Finland, Japan, Norway and Sweden.

6) Description of e-Infrastructure requirements.

The e-Infrastructure for EISCAT_3D is still under planning, and details are likely to change over the next two years. The overall structure for EISCAT_3D consists of:

- Antenna sites. Each of the antenna sites exports data at a high rate (5 Gbit/s)
- The Operations Centre collects the data from all the antenna sites for monitoring and production of multi-static data products. The centre keeps a five-month buffer of data from the antenna sites (20 PB).
- The Data Centre provides user access and archiving. It will be distributed over (at least) two sites for access and data redundancy. The data archive grows with 2 PB/year. Currently, the plan is to co-locate the Data Centre with existing e-Infrastructure centres in northern Scandinavia.

The numbers given above are for the first stage of EISCAT_3D. Given this structure, the needs for the capacities will be:

- Computing: The Operations Centre requires 500 Tflops/s, the Data Centre requires 50 Tflops/s.
- Disk Storage: The Operations Centre requires 20 PB, the Data Centre requires 2 PB.
- Tape storage: The EISCAT_3D data archive requires 2x 2 PB/year.
- Network: Dedicated lines of 10 Gbit/s are needed to each of the antenna sites, 20 Gbit/s in between the sites of the distributed Data Centre. Additional capacity is required for the Operation Centre is still under evaluation.

Since the EISCAT_3D will produce data over a long time (decades), there are obvious needs for curation, annotation, metadata services, etc. How to address these issues is under investigation, but will not increase the capacity requirements significantly during the initial stage of EISCAT_3D.

The most recent estimate assumes that up to 13 additional staff members will be needed during the commissioning phase of EISCAT_3D, many of them IT experts. The continued e-Infrastructure planning in preparation for the implementation phase is estimated to require 1-2 FTE.

7) Roadmap for implementation.

Since the construction time schedule is not known in detail, the time schedule for the e- Infrastructure needs remains uncertain. The table below assumes that the needs for network and computing will be active from 2017. The planning for the EISCAT_3D e-Infrastructure will continue during 2015 and the numbers in the table below are still preliminary.

A. Production requirements

	2015	2016	2017	2018	2019	Unit
CPU Op. Centre	0	0	500	500	500	Tflops/s
CPU Data Centre	0	0	50	50	50	Tflops/s
Storage Op. Centre	0	0	20 000	20 000	20 000	TeraByte Disk
Storage Data Centre	0	0	2 000	2 000	2 000	TeraByte Disk
Archive	0	0	2x 2000	2x 4000	2x 6000	TeraByte Tape
Support	1-2	1-2	5-10	5-10	5	FTE
Network	0	0	3x 10 + 20	3x 10 + 20	3x 10 + 20	Gigabit/s

B. Research requirements
The users' additional e-Infrastructure needs will be limited during the build-up phase of EISCAT_3D, but are likely to increase later. The details of this are currently unknown.

8) Sensitive data.
The EISCAT_3D infrastructure does not deal with sensitive personal data.

4.8 OSO - Onsala Space Observatory

Onsala Space Observatory, the Swedish National Facility for Radio Astronomy, provides scientists with equipment to study the Earth and the rest of the Universe. OSO operates several radio telescopes in Onsala, 45 km south of Göteborg, and takes part in international projects. The observatory is also geodetic fundamental station.

1) Scientific disciplines.
Astronomy and Space Geodesy/Geoscience

2) Coordinators.
John Conway (Director), Chalmers University of Technology.

3) Participating institutions.
OSO is hosted by Department of Earth and Space Sciences at Chalmers University of Technology, and is operated on behalf of the Swedish Research Council. The use of OSO telescopes and use of OSO support for international radio astronomy projects is open to anyone at a Swedish academic institution and users come from a wide range of universities. There are no user fees in astronomy and telescope time is allocated on the basis of peer reviewed observing proposals. OSO geophysical observations are not driven by individual user proposals but instead provide continuous time series data to global databases which are then used by scientists for geophysical research or to provide societal benefit (for such things as defining the Earth's terrestrial reference frame and monitoring global change).

4) Short description of the Research Infrastructure.
OSO presently operates three radio telescopes at Onsala with another under construction. The existing telescopes are the 20m diameter millimeter wave telescope, the 25m diameter centimeter wave telescope and low frequency LOFAR low frequency phased array telescope. These telescopes operate in both stand-alone mode and as part of international networks (so called VLBI-style observing). In the network mode, the combined peak data production rates of the astronomy instruments at OSO is presently 4 Gbit/s (expected to increase to at least 30 Gbit/s by the end of the decade). A new telescope facility, the Onsala Twin Telescope (OTT) funded by the Knut and Alice Wallenberg (KAW) foundation, to be used exclusively for geodetic observations, is presently being procured, will be constructed in 2015 and becomes operational in 2016. This new telescope will eventually by the end of the decade produce raw data at 30-80 Gbit/s rates. In addition to the dish telescopes at Onsala, OSO is also a partner in the APEX single dish submillimeter wave telescope in Chile. OSO is also involved in large international facilities where it provides support for use by Swedish astronomers, In addition to supporting LOFAR and astronomical VLBI users in this way, OSO hosts a Nordic Regional Support node for the ALMA submillimeter array in Chile. This support consists of expert help and local computing resources for the reduction and analysis of ALMA data for Nordic users. Finally, OSO is heavily involved in the design of the Square Kilometre Array (SKA) telescope to be sited in Australia and South Africa; this innovative meter and centimeter wavelength array radio telescope will start construction of its first phase in 2018 with first science being possible by 2020. Globally, SKA will provide the largest data handling challenge of any scientific instrument to be built in the coming decade (comparable or larger than future CERN needs).

The above existing and projected OSO affiliated instruments produce many data streams of various types, but data processing requirements are dominated by the large array interferometers in which OSO participates (specifically LOFAR, astronomy and geodetic VLBI, ALMA and in the future SKA).

Data/processing needs for these arrays comprise several different levels. Level 1 input data consists of Nyquist sampled electric field data taken from individual radio telescopes antennas which are then transferred to a correlator (a data stream multiplier usually built as a dedicated hardware device but sometimes a supercomputer) to be combined pair-wise between telescopes to produce data products called visibilities. Because at this stage data is averaged in time and frequency, total data volumes in visibilities are generally significantly smaller than the sum of the input electric field samples. The visibilities are related to the Fourier transform of the sky and form the input data at level 2. Within level 2, these visibility input data are processed to make images and image cubes (where the third dimension is frequency or equivalently, for a particular spectral line, Doppler recession velocity, for some applications time is a fourth dimension). These data cubes are the primary data product from interferometry arrays which are distributed to astronomy users. Level 3 processing consists of taking these cubes and performing automatic source identification, source class parameter estimation (i.e. typical flux density and angular size of source classes), formation of source catalogues and other forms of post processing. Level 4 consists of astrophysical simulations to produce simulated observations to compare with telescope data products at level 2 and 3.

Broadly, the OSO data and processing requirements at the different levels defined above can be described as follows. At level 1 since no large scale correlation is envisioned in Sweden, OSO needs are for large capacity (up to 100 Gbit/s) data links to send raw data from telescopes at Onsala (i.e. the 20m, 25m, LOFAR and OTT telescopes) to correlator centres in other countries. Level 2 processing, i.e. the stage of image cube formation from visibilities, is supported for present instruments (i.e. ALMA, LOFAR and astronomy VLBI) by a mixture of processing provided by the central project and either initial or re-processing locally at Onsala on a dedicated small (50 core) cluster. Level 3 (image analysis) is also provided in part by Onsala facilities and in part central project facilities. Because of increasing data rates from the above instruments, OSO anticipates the need for some SNIC resources in these areas in the next 5 years. In the era of SKA (post 2020), national resources will be required at level 3 and possibly level 2; and experience with current instrument support at these data processing levels will be essential for planning for national computer support for SKA post 2020. Future estimates of resource needs at levels 1 to 3 are described below in Sections 6A and 7A. Level 4 processing (data simulation) needs are described in Sections 6B and 7B.

5) References to recent contracts or other relevant documents from funding agencies.

OSO operations are funded for 2015 under a contract with the Swedish Research Council with a proposal being submitted next year to cover 2016 – 2021. OSO needs for high speed communications are listed in the Swedish Science Cases for e-infrastructure (2014). ALMA and possible future involvement in SKA are described in the the Swedish Research Council's Guide to Infrastructures (2012).

6) Description of e-Infrastructure requirements.

A. Production requirements

The production requirements for networks are dominated by the transfer of sampled electric field (level 1) data to correlators in other countries. The likely evolution of data link requirements for combined astronomy and geodetic applications is given in Section 7. At present, primary correlator destinations are in the Netherlands. In the future, data will likely also be sent to centres in the UK and Germany and occasionally further afield (USA, Japan). These link resources for level 1 data are solely requirements for SUNET/NORDUnet/GEANT; there is no need for SNIC processing or storage requirements at level 1.

The infrastructure needs for CPU and storage are dominated by the processing at level 2 and 3 of visibility data imported back to Onsala from the correlators attached to each international interferometer array (ALMA, LOFAR, VLBI, SKA). Specific production requirements include (1) pipeline reduction of the very largest ALMA and LOFAR data sets on SNIC platforms (feasibility tests required first), (2) implementing specialized modes for data reduction of LOFAR international baseline data (level 2) data for multi-direction visibility data pre-averaging, and (3) specialized data reduction/data analysis (visibility based stacking analysis). These tasks require relatively modest computing, initially perhaps 10,000 - 100,000 CPU hours in 2015, but potentially increasing to a million core hours by 2019 if these become default techniques applied to for LOFAR/ALMA data. Short term storage needs on the processing cluster are 5 to 50 TB for large ALMA and LOFAR projects respectively, with the data needing to be available at the CPU nodes only during the period of reduction (order a day). In the short term, the largest SNIC commitment to realise these modes would be in terms of personnel to implement production processing and may require up to 0.5 FTE. The level of engagement described above would provide experience and benchmarking of the type of national level 2 and 3 processing needed to support SKA post 2020. For SKA, once phase 1 is completed (around 2023), level 2 (image cube formation processing) will require globally around 100 Petaflop and 1 Exabyte per year of long term storage. Most of this processing and storage will be provided at the telescope site and regional data center level. Some level of processing capability (for image re-processing and data analysis) would be useful at the national level – the level required and the amount to be supplied by other potential SKA partner countries at the national level are presently not defined.

On more immediate timescales than SKA, there is a possible role for SNIC to join the Long Term Archive (LTA) project of LOFAR which stores both raw visibility data and processed image cubes. The LTA is presently supported by Dutch and German supercomputer centres and presently stands at 3 PetaByte of permanent storage after two years of LTA operation. Sweden as a partner in the international LOFAR project heavily involved in the largest data volume observational projects, especially Epoch of Reionisation (EoR) observations, should ‘do its bit’ to support the LTA (perhaps part funded by a redirection of the funding presently contributed by Sweden to LOFAR central operations via Onsala). The minimum useful commitment would likely be of the order of 1 PetaByte.

B. Research requirements (e.g. data analysis, simulation)

Needs for research are dominated by astrophysical simulations which create simulated data sets to compare with observations produced by the large radio interferometer arrays such as ALMA, LOFAR and eventually SKA. The largest computational load is likely from continuing simulations of radio observations of the era of galaxy formation in the early universe (EoR). These simulations take the results of structure formation models and calculate the physical effects of ionizing sources on atomic gas and then compute the resulting spectral line emission in the red-shifted 21cm wavelength line of hydrogen. This simulation work, led by Professor Garrelt Mellema, is carried out at Stockholm University. Mellema is a leading member of the EoR core science teams of LOFAR and SKA. Over the next five years, these simulations will have vital importance for inter-comparison with and interpretation of the observational results emerging from the LOFAR EoR key project. These simulations will also be used to inform the design and operational planning of SKA for EoR science.

EoR simulation work is already a major user of SNIC and PRACE[15] and this use is expected to expand in coming years.

Other simulation types more relevant to the ALMA infrastructure include spectral line radiation transport in molecular clouds and galaxies and large scale astro-chemistry simulations. The former problem is computationally difficult because emitted and absorbed spectral line emission depends on the rotational-vibrational quantum state of molecules, but those quantum states in turn depend on the radiation environment the molecules find themselves in, giving a complex non-linear coupled problem to solve. Complementary large scale astro-chemistry simulations analyze large networks of chemical reactions and molecule creation/destruction mechanisms in interstellar molecular clouds to produce (time and position dependent) estimates of the fractional abundance of different molecules within a molecular cloud or galaxy. When combined with the radiative transport simulations described above, it is possible to make a final estimate of the expected strength of spectral lines to be compared with ALMA observations.

---

[15] PRACE – Partnership for Advanced Computing in Europe. www.prace-ri.eu

In the past, simple radiation transport and astro-chemistry simulations for single molecules/simple chemical networks have been carried out at the university group level with local computer resources, but this is unlikely to be tenable in the future. One of the ways ALMA is revolutionizing millimeter wave astronomy is via its very broad bandwidths which allows hundreds or thousands of spectral lines to be observed simultaneously. ALMA has transformed millimeter wave astronomy from a data starved to a data rich subject in which success in producing the highest impact scientific conclusions will depend increasingly on access to high capacity computing simulation resources. As the use of ALMA is expanded in coming years, it is very likely that other simulations types (such as simulations of continuum polarization or simulated ALMA simulations of continuum and line emission from forming planetary systems) will become increasingly important.

It should be noted that the use cases described in the above paragraphs refer only to astrophysical simulations run to produce simulated observations to compare with data products from OSO related infrastructures. Other types of astrophysical simulations not producing simulated observations (for instance testing theoretical mechanisms for grain growth in proto-planetary disks) or simulations producing simulated data at non-radio/millimeter wavelengths are not covered here.

Additional notes on resource requirements are the following:

- Computing services. Simulations generally require general-purpose x86-based computing, possibly supported with GPU, Xeon Phi or similar.
- Data services. Larger simulations (often synthetic data cubes) are useful for many studies and should be stored long term (also should store copies of matching observational data image cubes from LOFAR for inter-comparison). Large scale simulations should be stored for a couple of years on SweStore or similar, to be replaced by new simulation results after that time.
- Long term persistent storage is required (tape or successor technology). The ability to download and access when preparing publications is required but does not require fast access, only days or weeks of notice.
- The expansion of large scale simulations work on SNIC resources to areas beyond EoR simulations requires significant help to install and adapt existing codes developed for research group platforms (i.e. single workstations or small clusters) to be used on SNIC supercomputer level platforms.
- Data rates are not a significant bottleneck. Requirements are for the transfer of simulation results from long term storage to local user computers when publications are being prepared or from transfer of comparison observational data cubes from LOFAR (Groningen) or ALMA data processing centers (Garching, Germany or Onsala) to the Swedish long term archive or directly to user computers at Swedish universities.

7) Roadmap for implementation.

A. Production requirements

In the production area, the network requirements for level 1 data transfer (from OSO to correlators in other countries) are the most well defined/critical requirement. The likely time development of these network requirements are given in the table below, these are upper limits and the speed of ramp-up may be delayed. As stated in section 6A, the computing and storage needs for production over the next 5 years are relatively modest, with CPU requirements starting at 10 000 core hour tests and perhaps increasing to 100 000 core-hours with an absolute maximum of a million core- hours/year by 2019; these requirements are very uncertain. To implement these modes operationally, an immediate commitment of SNIC manpower (0.5 FTE) would be required.

Developing such a CPU production capability could, as well as providing unique science return in the short term, benchmark and test observing modes which may be required at the national level in the SKA era (post 2020). There is also, on the short term, a possible role for Sweden via SNIC to join the Long Term Archive (LTA) project of LOFAR where a minimum commitment of 1 PetaByte of long term storage would be required.

	2015	2016	2017	2018	2019	unit
CPU	0.01	0.05	0.1	0.3	1.0	million core hours (high uncertainty)
Storage	1 000	1 000	1 000	1 000	1 000	TeraByte (if Sweden joins LOFAR LTA)
Support	0.5	0.5	0.5	0.5	0.5	FTE (to implement large scale CPU production)
Network	4	16	16	36	110	Gigabit/s

B. Research requirements

The estimates of resources for the simulation area in the table below are made based on current and past experience with EoR simulations using SNIC/PRACE and future anticipated needs for comparison with LOFAR and planning for SKA. Other simulation types are difficult to estimate, but are likely to be smaller in at least the short term. In order to give a best estimate of total future resource requirements, the estimates from EoR have been multiplied by a factor of 1.5 for 2015 and 2016 and by a factor of 2 for later years, these factors are similar to the level of uncertainty on the EoR simulation estimate. The estimates are likely accurate with a factor of two.

	2015	2016	2017	2018	2019	unit
CPU	35	35	60	70	100	million core hours
Storage	220	220	500	800	1 200	TeraByte
Support	0.25	0.25	0.25	0.25	0.25	FTE
Network	10	10	20	20	20	Gigabit/s

8) Sensitive data.

None of the data requirements in this application area relate to personal data or ethical considerations.

5 CONCLUDING REMARKS AND NEXT STEPS

This report includes an initial inventory of the large scale needs for compute and storage infrastructure by a number of Swedish research infrastructures. It is recognized that this inventory is not necessarily complete and contains resource estimates that will be subject to change. SNIC therefore proposes that this inventory is updated at regular intervals and at least once per year.

The detailed requirements and the best mechanisms for addressing, evaluating or implementing them require further study involving experts from the research infrastructures, user communities and from the e-Infrastructures. Detailed studies and specifications of the required e-Infrastructure may span a number of years. SNIC and its partner centers are involved in a number of pilot projects for prototyping solutions with some of the infrastructures, mostly through the advanced user support (or application expert) positions that exist within SNIC. In discussions with some of the other research infrastructures, they expressed interest to work together on short-term pilot or competence projects to explore key functionality that satisfies agreed requirements. Such pilot projects serve multiple purposes. The projects should create working prototypes that are suitable for multiple user communities and research infrastructures. These could eventually lead to the introduction of new functionality and services in the production systems of the research infrastructure and national e-Infrastructures. In addition, the pilot projects should help refining the definition of the e-Infrastructure requirements and roadmaps for implementation for the research infrastructures.

Where appropriate, SNIC proposes that such joint pilot and competence projects are identified and planned with multiple research infrastructures to ensure maximum coverage and impact. For national research infrastructures that form a node in international initiatives, this collaboration should also be pursued in international context (e.g. EGI[16], NeIC[17], EUDAT[18]).

Besides piloting new functionalities and services, SNIC aims to define a framework for longer-term collaboration with other research infrastructures in order to:

- *Coordinate:* Agreements can be made between the parties to define how the infrastructures will interact, which aspects are in need for coordination, and who is responsible for which tasks and services to ensure that a complete, agile and state-of-the-art compute and storage infrastructure is provided and corresponding support for research data.
- *Plan:* Create a roadmap for the implementation of the e-Infrastructure in which SNIC can provide adequate infrastructure for computing and storage to the other infrastructure. In particular, establishing a common understanding of the needs for computing and storage infrastructure for the research infrastructure, and understanding the required planning that is needed to deploy this.

Such collaboration should maintain a rolling 3-5 year plan that describes at a high-level the roadmap for compute and storage infrastructure that is needed by the research infrastructure. This should be accompanied by an annual plan that describes in detail the required compute and storage infrastructure for the next 12-24 months, and which is in agreement with the rolling plan. The multi- year roadmap and annual plan are revised each year to allow all parties to plan and budget properly.

In addition to this framework, specific agreements can be made between the parties that define objectives, rights and obligations, and cost-sharing that the parties commit to for the provisioning and usage of the overall e-Infrastructure. Where appropriate, these agreements also define service levels, for example concerning user access, resource allocation and the availability of data that is stored on SNIC resources.

[16] EGI - European Grid Infrastructure, www.egi.eu

[17] NeIC - Nordic e-Infrastructure Collaboration, neic.nordforsk.org

[18] EUDAT – European Data Infrastructure, www.eudat.eu

6 APPENDIX: INSTRUCTION FROM THE SWEDISH RESEARCH COUNCIL

Uppdragsbeskrivning till SNIC avseende kartläggning av andra infrastrukturers behov av storskaliga datorresurser för beräkning och lagring

Bakgrund och utgångspunkt för kartläggningen
RFI har under åren beviljat bidrag för uppbyggnad och drift av flera stora infrastrukturer som har, eller kommer att få, stora behov av HPC-resurser och lagring. Dessa behov är till stor del ospecificerade och därför finns de inte med i budgeten, vare sig hos SNIC eller hos RFI. Detta har uppmärksammats från flera håll, t.ex. i bg4:s diskussion av rapporten *Utvärdering av svenska forskares behov av e-infrastruktur* och genom en skrivelse från SNICs föreståndare. Det är således av stor vikt att andra infrastrukturers behov av storskaliga datorresurser kartläggs.

Behoven kan komma från två håll:

I. dels direkt från infrastrukturens egen verksamhet,
II. dels från infrastrukturens användare.

En rimlig grundprincip bör vara att infrastrukturer (I ovan) som redan har beviljats bidrag av RFI också får tillgång till nödvändiga resurser hos SNIC. Likaså bör användare (II ovan) som beviljats tillgång till en infrastruktur för t.ex. ett experiment också få tillgång till de SNIC-resurser som krävs för att analysera data från experimentet.

Kartläggningen blir en viktig del i rådets och SNICs budgetplanering, och även i rådets strategiarbete (bl.a. i arbetet med *RFIs guide till forskningsinfrastruktur*).

Uppdrag
RFI uppdrar till SNIC att utföra kartläggningen som beskrivs ovan. Kartläggningen ska beskriva infrastrukturernas behov av HPC och lagring under de närmsta fem åren, samt vilka eventuellt ökade kostnader detta medför för SNIC.

Underlaget till RFI ska vara strukturerat så att det går att utläsa vilka behov som uppkommer från vilken infrastruktur, likaså vilka behov som rör (I) respektive (II) ovan. Behoven ska beskrivas i text och om SNIC anser det nödvändigt att äska om ytterligare resurser för att möta de infrastrukturernas behov så ska även en budget bifogas.

I arbetet med kartläggningen förväntas SNIC ha en tät dialog med andra infrastrukturer som RFI finansierar. De andra infrastrukturerna förväntas å sin sida att bistå SNIC i arbetet med kartläggningen. Det är ett gemensamt ansvar för SNIC och de andra infrastrukturerna att RFI får ett bra underlag.

Tidsplan
RFI ämnar diskutera kartläggningen vid sitt septembermöte. Allt material måste därför finnas forskningssekreteraren för e-infrastruktur tillhanda senast den 15:e augusti 2014. RFI anlitar ofta internationella ledamöter i beredningsgrupperna och kartläggningen bör därför redovisas på engelska.

The report presents a survey of the needs for supporting e-infrastructure from eight large data generating infrastructures. The needs for e.g. digital communication, storage and computing capacity to support operation of the infrastructures is estimated both in terms of technical specifications and costs. The report was produced by the Swedish National Infrastructure for Computing (SNIC) by invitation from the Council for Research Infrastructures (RFI). It offers RFI and other stakeholders guidelines for strategic planning.

Swedish Research Council | Västra Järnvägsgatan 3 | Box 1035 | SE-101 38 Stockholm, Sweden | +46 8 546 44 000 | vetenskapsradet@vr.se | www.vr.se