

Dellingr Phase 2: Deliverable 3

Pilot projects: NLPL and additional pilots

Dellingr team *

Wednesday March 25th 2020

*Mathias Brännvall mathias.brannvall@it.uu.se; Juha Fagerholm juha.fagerholm@csc.fi; Jens Svalgaard Kohrt svalgaard@sdu.dk; Ivar Koppel ivar.koppel@ut.ee; Bjørn Lindi bjorn.lindi@ntnu.no; Ilja Livenson ilja.livenson@ut.ee; Petri Nikunen petri.nikunen@csc.fi; Anders Sjöström Anders.Sjostrom@lunarc.lu.se; Máni Maríus Viðarsson mani@hi.is; John White john.white@cern.ch;

Contents

1	Executive Summary	3
2	Purpose of the Document	3
3	Dellingr Pilot	3
4	Nordic Language Processing Laboratory	4
5	Approaches to support the Dellingr pilot	5
5.1	Denmark	5
5.2	Estonia	6
5.3	Finland	6
5.4	Iceland	6
5.5	Sweden	7
6	Approaches to support the NLPL	7
6.1	Finland	7
6.2	Norway	7
7	Summary	7
7.1	NLPL	7
7.2	Dellingr pilot	8
	Appendix A Author Information	12

1 Executive Summary

The NeIC Dellingr project is investigating how a lightweight framework for sharing **H**igh **P**erformance **C**omputing (HPC) resources¹ can be implemented between participating countries². These resources will be open to eligible researchers³ from the participating countries who wish to access resources in other participating countries. A feature of this resource sharing includes the case where the computing project is performed in an HPC centre outside the home country of the researcher.

A resource-sharing pilot was run from September 2017 to June 2018 in which researchers within the participating countries could apply for a limited amount of computing. The resources in the pilot were assigned on a “first-come first-served” basis once the requests were verified to come from bone-fide academic applicants.

This document details describing the approaches that have been adopted in the national computing providers in each country to support the Dellingr pilot project and the NLPL. In the phase 1 of the Dellingr project, the Deliverable 2 [1] described the processes in place for users to access and use resources at each national provider. Also, this document summarizes the possibilities and issues related to sharing High Performance computing resources within the participating countries.

2 Purpose of the Document

The purpose of this document is to describe, to NeIC and the national providers of HPC computing e-infrastructure in the participating countries, the work necessary to share computing resources across borders. This document is to be used as a basis for further discussions on HPC resource sharing leading to a general agreement.

3 Dellingr Pilot

In the Dellingr project Phase 1 a resource-sharing pilot was proposed and approved, see Phase 1 Delivery Object 1 [2]. The first pilot was a test to ensure that resources could be efficiently shared across the national boundaries of: DK, EE, FI, IS, NO⁴, SE. Also, the procedures for the national HPC providers were tested as given in the Phase 1 Deliverable 2 [1]. The initial amounts of resources donated to the pilot are given in [1], these were subsequently increased in some cases⁵.

¹A resource is defined in this instance as the compute, network and storage infrastructure of a shared system.

²Currently this includes the Nordic countries and Estonia.

³At the moment we do not address commercial users that some countries have agreements to provide/use resources.

⁴Norway did not directly contribute resources to the pilot but Norwegian researchers were granted access to pilot resources.

⁵Estonia joined in December 2017 with approximately 1M core-hours.

The representatives of the National e-infrastructure providers informed research communities and groups within their countries via email about the existence of the pilot and opportunity to apply for computing resources. The rule for this pilot is that prospective applicants had to be eligible to run on their national e-infrastructure provider resources in order to be considered to run abroad. These potential users were directed to an explanatory page on the NeIC Dellingr welcome page [3] which led to a Typeform application page [4]. The process for users to apply for resources in this pilot was as follows:

1. Users applied to the Typeform page to gather basic information on the project: Name, contact details, affiliation, subject, resource request, special resources required;
2. The Typeform page forwarded the request to the Project Manager (PM);
3. PM consulted the national provider reps to host the request in a country other than originator;
4. PM consulted whether shared resources had been consumed fully;
5. If agreed, the user was put into direct contact with the national provider;
6. User applied using the national provider procedure(s);
7. User (or their delegate) gains access to shared resource and is then free to setup and run their software.

The accounting of shared resources was done by the national providers using their own procedures. Overall, the first Dellingr pilot served 20 users with 2.4M core-hours⁶ granted over 9 months.

There were 23 applications in total, 3 were not eligible for the pilot as their hardware requirements could not be met or the application originated from a non-eligible institution. In more detail: one application required a minimum of 1.5 TB of RAM and originated from an institute funded by a different ministry than that of their national e-infrastructure provider. In this case they were not, at the time, eligible to run on their national e-infrastructure provider and therefore could not “avoid” the rules by running in another country under this pilot. Two applications were received from a government-funded research institute (not academic) and therefore did not meet the criteria of the pilot to run abroad.

4 Nordic Language Processing Laboratory

The NeIC Nordic Language Processing Laboratory [5] (NLPL) project has installed software and data collections on two HPC systems: in Norway on Abel [6]; and in Finland on Taito [7]. The

⁶A process consumes a core-hour if it runs for a wall-time hour on a single core of a CPU.

guaranteed computing allocations are 1 million core-hour on Abel and three million on Taito per year, over the current project three-year period of 2017–2019. The on-line storage on Abel has been extended to 2 TB.

On Abel the project has 1 million CPU-hours over the project period (3 years). On Taito the project has 3 million CPU-hours over the project period.

There are different mechanisms in play for administrating these allocations. On Taito, a local project is created for each Principal Investigator which requests NLPL compute resources. Individual users are then added to these local projects by confirmation of the PI. All the paperwork is handled by CSC. In Norway there is one project with one PI who adds the users who will have access. Usage statistics are gathered with the tools which are available at CSC [8] and UNINETT Sigma2 [9].

In addition the project has storage resources: 2 TB on Abel; 25 TB on Taito; and, as a back-up solution, 60-100 TB on the Norwegian National Infrastructure for Research Data (NIRD) [10]. The NIRD storage is detached from the HPC-resources, hence a transfer mechanism (scp,rsync or similar) must be in place to make use of the largest storage. The “storage layout” is a challenge for the project due to:

- A missing abstraction layer, the project must implement its own transfer mechanisms for making use of the largest storage area;
- Though large in total size, the somewhat small volume on Abel means not all data on Taito is available on Abel. Data collections are replicated, but there are indices and metadata which cannot be stored on Abel.

5 Approaches to support the Dellingr pilot

5.1 Denmark

Several projects were delegated to Denmark. All projects were allocated to run on Abacus 2.0 at The DeIC National HPC Centre [11], SDU at the University of Southern Denmark.

For most projects, the amount of resources allocated were reduced to allow for more projects to gain access. Projects were allocated in the same way as local Danish/SDU projects, except that users were required to login to the Abacus 2.0 admin web page using eduGAIN [12] instead of WAYF [13]. For users coming from a previously unknown university, eduGAIN had to be setup but otherwise this worked nicely.

For one of the pilot projects, a crucial part of the project was to be able to access licensed software (MATLAB). After additional effort, it was determined that the already available licenses at Abacus could also be used by the pilot project.

5.2 Estonia

Four projects were allocated to Estonia. All the users were able to log in using eduGAIN and self-provision required allocations and accounts through Estonian self-service [14], based on Waldur [15]. 3 of the 4 users obtained accounts but did not perform any actual work computing, the Danish user has jobs still running. All of the accounts were allocated to the computing centre at the University of Tartu.

5.3 Finland

Five projects were assigned to Finland. They were run at CSC - IT Center for Science. The users from the Dellingr pilot could choose freely what services to use – HPC, cloud or storage. One project was given less resources than they requested. Others were supported in full. The procedure for provisioning accounts was the following. First the Finnish representative of the Dellingr team (CSC staff) asked the users to register at CSC. When they had registered, the project was created manually and the users were added.

Resources for the projects were allocated in billing units (BU), which is a common currency for all resources allocated by CSC. The CPU grants were converted to BUs using the conversion factor 2.0, so that for each core-hour granted they got two BUs. Standard HPC resources consume two BUs per core-hour. Specialized systems such as high-memory nodes or GPUs, consume more BUs per unit time. All users had a standard allocation of “personal” disk space to store their data. However, for one project this was not enough, so project space was allocated for them. This, as outlined above, consumed their resources with a rate of 3.5 BU per TB per hour.

The resource usage was based on Slurm [16] logs and it was monitored regularly using the local accounting reporting system. The projects used resources variably: 0 %, 3 %, 100 %, 101 % and 103 % of the granted allocations.

5.4 Iceland

Two pilot applications were delegated to the Icelandic HPC providers during the pilot phase. Resources requested by one of the projects were renegotiated to a lower amount to better coincide with the resources pledged of the Icelandic HPC center [17] (IHPC).

Since the account creation process on the IHPC systems is still quite rudimentary, utilizing no federated authentication, the account creation process was no different from any local user. In short the PIs sent emails to the IHPC system administrator requesting accounts and they were in turn supplied logins for them and other project members as needed. Further support was provided as requested which mostly was limited to getting acclimatized to the IHPC environment. Resource utilization was tracked using the standard Slurm reporting tools.

5.5 Sweden

Five pilot applications were delegated to the Swedish SNIC [18] HPC resource at LUNARC [19]. All projects were guided towards the project and account creation portal SUPR [20], used for all Swedish HPC projects within SNIC. The procedures for granting the non-Swedish users access to LUNARC resources did not have to be modified from those described in [1]. Resource utilization was tracked using the standard Slurm reporting tools and reported back to the SNIC portal SUPR.

Of the five applications, one (from DK), was deemed not eligible as the applying entity was not a Danish university institute but rather a private one. Two of the assigned projects (DK, FI) failed to apply for the assigned resource for unknown reasons. The final two projects (DK, FI) have been using the assigned resource moderately with an initial burst of usage and with the usage tapering off as the project progressed.

6 Approaches to support the NLPL

6.1 Finland

Five pilot applications were delegated to Finland to be run on HPC at CSC. All projects were allocated to the Taito super-cluster or the Pouta [21] cloud environment. The account creation process was the same as for other users working outside Finland: First the NLPL team member of CSC asked the users to register for CSC. When they had registered, the project was created manually and the users were added.

6.2 Norway

In Norway the NLPL is only one research project in the NOTUR [22] context. Stephan Oepen is the project manager for the project. Users are added and removed to the project by the project manager. Hence, there is Stephan which is the resource manager. Once a user has access to the project resources, the user may consume all resources as they find fit. This can lead to over-consumption, or to much resources used in one direction. There is no resource control or mechanism for resource control on the user level. So far this has not been a problem.

7 Summary

7.1 NLPL

The NeIC NLPL project demonstrates that resources can be shared across borders when used for a specific purpose (language processing project). The national procedures are respected within NLPL

by arranging beforehand for national PIs to act as a host proxy for “foreign” users. The eligibility of users within NLPL is predetermined as they are, by default, members of the project.

For the storage issues, maintaining parallelism of the two instances of the NLPL environment is challenging due to differing storage philosophies: the generous project directory on Taito is not backed up; whereas the much smaller space on Abel is backed up. The project created back-up scripts (nightly rsync of everything to NIRD, on both Abel and Taito) and nightly rsync replication scripts of selected data sub-directories from one of the two systems to the other one. The NLPL project is moving towards maintaining the master copy directly on NIRD and replicating (parts of it) onto both Abel and Taito. This has cost the project almost one person month to develop these strategies and maintain them. This has been accomplished since a NeIC project with such a large and diverse membership inevitably has participants with practical computing knowledge.

The project has the ambition to grow in terms of sites/systems that offer the NLPL software and data stack. This should reduce the need for backup. As NLPL starts to use new systems, the project may be able to use the NIRD storage in a better way. This is a task the the project probably needs to solve. The NLPL project notes that there is room for more general and high-level solutions to sharing data resources across multiple systems.

7.2 Dellingr pilot

Overall the first pilot project of Dellingr Phase 2 proved that HPC resources can be shared between national providers and additional work required by national HPC providers proved to be minimal to accommodate this project.

As for the scalability of the resource-sharing scheme used in this pilot. The process used in this pilot, described in Section 3, to share resources is scalable to some level. The frequency of the applications was such that one person (in this case the PM) was able to deal with them. If a larger resource-sharing scheme was envisioned, with multiple applications per day, then this process would not scale. More automation for requesting and up-to-date accounting of resources and a more distributed assignment process involving national provider members would be required. If there are a large number of applications then the resources to be shared will run out (assigned) before any long-term scalability issues are seen.

Possible options to streamline future resource sharing projects include communicating clearer with possible applicants eligibility rules and advocating for a more widespread usage of federated logins.

A more widespread usage of federated logins (eduGAIN and/or other federations) would ease determining the affiliation and eligibility of the applicants. One example, seen in this first Dellingr pilot, is determining the status of an applicants’ institution i.e. is it a publicly-funded university or a private institute. This information is needed when making a decision on where the computing work can be hosted. Proprietary software licensing rules either need to be investigated further, as made

clear by one of the project utilizing the Danish HPC providers. An option could be to reduce the the scope of future resource sharing projects to accommodate only free and open software.

References

- [1] Dellingr Phase 1 Deliverable 2. URL: <https://wiki.neic.no/wiki/Dellingr#D02>, doi:10.5281/zenodo.3965751.
- [2] Dellingr Phase 1 Deliverable 1. URL: <https://wiki.neic.no/wiki/Dellingr#D01>, doi:10.5281/zenodo.3964423.
- [3] NeIC Dellingr Welcome page.
<https://dellingr.neic.no/>
- [4] Typeform.
<https://www.typeform.com/>
- [5] NeIC Nordic language processing laboratory project.
<https://neic.no/nlpl/>
- [6] Abel supercomputer.
<https://www.uio.no/english/services/it/research/hpc/abel/more/>
- [7] Taito supercomputer.
<https://research.csc.fi/taito-supercluster>
- [8] CSC.
<https://www.csc.fi/>
- [9] Uninett Sigma2.
<https://www.uninett.no/en/sigma>
- [10] National Infrastructure for Research Data.
<https://documentation.sigma2.no/storage/nird.html>
- [11] Abacus.
https://www.deic.dk/en/HPCCenter_SDU
- [12] eduGAIN - enabling worldwide access.
<https://edugain.org/>
- [13] Where are you from (WAYF).
<https://www.wayf.dk/en>
- [14] Estonian self-service.
<https://minu.etais.ee>

- [15] Waldur cloud brokerage.
<https://waldur.com/>
- [16] Slurm workload manager.
<https://slurm.schedmd.com/>
- [17] Icelandic High Performance Computing.
<http://ihpc.is/>
- [18] Swedish National Infrastructure for Computing.
<https://www.snic.se/about/>
- [19] Center for scientific and technical computing at Lund University (LUNARC).
<http://www.lunarc.lu.se/>
- [20] SNIC User and Project Repository.
<https://supr.snic.se/>
- [21] CSC Pouta cloud services.
<https://research.csc.fi/cloud-computing>
- [22] NOTUR project.
<http://www.notur.no/>

Appendix A Author Information

Name	email address	ORCID
Mathias Brännvall	mathias.brannvall@it.uu.se;	0000-0003-4979-4123
Juha Fagerholm	juha.fagerholm@csc.fi	0000-0002-9972-4468
Jens Svalgaard Kohrt	svalgaard@sdu.dk	0000-0002-3104-0406
Ivar Koppel	ivar.koppel@ut.ee	0000-0002-5617-4785
Bjørn Lindi	bjorn.lindi@ntnu.no	
Ilja Livenson	ilja.livenson@ut.ee	0000-0002-4011-8367
Petri Nikunen	petri.nikunen@csc.fi	0000-0003-0759-6372
Anders Sjöström	Anders.Sjostrom@lunarc.lu.se	0000-0003-2213-2138
Hjörleifur Sveinbjörnsson	hs@hi.is	0000-0002-4120-1234
John White	john.white@cern.ch	0000-0001-5614-0895