

Reimagining research computing

Wednesday, 22 May 2019 13:30 (1h 30m)

In modern times, computation power is becoming more and more important. However, at the same time, the rest of the world is becoming consumerized: while the general expectation is that information technology is easier to use, the design of high-performance computing (HPC) systems has not kept up with modern developments in computer usability. There are many historical artifacts of how HPC systems are set up: HPC systems are often optimized for data transfer over scp, while users often prefer solutions where remote drives are mounted. We expect computations to fit into nice “rectangular” boxes of number of cores \times time \times memory, while with modern data science workflows, the time and memory can be unknown at the start of a job, and, in particular, interactive usage leads to highly intermittent CPU and memory requirements. Why is knowing Linux shell scripting a requirement for every job when we want our facilities to be usable by anyone? How can we empower users to have more control over their software stack?

In this workshop, we will explore the largest usability barriers in HPC systems, existing solutions, and create a joint vision of a modern HPC system. The first talks will be presentations on vision and usability from invited speakers from both HPC and human-computer interaction (HCI). After that, there will be brainstorming sessions (guided, in small groups, unconference, or panel discussions) where we identify the biggest pain points. Then, there will be group discussions in a speed-blogging format to create a shared vision document which will be the result of this workshop. After this workshop, there should be additional Nordic infrastructure co-operation to improve the accessibility, and possibly standardization, of large computational resources beyond those who traditionally use them.

“Homework”: This is an interactive workshop, so please come prepared. Talk to people at your institution and/or other meeting at NeIC. Poll the people around you: what are the biggest issues with using your institution’s computational facilities? Issues can be both general and specific, e.g. “all files have to manually be transferred, but due to the use of ssh proxy hosts there it is difficult from outside the campus network” or “it is easier to pay Amazon than pay us”.

Workshop outline

- **HPC and accessibility** (30 min, speaker TBA): Inspiring talk about why accessibility of resources is important.
- **Strides towards accessibility at Aalto University** (10 min, Richard Darst): Demo of some of our recent ideas and problems we have seen.
- **Strides towards accessibility** at [TBA] (10 min, speaker TBA)
- **Unconference planning**, divide into layers and discuss (10 min)
- **Unconference** (30 min)
- lunch break
- **Unconference** (30 min)
- **Presentations by groups** (30 min)
- **Panel follow-up** (30 min) (panel members TBA)

Notes

- Resources first or usability first. Traditional computing center strategy is resources first, then think about how to use them. Google/Amazon strategy is make them usable, then scale up (credit: jh)
- Google/AWS cluster as a service as a baseline.
- Intro talk points: we know that there are solutions, but they are all different. Not as useful to have many different solutions. Life isn’t just HPC, but it is a standard baseline.

Primary authors: BAST, Radovan (UiT/NeIC); DARST, Richard (Aalto University); RAZICK, Sabry (UiO/NeIC); WIK-FELDT, Thor (KTH/NeIC)

Presenter: DARST, Richard (Aalto University)

Session Classification: Workshops I