

NDGF Site Report

2016-10-18, HEPiX Fall 2016, Berkeley

Mattias Wadenstein



norden

NordForsk



Nordic e-Infrastructure
Collaboration

Overview

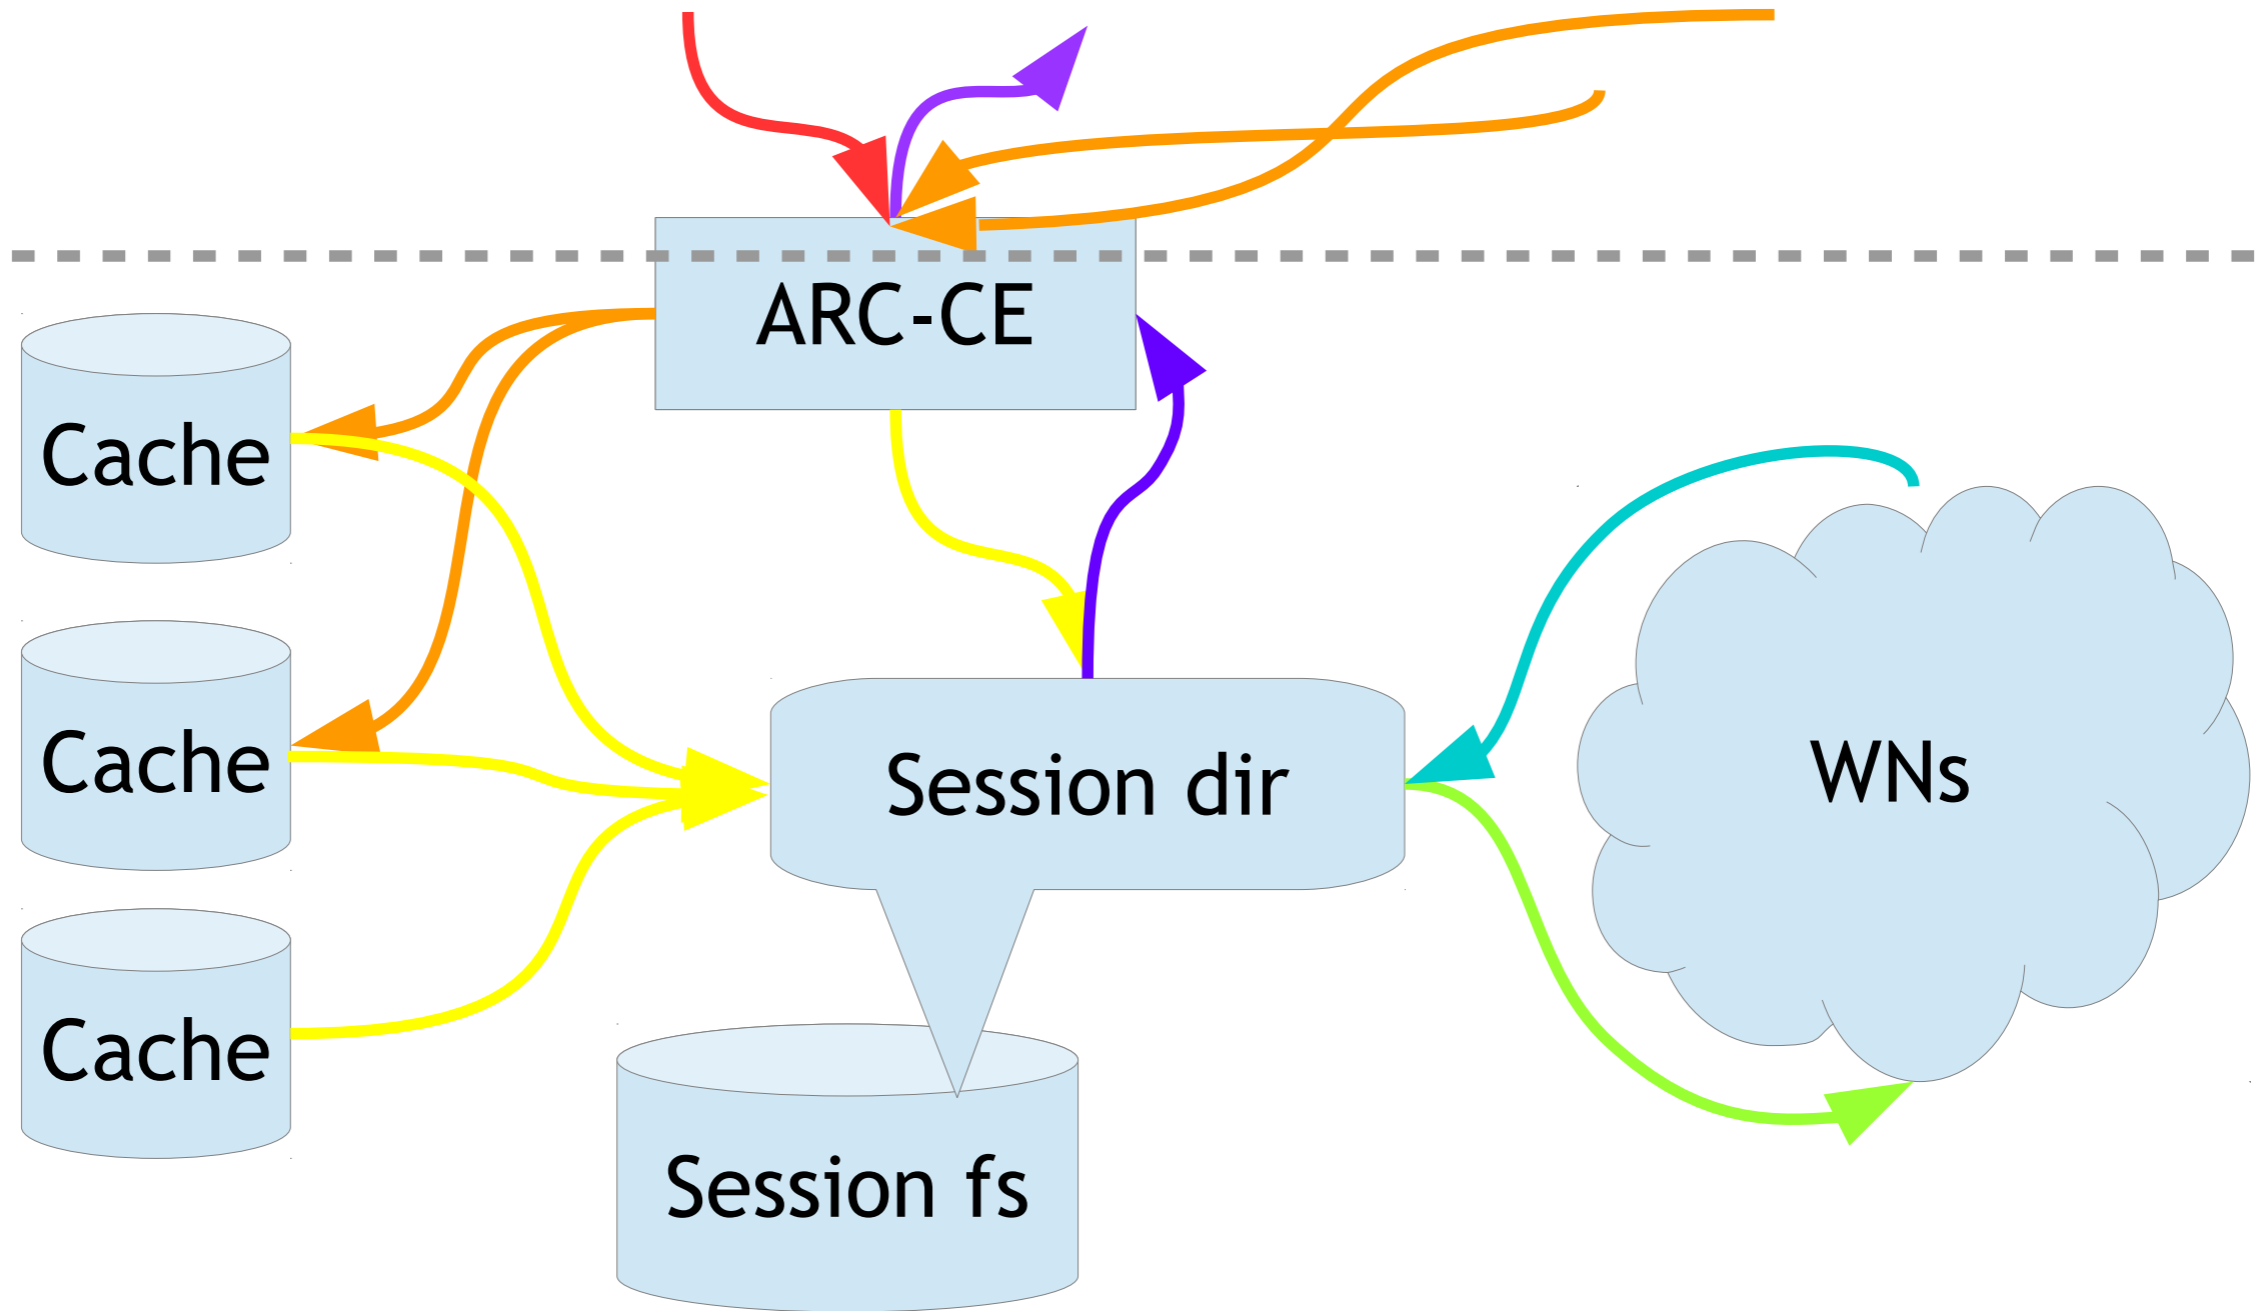
- NDGF overview
- Network upgrade
- Staff changes
- dCache update
- Site reports
- NeIC 2017



NDGF Overview

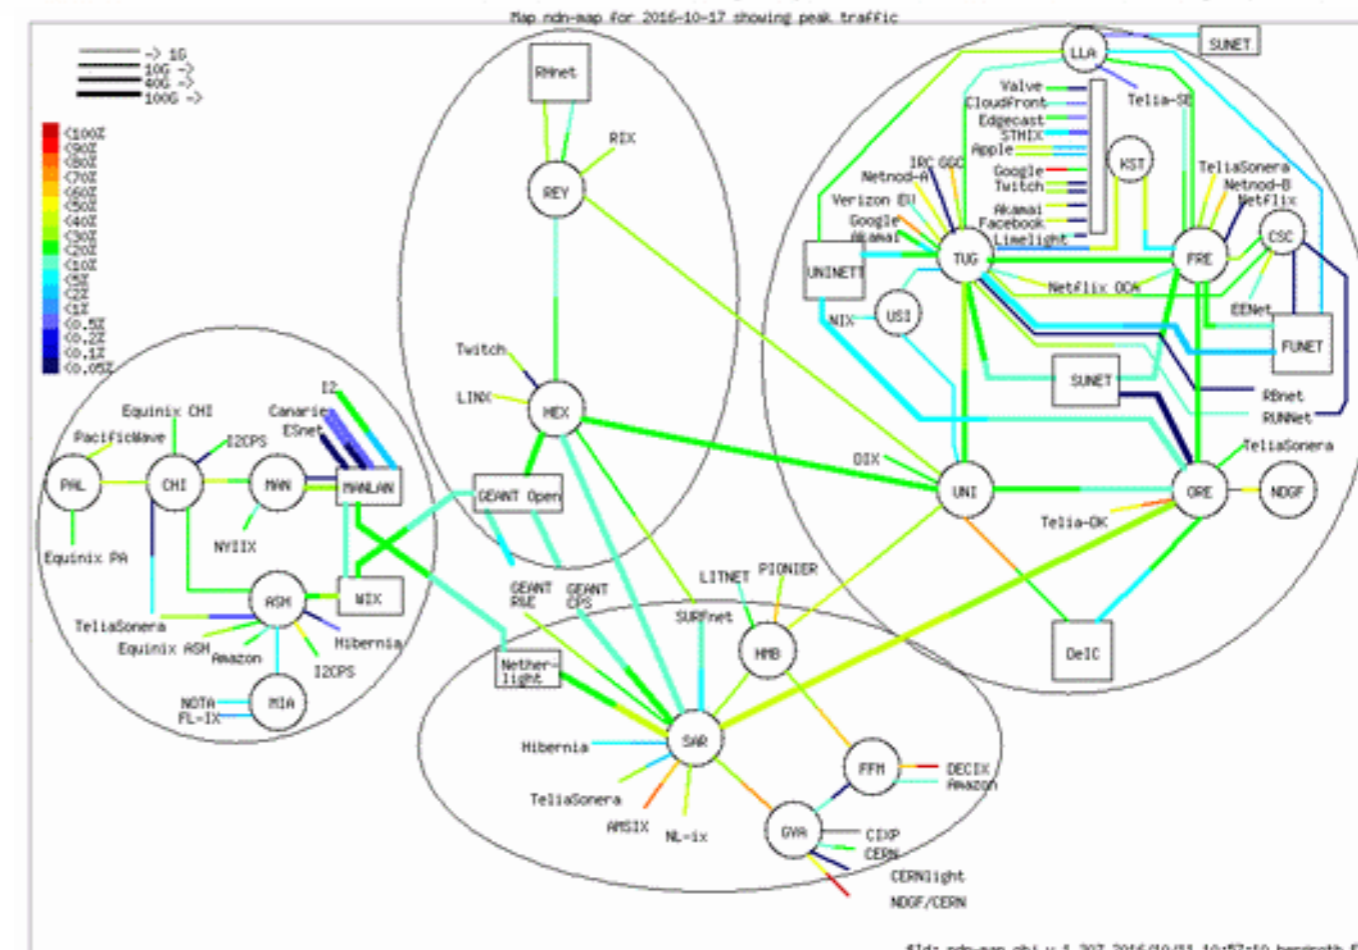
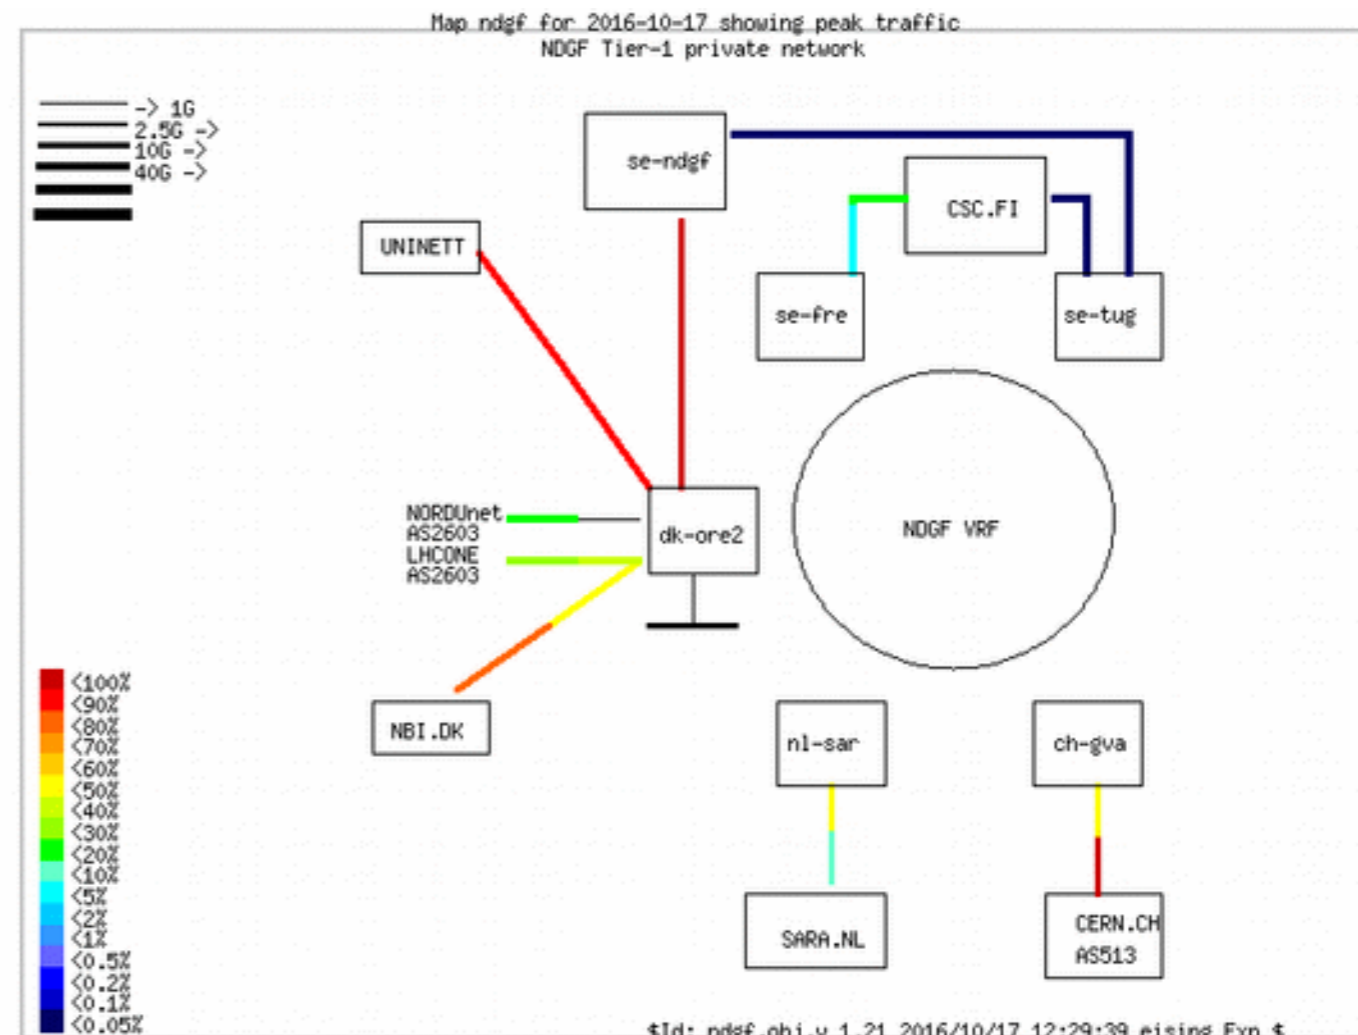
- Distributed WLCG tier-1 site
 - 6 Nordic academic HPC sites with dCache pools and ARC-CEs
 - And disk from IJSs T2 in Slovenia
- Supports ATLAS and ALICE
 - Funding agency targets:
 - 6% of ATLAS tier-1 resources
 - 9% of ALICE tier-1 resources
- Thanks to the disk space consolidated into the tier-1 from the tier-2s SI-SIGNET-T2 and SE-SNIC-T2 we currently have second largest ATLASDATADISK among all tier-1 sites

NDGF Overview: ARC-CE with data staging




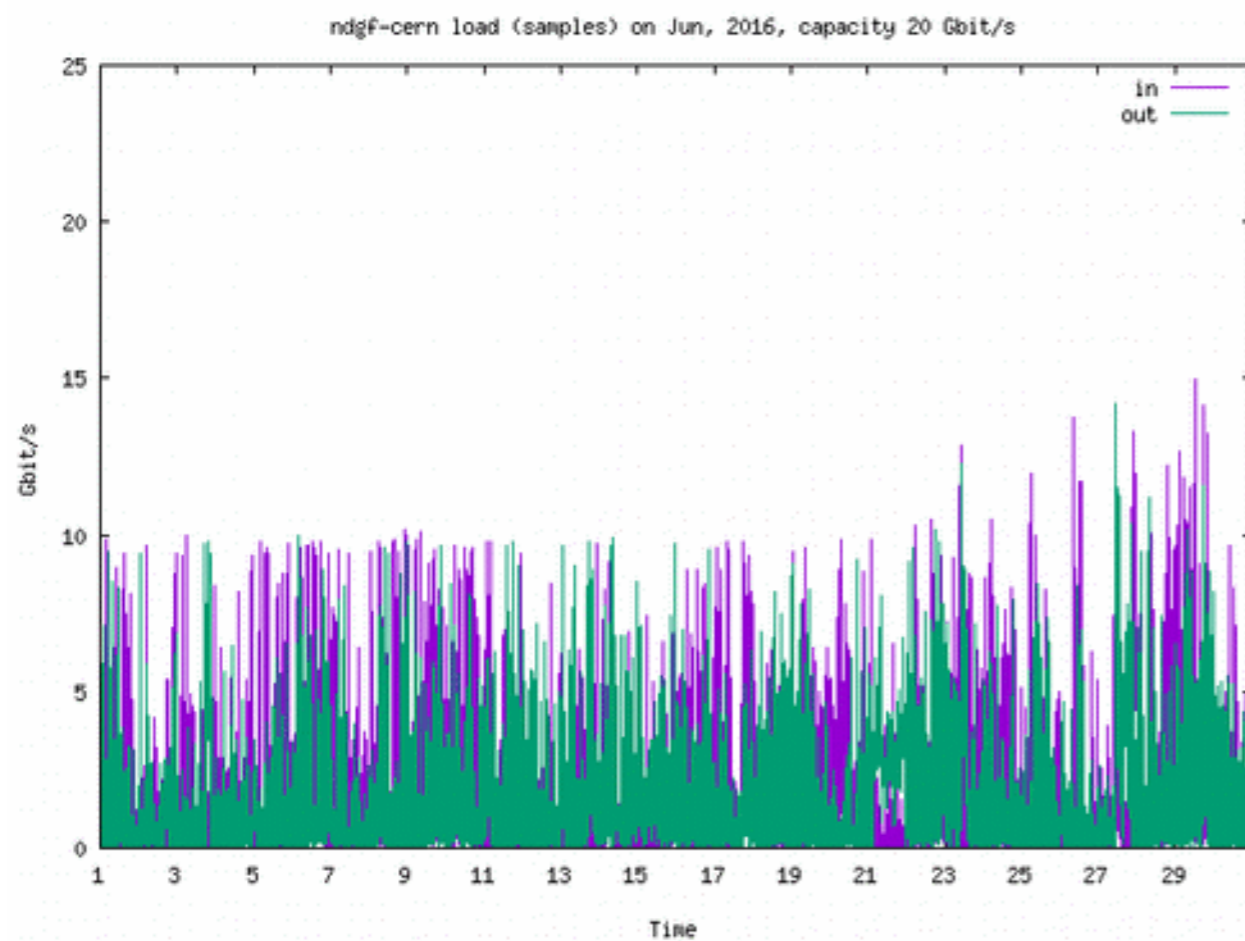
Network upgrade

- Moving the LHCOPN away from dedicated optical links
- New solution is using VRF and MPLS over the “normal” links in NORDUNet
- Typically 20Gbit/s guaranteed to each country, bursting above that “less than best effort”



Network upgrade

- Added another 10Gbit/s to the CERN link for 20Gbit/s total



Staff changes as NDGF-T1

- Jon Kerr Nilsen, our ARC developer moved on to new challenges. Replacement hired: Maiken Pedersen
- Gerd Behrmann, our dCache developer is moving on to new challenges. Replacement position open on www.neic.no - if you know someone that'd fit in his shoes and is interested, please point them our way.

dCache update

- We are now running dCache 3.0 in production with redundant headnodes and haproxy + ucarp in front of most doors
- See separate talk on this tomorrow
- Starting to move more and more aspects of the central nodes into ansible for management

Site reports: University of Oslo

- Twice the space for ARC grid cache for data staging
- Lessons learnt: Not all cables that can fit physically are equivalent, took days to figure out that the ARC CE frontend needed a Cisco branded infiniband cable.



Site reports: CSC


- Network upgraded to 2x10G with WAN transport on a virtual (MPLS) link on the NORDUNet 100G links.
- NDGF-T1 hardware in Finland has been renewed from:
 - Half a rack of DL360G7 + D2600 shelves
 - to
 - 3x HP SL4510 Gen9
- For hot pix, see Johan Guldmyr's upcoming T2_FI_HIP site report, same hardware

Site reports: HPC2N

- New HPC cluster: Kebnekaise



Site reports: HPC2N



Site reports: HPC2N

- Total budget including staff 75MSEK (~8.5MUSD) for 4.5 years



Node Type	#nodes	Cores	Memory
Compute	432	2x14	128 GB/node
LargeMem	20	4x18	3072 GB/node
2xGPU	32	2x14 + 2xK80	128 GB/node
4xGPU	4	2x14 + 4xK80	128 GB/node
KNL	36	68	192 GB/node
=SUM in 13 racks	544	17552 + 80xK80	125 TB

Site reports: HPC2N

- FDR and EDR Infiniband
 - LargeMem and switches on EDR
- Fairly power dense
 - 2x32A 3-phase per rack
- Upcoming recruiting
 - Know someone that wants to admin Kebnekaise in Umeå, northern Sweden? Talk to me.
 - 1 sysadmin & scientific software
 - 1 sysadmin & storage (lustre)



NeIC 2017 Conference Umeå 2017-05-29 - 2017-06-01

- 2 days of workshops
 - dCache, sensitive data, data management, KNL programming, ...
- 2 days of conference
 - On the e-Infrastructure needed to enable science
 - Sessions on: Society, Science, Technology, Industry, Online computing
- Excellent location :)
- More info on: <http://neic2017.nordforsk.org>

Questions?

