# Operations in NDGF-T1 and SE-SNIC-T2

Erik Edelmann

Nordic e-Infrastructure Collaboration (NeIC) / CSC – IT center for Science

ALICE Tier-1/Tier-2 workshop 2019

# NDGF & NeIC: Background

- Nordic DataGrid Facility (NDGF) was founded 2002 to coordinate the cooperation between Finland, Sweden, Norway and Denmark to create a Tier-1 for ATLAS and ALICE.

# NDGF & NeIC: Background

- Nordic DataGrid Facility (NDGF) was founded 2002 to coordinate the cooperation between Finland, Sweden, Norway and Denmark to create a Tier-1 for ATLAS and ALICE.
- A few years ago, NDGF was reorganized into NeIC (Nordic e-Infrastructure Collaboration)
  - Maintaining NDGF-T1 is a project within NeIC

# NDGF & NeIC: Background

- Nordic DataGrid Facility (NDGF) was founded 2002 to coordinate the cooperation between Finland, Sweden, Norway and Denmark to create a Tier-1 for ATLAS and ALICE.
- A few years ago, NDGF was reorganized into NeIC (Nordic e-Infrastructure Collaboration)
  - Maintaining NDGF-T1 is a project within NeIC
- In addition to the NDGF-T1, there's a few T2:s in the Nordics
  - FI-HIP-T2 for CMS
  - SE-SNIC-T2 for ALICE and ATLAS
  - Sometimes hard to distinguish from NDGF-T1.

# NDGF & NeIC: Background

- Nordic DataGrid Facility (NDGF) was founded 2002 to coordinate the cooperation between Finland, Sweden, Norway and Denmark to create a Tier-1 for ATLAS and ALICE.
- A few years ago, NDGF was reorganized into NeIC (Nordic e-Infrastructure Collaboration)
  - Maintaining NDGF-T1 is a project within NeIC
- In addition to the NDGF-T1, there's a few T2:s in the Nordics
  - FI-HIP-T2 for CMS
  - SE-SNIC-T2 for ALICE and ATLAS
  - Sometimes hard to distinguish from NDGF-T1.
- NDGF-T1 is sometimes also referred to as the NT1 ("Nordic Tier 1")

# ALICE sites in NDGF-T1 & SNIC-T2

# Sites: Finland

- HIP
  - Run by Helsinki Institute of Physics (HIP)

# Sites: Finland

- HIP
  - Run by Helsinki Institute of Physics (HIP)
  - Bunch of VMs, running on CSC's openstack system cPouta

# Sites: Finland

- HIP
  - Run by Helsinki Institute of Physics (HIP)
  - Bunch of VMs, running on CSC's openstack system cPouta
  - A "quick-and-dirty" setup (read: an embarassingly ugly mess of ansible and python scripts), but at least it's operational

# Sites: Finland

- HIP
  - Run by Helsinki Institute of Physics (HIP)
  - Bunch of VMs, running on CSC's openstack system cPouta
  - A "quick-and-dirty" setup (read: an embarassingly ugly mess of ansible and python scripts), but at least it's operational
  - Increased from 296 cores $\times$ 19.54 HS06 to 312 cores $\times$ HS06 19.62 in January

# Sites: Finland

- HIP
  - Run by Helsinki Institute of Physics (HIP)
  - Bunch of VMs, running on CSC's openstack system cPouta
  - A "quick-and-dirty" setup (read: an embarassingly ugly mess of ansible and python scripts), but at least it's operational
  - Increased from 296 cores $\times$ 19.54 HS06 to 312 cores $\times$ HS06 19.62 in January
  - Still on CentOS 6 and SGE

# Sites: Finland

- HIP
  - Run by Helsinki Institute of Physics (HIP)
  - Bunch of VMs, running on CSC's openstack system cPouta
  - A "quick-and-dirty" setup (read: an embarassingly ugly mess of ansible and python scripts), but at least it's operational
  - Increased from 296 cores $\times$ 19.54 HS06 to 312 cores $\times$ HS06 19.62 in January
  - Still on CentOS 6 and SGE
    - TODO: Upgrade to CentOS 7 & Slurm

# Sites: Sweden 1/2

- SNIC
  - National Supercomputer Center, Linköping, Sweden

# Sites: Sweden 1/2

- SNIC
  - National Supercomputer Center, Linköping, Sweden
  - Also takes care of the SE-SNIC-T2 storage.
  - Backend: Slurm
  - HEPSPEC06: 15.925
  - Shared with ATLAS

# Sites: Sweden 1/2

- SNIC
  - National Supercomputer Center, Linköping, Sweden
  - Also takes care of the SE-SNIC-T2 storage.
  - Backend: Slurm
  - HEPSPEC06: 15.925
  - Shared with ATLAS
  - Had a long break during last summer-autumn

# Sites: Sweden 1/2

- SNIC
  - National Supercomputer Center, Linköping, Sweden
  - Also takes care of the SE-SNIC-T2 storage.
  - Backend: Slurm
  - HEPSPEC06: 15.925
  - Shared with ATLAS
  - Had a long break during last summer-autumn
    - The old cluster Triolith "replaced" by "new" cluster Bluegrass

# Sites: Sweden 1/2

- SNIC
  - National Supercomputer Center, Linköping, Sweden
  - Also takes care of the SE-SNIC-T2 storage.
  - Backend: Slurm
  - HEPSPEC06: 15.925
  - Shared with ATLAS
  - Had a long break during last summer-autumn
    - The old cluster Triolith "replaced" by "new" cluster Bluegrass
  - Last few months $\sim$ 600 cores for ALICE, may shrink a bit in the future.

# Sites: Sweden 2/2

- LUNARC
  - Center for scientific and technical computing for research at Lund University, Lund, Sweden
  - SE-SNIC-T2 CE
  - Backend: Slurm

# Sites: Sweden 2/2

- LUNARC
  - Center for scientific and technical computing for research at Lund University, Lund, Sweden
  - SE-SNIC-T2 CE
  - Backend: Slurm
  - $\sim$ 5.3 kHS06 CPUs shared by ALICE and ATLAS.

# Sites: Sweden 2/2

- LUNARC
  - Center for scientific and technical computing for research at Lund University, Lund, Sweden
  - SE-SNIC-T2 CE
  - Backend: Slurm
  - $\sim$ 5.3 kHS06 CPUs shared by ALICE and ATLAS.
    - Combined CPU pledge for ALICE and ATLAS: 7.9 kHS06 CPUs ...

# Sites: Denmark

- DCSC/KU
  - Københavns Universitet
  - Backend: ARC / Slurm
  - $\sim$ 13.4 kHS06 CPUs shared by ALICE and ATLAS
  - Heterogeneous cluster; oldest node $>$ 10 years old.
  - Average HEPSPEC06: 14.43

# Sites: Denmark

- DCSC/KU
  - Københavns Universitet
  - Backend: ARC / Slurm
  - $\sim$ 13.4 kHS06 CPUs shared by ALICE and ATLAS
  - Heterogeneous cluster; oldest node $>$ 10 years old.
  - Average HEPSPEC06: 14.43
  - Will grow in not too distant future, exact schedule unknown.

# Sites: Norway

- UiB (Bergen)
  - Backend: Slurm
  - $\sim$ 6.7 kHS06 CPUs
  - HEPSPEC06: 12.12

# Sites: Norway

- UiB (Bergen)
  - Backend: Slurm
  - $\sim$ 6.7 kHS06 CPUs
  - HEPSPEC06: 12.12
  - Future: New hardware arrived, production use should start soon.
    - 24 Dell compute nodes with dual socket AMD EPYC
    - 25 Dell disk server nodes with 140 TB (raw)
    - Everything 10Gb/s
    - Hardware is part of the Norwegian research cloud installation in Bergen.
    - Tape provided via dedicated TSM machine (4 drives)

# Sites: Norway

- UiB (Bergen)
  - Backend: Slurm
  - $\sim$ 6.7 kHS06 CPUs
  - HEPSPEC06: 12.12
  - Future: New hardware arrived, production use should start soon.
    - 24 Dell compute nodes with dual socket AMD EPYC
    - 25 Dell disk server nodes with 140 TB (raw)
    - Everything 10Gb/s
    - Hardware is part of the Norwegian research cloud installation in Bergen.
    - Tape provided via dedicated TSM machine (4 drives)
- UiO (Oslo)
  - Backend: Slurm
  - 10 cores for ALICE
  - HEPSPEC06: 17.63
  - Future: Will most likely disappear by the end of this year.

# NDGF-T1 CPU resources, compared to pledge

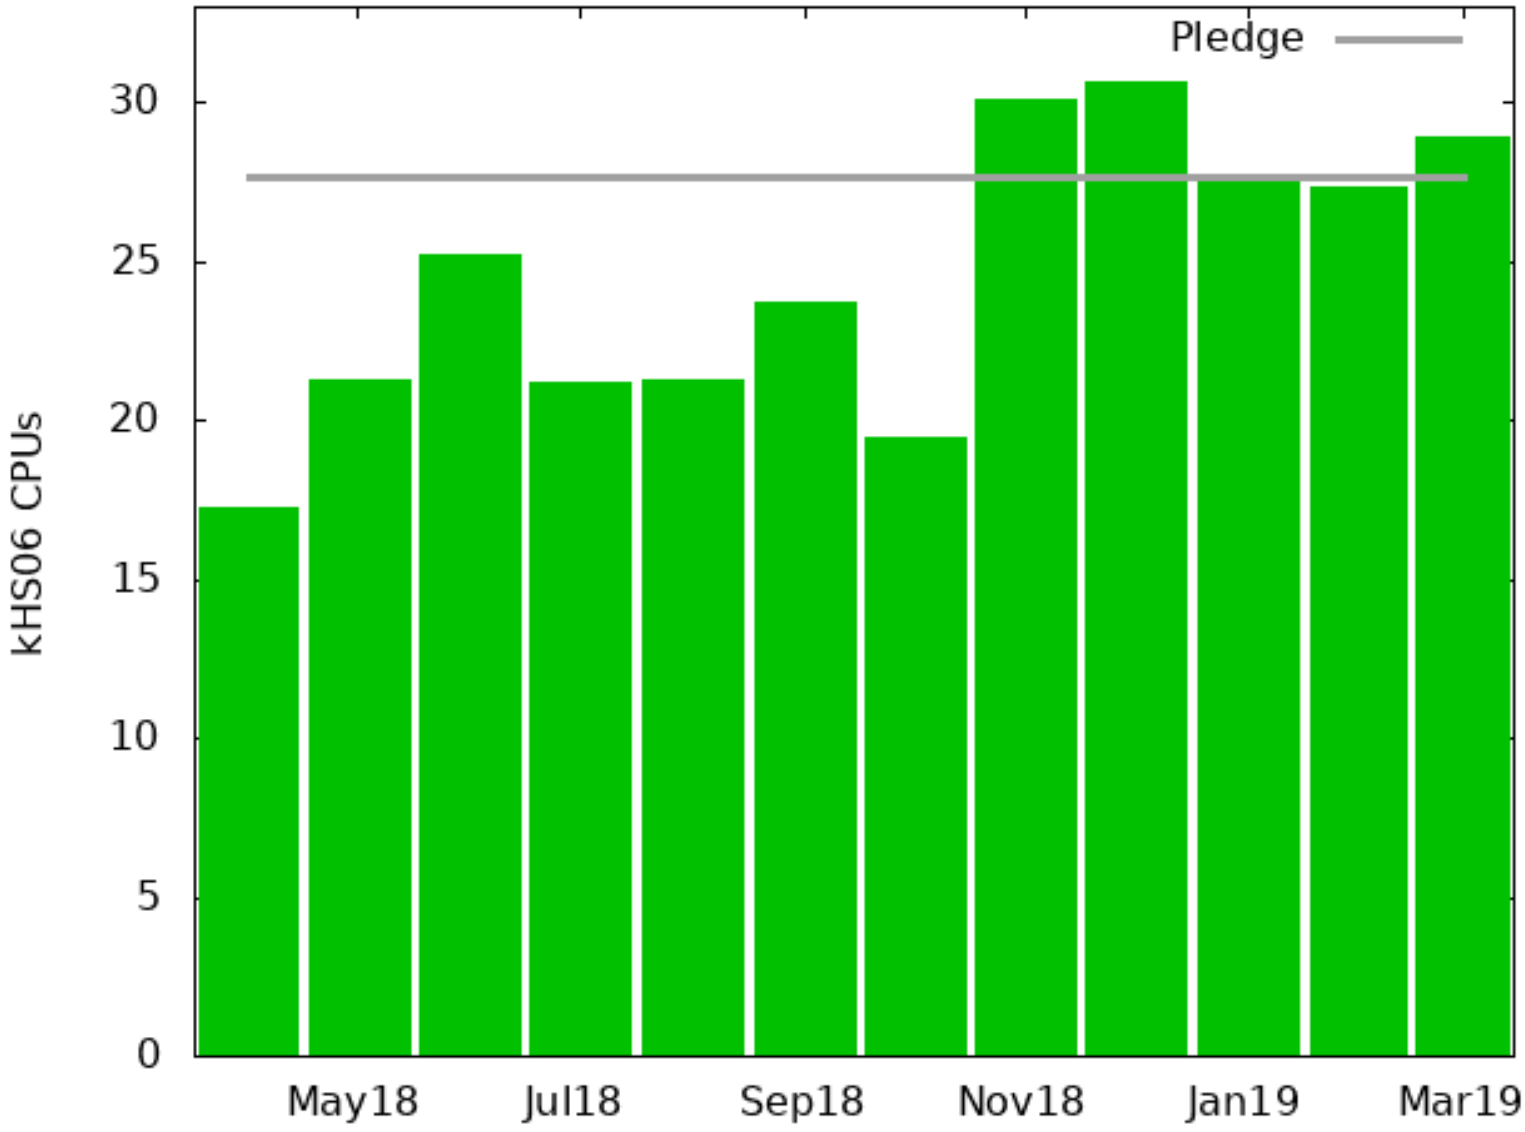

# SNIC-T2 CPU resources, compared to pledge

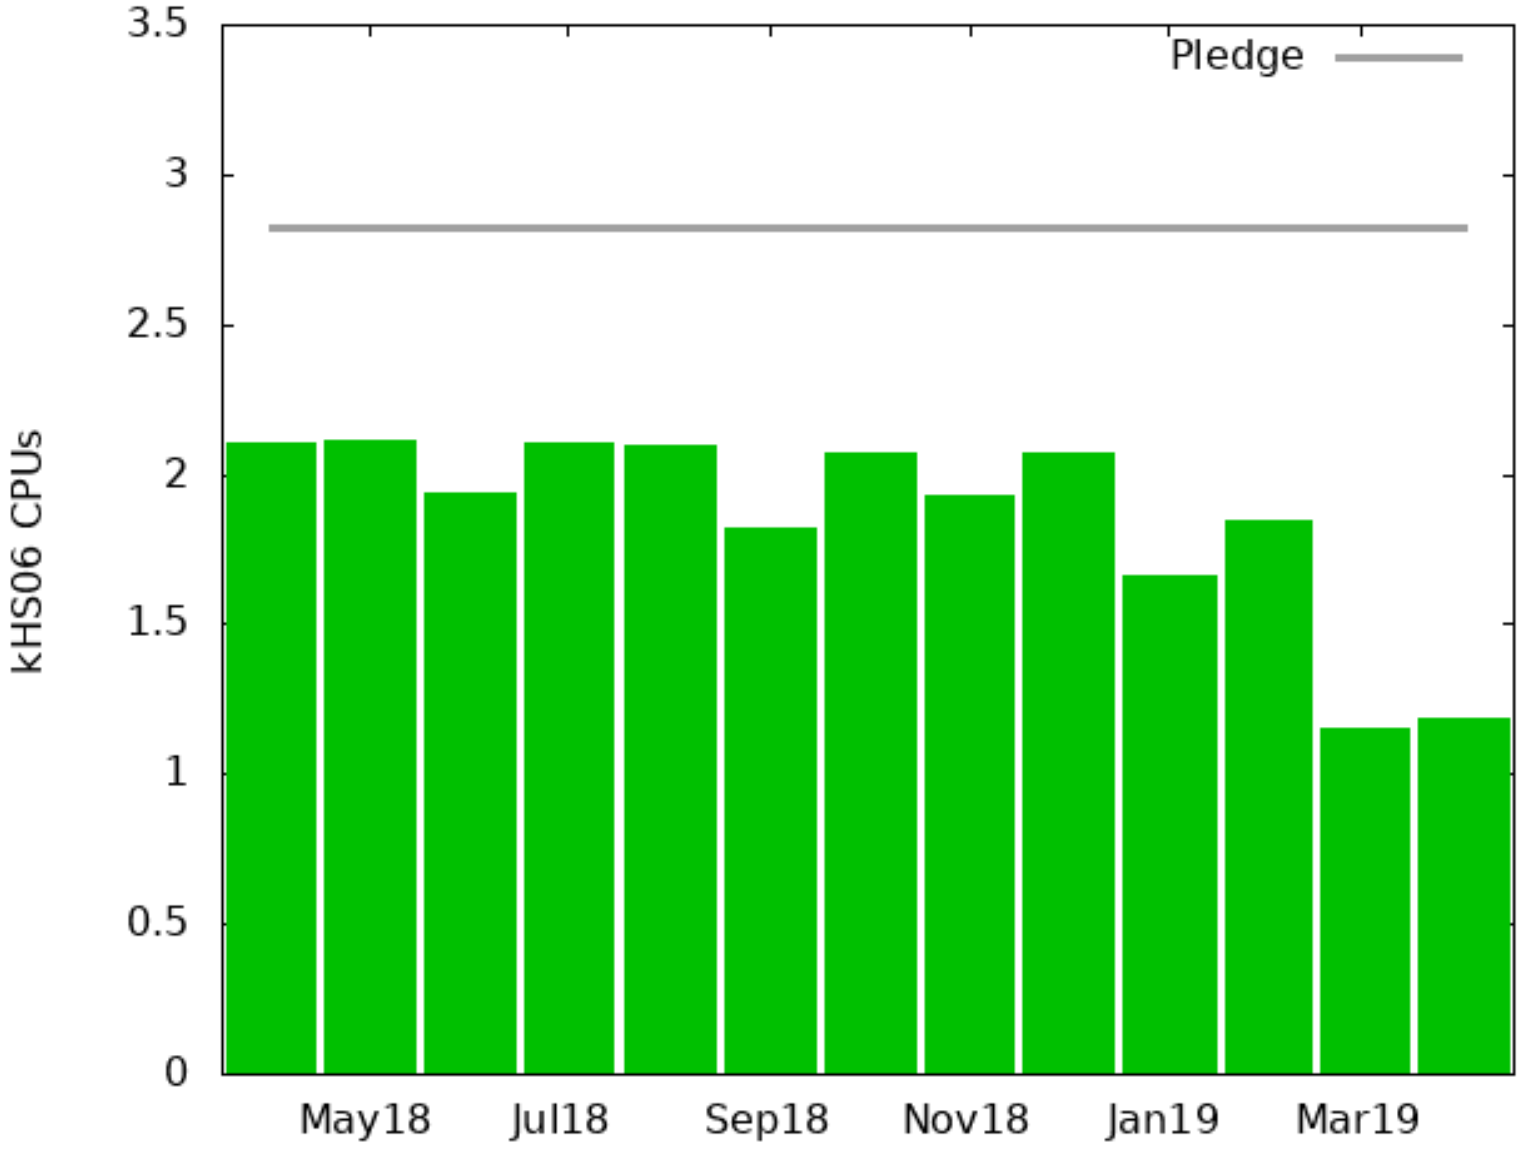

# Storage

- Many small sites looks like one big site, using dCache

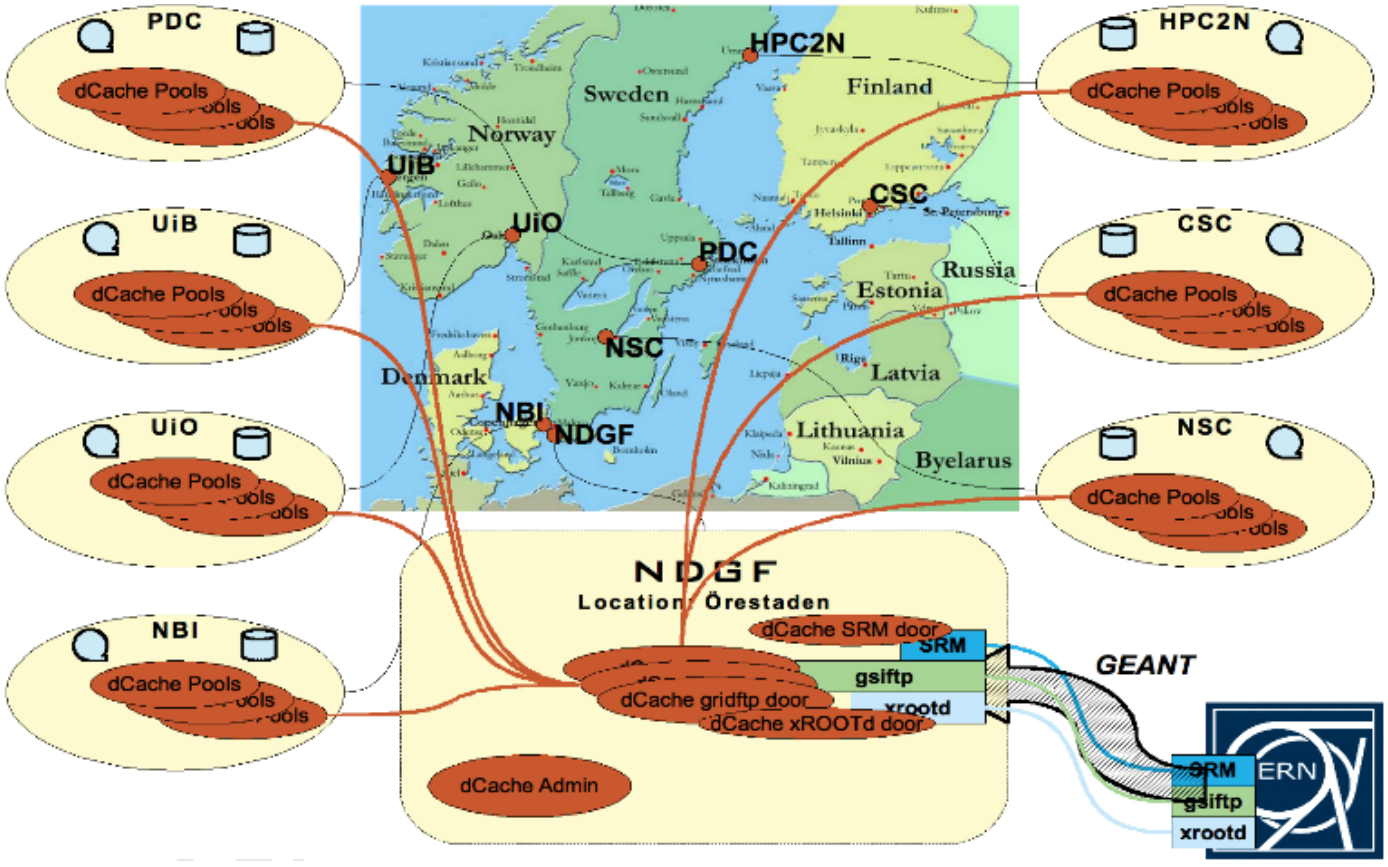


- Lets us hide problems on one site by keeping data on other sites.
  - (at least, almost . . . )

# Storage in numbers

NDGF-T1

- Disk:
  - Pledge: 4240 TB (2510 TB last year)
  - Currently available: 2582 TB
- Tape:
  - Pledge: 2.99 PB (1.86 PB last year)
  - Currently available: 2.02 PB

More a-coming!

# Storage in numbers

SE-SNIC-T2

- Disk:
  - Pledge: 400 TB
  - Currently available: 407 TiB

# Storage in numbers

SE-SNIC-T2

- Disk:
  - Pledge: 400 TB
  - Currently available: 407 TiB
  - Future: Disks may be replaced, but capacity unlikely to change

# IPv6 readiness

- Yes:
  - All our SEs
    - NDGF::dCache
    - NDGF::dCache_tape
    - SNIC::dCache
  - Some CEs
    - SNIC
    - UiO

# IPv6 readiness

- Yes:
  - All our SEs
    - NDGF::dCache
    - NDGF::dCache_tape
    - SNIC::dCache
  - Some CEs
    - SNIC
    - UiO
- Not yet:
  - LUNARC
  - UiB
  - HIP (Not supported by the underlying cloud system yet)
  - DCSC/KU (On the TODO list)

# Questions?