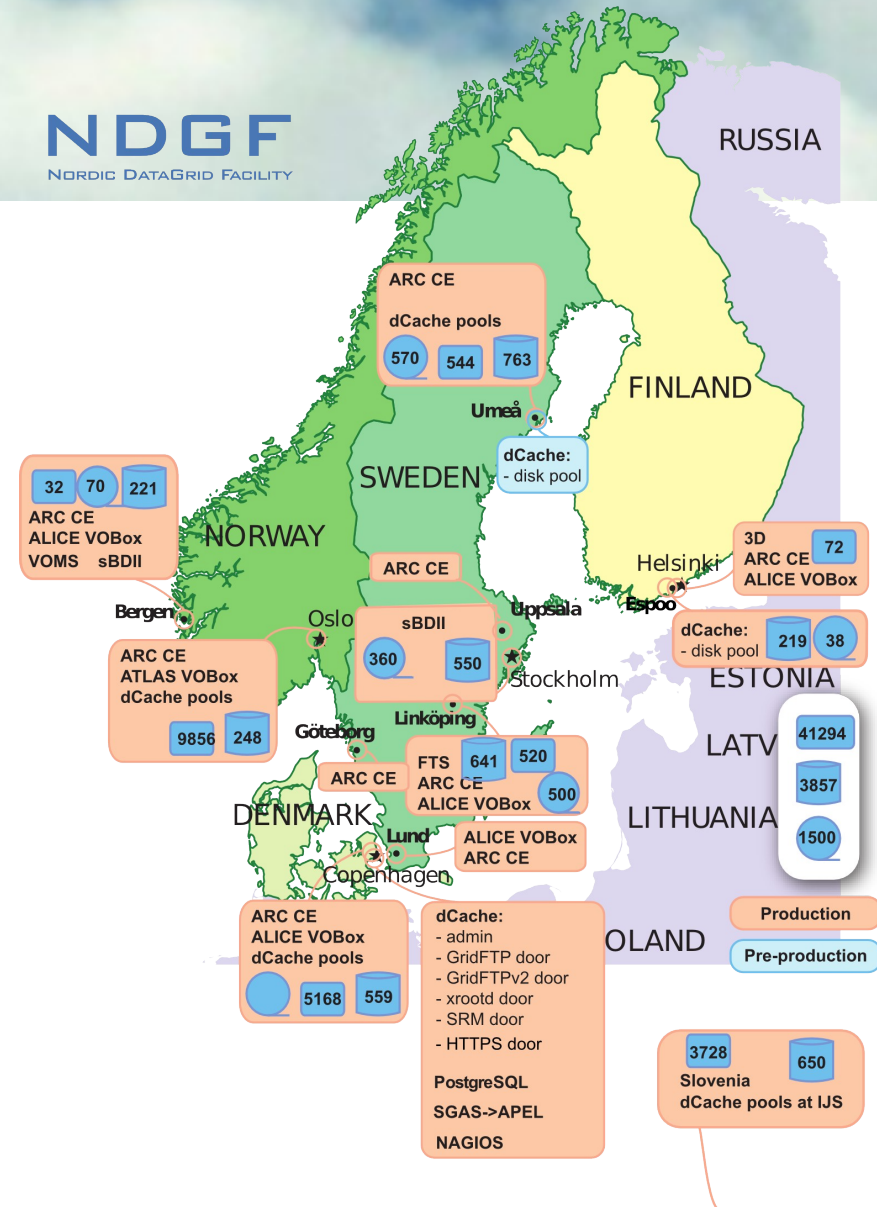




# **NDGF Site Report**

Mattias Wadenstein  
Hepix 2012 Fall Meeting  
2012-10-15, Beijing

- NDGF is a distributed site
- Most of the big HPC centers in Norway, Sweden, Finland and Denmark
- Tier1 storage at 7
  - Plus IJS in Slovenia
- Most numbers 2012
  - CPU is total cpu, not WLCG allocated



- NDGF Tier-1 is one project in NeIC
  - Headed by me
- Nordic e-Infrastructure Collaboration
  - Headed by Gudmund Høst
- Has seen a year of painful reorganization
- Now looking to recruit and build a sustainable WLCG Tier-1
- Currently ~10 people, average 50%

## ATLAS Grid Monitor

2012-10-13 CEST 10:50:31



Processes: ■ Grid ■ Local

Country	Site	CPUs	Load (processes: Grid+local)	Queueing
<span style="color: red;">■</span> <b>Denmark</b>	<b>Steno Tier 1 (DCSC/KU)</b>	5168	<div><div style="width: 100%;"></div></div> 864+2935	<b>1296</b> +1
<span style="color: red;">■</span> <b>Norway</b>	<b>Tier1 (BCCS/UiB)</b>	372	<div><div style="width: 100%;"></div></div> 0+234	<b>0</b> +1
	<b>Titan A (UiO/USIT)</b>	9856	<div><div style="width: 100%;"></div></div> 886+2216	<b>282</b> +8
	<b>Titan C (UiO/USIT)</b>	9856	<div><div style="width: 100%;"></div></div> 0+3103	<b>0</b> +8
<span style="color: red;">■</span> <b>Slovenia</b>	<b>Arnes</b>	1636	<div><div style="width: 100%;"></div></div> 1558+0	<b>1003</b> +0
	<b>SiGNET</b>	2092	<div><div style="width: 100%;"></div></div> 2027+0	<b>684</b> +0
<span style="color: blue;">■</span> <b>Sweden</b>	<b>Alarik (SweGrid, Luna&gt;</b>	3328	<div><div style="width: 100%;"></div></div> 114+2315	<b>52</b> +0
	<b>Grad (SweGrid, Uppmax)</b>	512	<div><div style="width: 100%;"></div></div> 353+0	<b>59</b> +2
	<b>Ritsem (SweGrid, HPC2&gt;</b>	544	<div><div style="width: 100%;"></div></div> 322+0	<b>243</b> +0
	<b>Siri (SweGrid, Lunarc)</b>	512	<div><div style="width: 100%;"></div></div> 332+138	<b>240</b> +50
	<b>Smokerings (NSC)</b>	520	<div><div style="width: 100%;"></div></div> 416+80	<b>2525</b> +0
	<b>Smokerings TEST (NSC)</b>	520	<div><div style="width: 100%;"></div></div> 0+496 (queue inactive)	<b>0</b> +2519
<span style="color: red;">■</span> <b>Switzerland</b>	<b>Bern ATLAS T3</b>	532	<div><div style="width: 100%;"></div></div> 456+0	<b>150</b> +0
	<b>Bern UBELIX T3</b>	1216	<div><div style="width: 100%;"></div></div> 448+0	<b>396</b> +1
	<b>Geneva ATLAS T3</b>	278	<div><div style="width: 100%;"></div></div> 27+193	<b>77</b> +357
	<b>Manno PHOENIX T2</b>	2176	<div><div style="width: 100%;"></div></div> 252+1916	<b>53</b> +403
	<b>Manno PHOENIX T2</b>	2176	<div><div style="width: 100%;"></div></div> 245+1921	<b>54</b> +402
<b>TOTAL</b>	<b>17 sites</b>	<b>41294</b>	<b>8300 + 15547</b>	<b>7114 + 3752</b>


- CSC
- HPC2N
- DCSC-KU
- IJS

- Upgraded to dCache FHS compliant 2.2.0
  - Installed and currently benchmarking new HP DL360G7 (with P411, 1GB cache and 4 x D2600) pool nodes
  - Disks in Sun X4540 approaching three years old and we are seeing about one 1TB disk failure every two weeks between two machines

- Added two new pool nodes, also HP DL360G7, P411 and 3 x D2600 (12x2TB disks)
- Enabled possibility to move over nodes to another Internet connection if LHCOPN is unavailable
- For the Jade.hip.fi cluster
  - Upgrade of firmware on HP's 1:10Gb Ethernet modules did not fix problem that sometimes a server causes loss of Ethernet connectivity for a whole c-class enclosure (fix is to reload drivers/restart the right node).




- New dCache pool hardware
  - HP DL180G6, 12x3T SATA HDDs, 4xGigE
  - Network sizing 1 GigE per 10T effective space
    - driving requirement is being able to evacuate a pool in 24 hours
- First retirement of old dCache pool hardware :)





- Recent HP ProCurve 5400zl firmware finally supports L4 based LACP/802.1ad load balancing, and it seems to work




- No LanFree
  - dCache tape pools are migrated to not use LanFree, much less hassle and since LTO keeps missing the performance wrt the roadmap we see no reason given our limited amount of tape drives used for dCache tape pools
- Tape library
  - IBM seems to make progress on fixing LTO firmware, no spurious dumps without legit cause for quite a while now.



- In the process of refreshing infrastructure servers
  - FYI: The HP Gen8 SmartDrive tray has TWO firmware blobs on it!




- mds600
  - dCache pools
  - 3x HP MDS600 with 3x70=210 2GB SATA disks
  - 6x HP DL360 G6 with P411 controllers
  - Have been in production for over a year, however, now with dual gigabit connection






- dl180
  - dCache pools
  - 7x HP DL180 G6 with 14x 3GB SAS disks on a P212 controller
  - 2 of the disks are mounted in the back (dl180-back).
  - Have been in production for some months and lately with dual gigabit connection





- dl165
  - tier1 compute
  - 27x HP DL165 G7 with 2x AMD 6276 CPUs (32 cores), 96GB memory and 4x300GB 10k SAS disks
  - Have been in production since april



- Both big production clusters moving to SLURM this fall
  - Infosys unavailability due to torque stalls are gone after switch to SLURM even under extreme loads, like staging 2GByte/s
- Arnes on LHCONE, Pikolit will follow as soon as they have time
- Arnes cache nfs4 over infiniband
- Both have sl6 environments ready for production as soon as Atlas blesses sl6
- More sites joining Sloveninan grid, not WLCG though



# Questions?