

# NDGF Site report

News and updates



norden

NordForsk



Nordic e-Infrastructure  
Collaboration



norden


NordForsk



Nordic e-Infrastructure  
Collaboration

# Overview

- Other NeIC activities
- News in the organisation
- Central services
- Computing
- Storage
- News from our sites



## NeIC activities outside of HEP

- NeIC 2015 conference
  - e-Infrastructures for science in the Nordics
  - Espoo, Finland, 5-8th May 2015.
- Glenna cloud project
  - Resource sharing and federation of Nordic science clouds
- Nordic accounting workshop 16th-17th of October
  - Discussion on energy accounting
  - [https://wiki.neic.no/wiki/Nordic\\_accounting\\_workshop\\_2014](https://wiki.neic.no/wiki/Nordic_accounting_workshop_2014)
- Norway and Denmark plan on leaving EGI

## Organisation news

- Two new half-time sysadmins
  - Taking on tasks previously funded by EGI and national initiatives (NGI stuff), as well as real work
- Gerd re-implemented part of the TSM integration scripts into a java implementation inside dCache
  - But then had to step in as PHB March-October due to...

## Organisation news

- Me doing this instead of Annecy:
- Cured now and back in decent shape
- Felt weird not travelling for 6 months...
  - But I was kept busy with important stuff.
  - And I did have lots of time to keep up with my email, unlike now...
- Lots of “interesting” things learnt











## Central services

- Ganeti for handling redundant VMs working pretty well
  - But a few rough corners getting started
- New hardware great for dCache Chimera DB
  - Have yet to induce noticeable load from dCache, even doing high rate removals etc
  - Deleting at 450Hz, feeding enough removals into dCache was the limiting factor, not the backend
- Migrating old servers to new VMs is a slow process
  - Especially since now would be a good time to do a new base OS instead of ancient Ubuntu 10.04 and CentOS 5.

# Computing

- Doing OK
- Could be better at some sites, struggling to get appropriate share in some clusters for various reasons.
- Test deploying ARC remote cache reading
  - CE->CE direct transfer of cached file
  - As fallback or preferred for performance reasons
  - Looks promising, but with a few small issues to be fixed

Country	Site	CPUs	Load (processes: Grid+local)	Queueing
 <i>China</i>	BOINC Cluster	1176	0+0	28+0
 <i>Denmark</i>	Steno Tier 1 (DCSC/KU)	0		0+0
 <i>Germany</i>	LRZ-C2PAP	4072	0+3671	136+0
	LRZ-LMU	728	377+111	218+0
	LRZ-LMU lcg-lrz-ce0	1088	585+0	235+0
	LRZ-LMU lcg-lrz-ce3	1040	2+787	0+371
	RZG ATLAS HYDRA	167848	0+149654	8+8
	wuppertalprod	3644	185+564	213+1835
 <i>Norway</i>	Abel C1(UiO/USIT)	10872	1195+8778	195+247
	Abel C3(UiO/USIT)	10872	179+9822	58+385
 <i>Slovenia</i>	Arnes	2304	2046+0	5305+0
	SiGNET	2834	1844+363	272+898
 <i>Sweden</i>	Abisko (HPC2N)	15936	16+14984	33+0
	Alarik (SweGrid, Luna>	3776	63+3128	26+2
	Triolith - Atlas (NSC)	25472	254+21431	56+0
 <i>Switzerland</i>	ATLAS BOINC	11167	2184+1193	352+34
	Bern ce01 (UNIBE-LHEP)	1368	1+29	0+0
	Bern ce02 (UNIBE-LHEP)	776	1+0	1+0
	Bern UBELIX T3	2672	0+2300	72+214038
	Geneva (UNIGE-DPNC)	688	80+0	23+0
	LHEP HPC TEST	4288	0+2928	0+0
	Lugano PHOENIX T2	3226	56+569	120+163
	Lugano PHOENIX T2	3226	4+621	80+203
 <i>UK</i>	arc-ce01 (RAL-LCG2)	13902	2010+10990	149+0
	arc-ce02 (RAL-LCG2)	13902	1243+11728	366+0
	arc-ce03 (RAL-LCG2)	13902	1524+11499	391+0
	cetest01 (UKI-LT2-IC->	4	260+3622	98+1178
	t2arc01 UKI-SOUTHGRID>	168	789+35	1+0
<b>TOTAL</b>	<b>28 sites</b>	<b>320951</b>	<b>14818 + 258727</b>	<b>8436 + 219362</b>



# Storage

- Bit of oops when doing ALICE dark data cleanup
  - “Here's a list of all files at NDGF.”
    - “These are the files we have no record of at NDGF.”
  - “So we can delete all these?”
    - “Yup, go ahead and delete them!”
  - “Are you sure?”
    - “Yes, we're sure.”
  - “No undeletes available, really sure?”
    - “Yes, yes, go ahead!”
  - “Ok then!”

# Storage

- Bit of oops when doing ALICE dark data cleanup
  - Turns out that when we say “NDGF”, we mean all namespace entries in the srm.ndgf.org dCache instance under /alice/
  - ALICE mean the ALICE::NDGF catalog
  - Which is distinct from ALICE::NDGF\_tape
  - So, we, uhm, deleted all ALICE tape data...
  - Noticed when a tape library admin asked “huh, didn't we have more tape used in T1ALICET storage class than 1 half-full one?” a week or two later
  - Currently resubscribing 200+TB tape data at trickle pace.

## News from sites



# Something is burnt in the state of Denmark



## NSC, Linköping

- New central routers bought. Two Juniper QFX5100, replacing two Catalyst 4506E. Planned production later this autumn.
  - Planning to run BGP, for peering between LHCOPN and Internet (but probably not until early next year).
- New cluster for weather forecasts and climate simulation ordered. 641 nodes, 10256 cores (Haswell). 800 Tbyte Lustre storage. Installation in November, production in spring.
- New storage system for academic use: 3 IBM GSS26; 2.8 Pbyte user-accessible (4 Pbyte raw disk). In production in November.
- Expanded Lustre storage for climate data with another 1.7 Pbyte.

## HPC2N, Umeå

- New Lustre filesystem for academic HPC use
  - DDN storage
  - Infiniband networking
  - 1 PB, >20GB/s from 300 clients doing mixed read+write
- New tape frontend dCache pools
- Found an “interesting”  $O(n^2)$  performance bug in TSM's dsmd retrieve (from archived files with >1 link)
  - Noticable at 1k files, hurts at 10k, 100k impossible
  - Ongoing case with IBM progressing during the meeting

# Questions?

