

NDGF

The Distributed Tier1 Site



norden

NordForsk



Nordic e-Infrastructure  
Collaboration

# Overview

- Reason
- Design
- Reality
- Experience

## Reason

- No Nordic country is big enough to host a Tier-1
- All the Nordic countries together are big enough to warrant a Tier-1
- 6% of ATLAS and 9% of ALICE Tier-1 resources
- Currently distributed among seven HPC sites from Copenhagen in the south to Umeå in the north
  - Plus Slovenian Tier-2 disk installed technically as NDGF T1 disk

# Design

- Computing on ARC
  - It's a grid, so just run jobs wherever
    - Mostly on shared HPC resources
  - Form the Tier-1 and Tier-2 sites in accounting and BDII
- Storage in distributed dCache
  - Took a bit of development effort to make it appear as one endpoint
  - Central headnodes, pools out at sites
- Ops/Devel staff mix
  - Fix software problems in the software projects instead of creating procedures to handle them in operations

# Reality - ARC

2013-10-27 CET 19:22:31

● Tier-1 resources

● Associated Tier-2 resources

ARC as seen from ATLAS



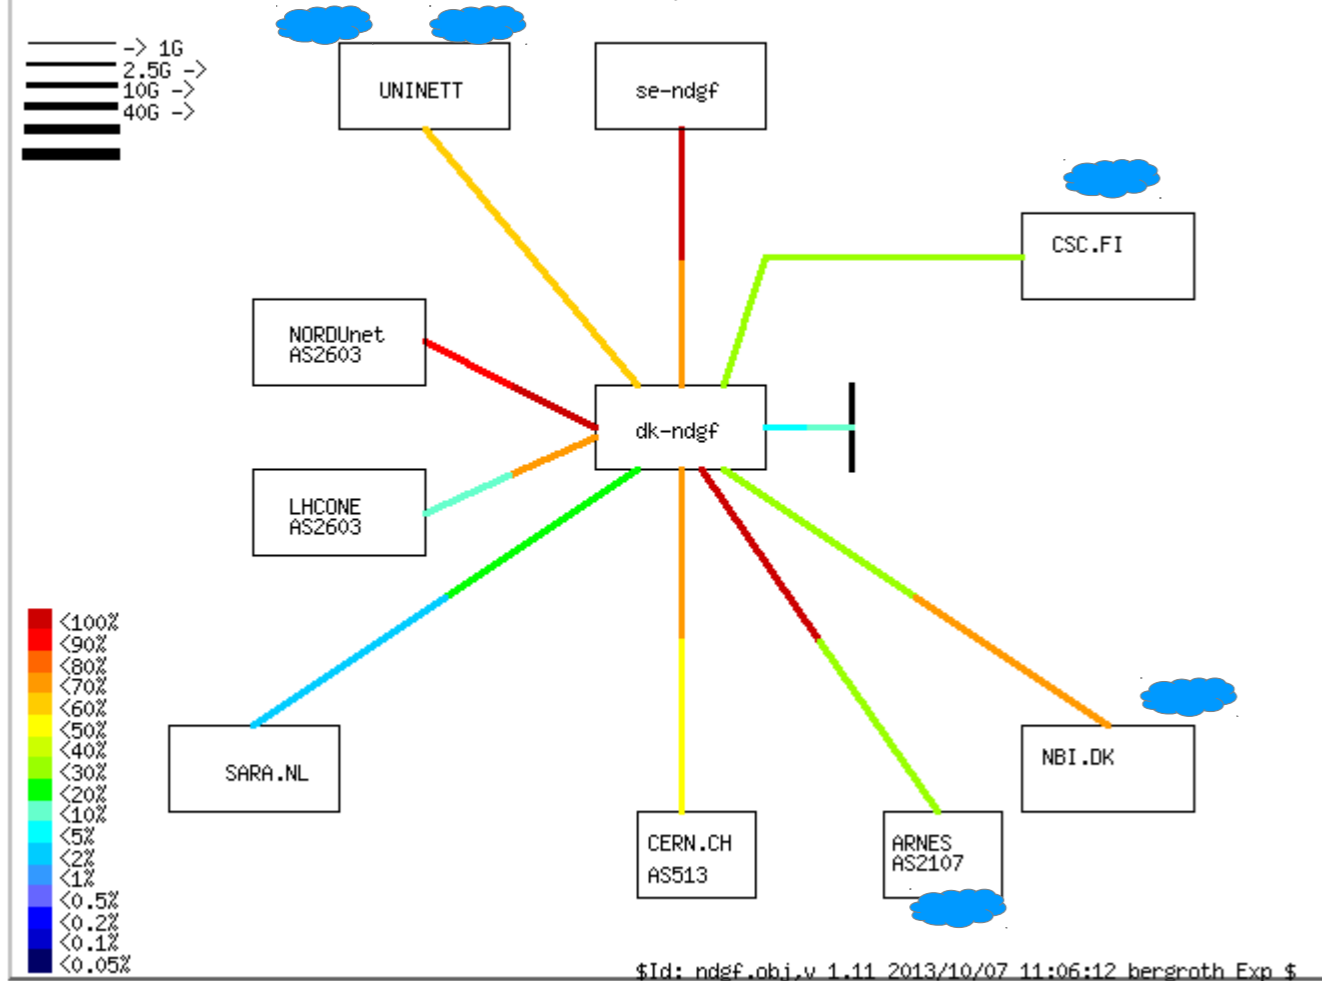
Processes: ■ Grid ■ Local

Country	Site	CPUs	Load (processes: Grid+local)	Queueing
Denmark	● Steno Tier 1 (DCSC/KU)	5888	<span style="color: green;">■</span> 832+3502 <span style="color: gray;">■</span>	465+0
	LRZ-LMU	2312	<span style="color: green;">■</span> 1128+769 <span style="color: gray;">■</span>	245+101
Germany	LRZ-LMU	2312	<span style="color: green;">■</span> 774+1124 <span style="color: gray;">■</span>	87+210
	RZG ATLAS HYDRA	157296	<span style="color: gray;">■</span> 0+148344	10+0
	wuppertalprod	3268	<span style="color: green;">■</span> 2777+953 <span style="color: gray;">■</span>	338+11962
Norway	● Abel C1(UiO/USIT)	10864	<span style="color: green;">■</span> 875+7392 <span style="color: gray;">■</span>	90+0
	● Abel C3(UiO/USIT)	10864	<span style="color: green;">■</span> 902+7380 <span style="color: gray;">■</span>	4+0
Slovenia	● Arnes	1924	<span style="color: green;">■</span> 1236+0 <span style="color: gray;">■</span>	243+0
	● SiGNET	2194	<span style="color: green;">■</span> 1668+10 <span style="color: gray;">■</span>	389+44
Sweden	● Abisko (HPC2N)	15456	<span style="color: green;">■</span> 442+14142 <span style="color: gray;">■</span>	165+0
	● Alarik (SweGrid, Luna>	3776	<span style="color: green;">■</span> 336+1813 <span style="color: gray;">■</span>	81+0
	● Ritsem (SweGrid, HPC2>	544	<span style="color: green;">■</span> 456+1 <span style="color: gray;">■</span>	105+0
	● Siri (SweGrid, Lunarc)	448	<span style="color: green;">■</span> 312+73 <span style="color: gray;">■</span>	131+50
	● Smokerings (NSC)	504	<span style="color: green;">■</span> 302+202 <span style="color: gray;">■</span>	68+0
	● Tintin (SweGrid, Uppm>	2624	<span style="color: green;">■</span> 11+2292 <span style="color: gray;">■</span>	0+20
Switzerland	● Bern ATLAS WLCG T2	1352	<span style="color: green;">■</span> 1135+0 <span style="color: gray;">■</span>	195+0
	● Bern UBELIX T3	2272	<span style="color: green;">■</span> 600+977 <span style="color: gray;">■</span>	135+1104
	● Geneva ATLAS T3	402	<span style="color: green;">■</span> 243+61 <span style="color: gray;">■</span>	34+0
	● Manno PHOENIX T2	2490	<span style="color: green;">■</span> 742+1238 <span style="color: gray;">■</span>	337+246
	● Manno PHOENIX T2	192	<span style="color: green;">■</span> 122+20 <span style="color: gray;">■</span>	416+384
UK	arc-ce01 (RAL-LCG2)	5426	<span style="color: green;">■</span> 1261+3675 <span style="color: gray;">■</span>	465+7
	arc-ce02 (RAL-LCG2)	5423	<span style="color: green;">■</span> 1201+3732 <span style="color: gray;">■</span>	414+1
	arc-ce03 (RAL-LCG2)	5425	<span style="color: green;">■</span> 1301+3633 <span style="color: gray;">■</span>	423+1
	cetest01 (UKI-LT2-IC->	4	<span style="color: green;">■</span> 193+2656 <span style="color: gray;">■</span>	184+1075
<b>TOTAL</b>	<b>24 sites</b>	<b>243260</b>	<b>18849 + 203989</b>	<b>5024 + 15205</b>

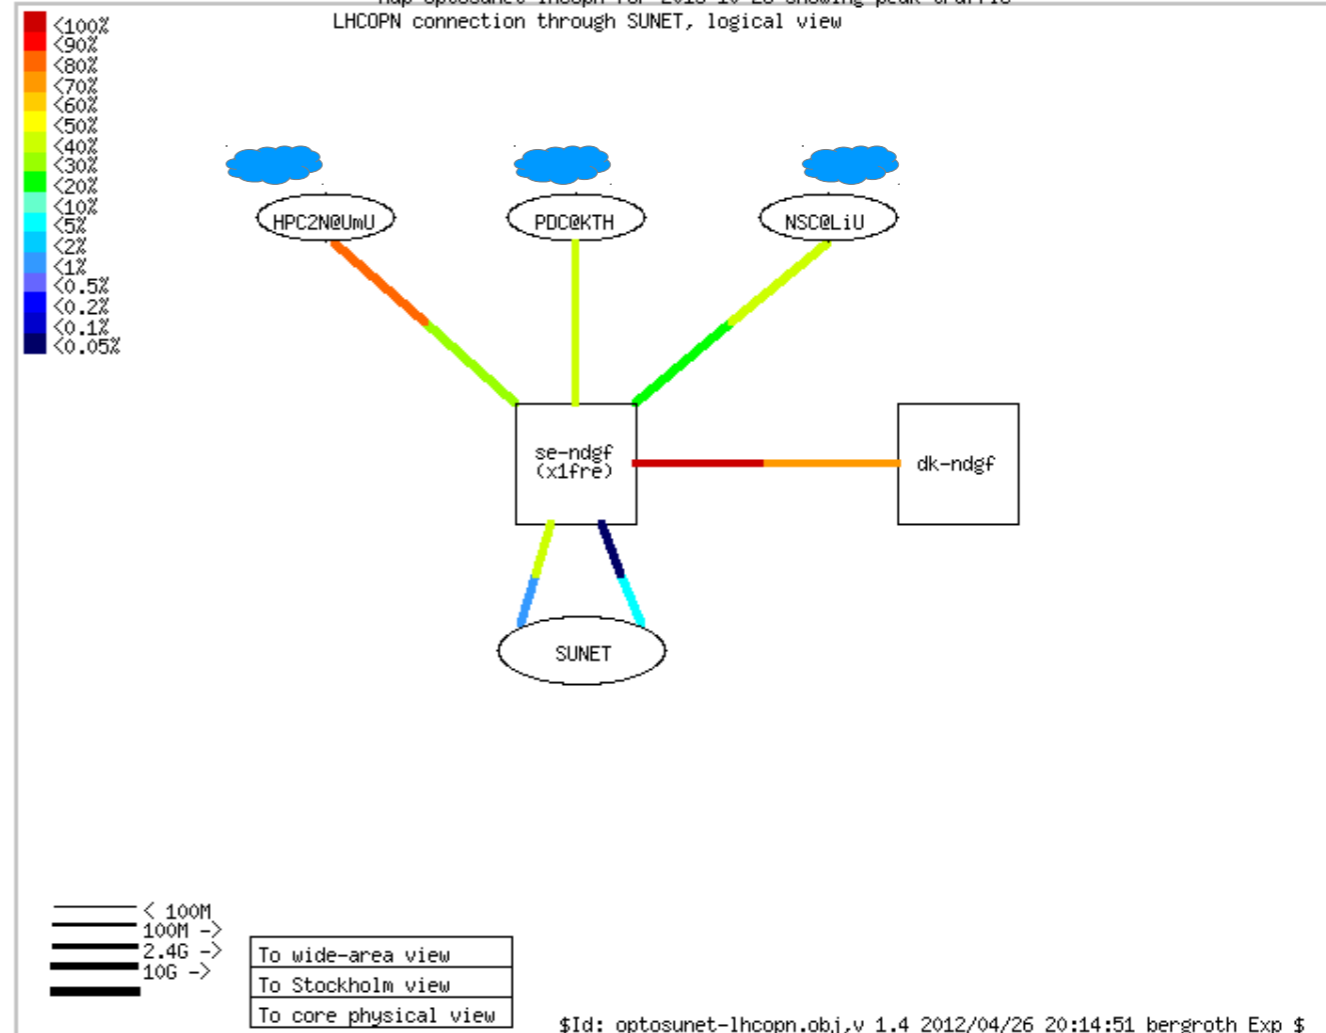
# Reality - Networking

- The clouds point to where there are dCache pools

Map ndgf for 2013-10-26 showing peak traffic  
NDGF Tier-1 private network







Map optosunet-lhcopn for 2013-10-26 showing peak traffic  
LHCOPN connection through SUNET, logical view





# Reality - Storage

- Pools at 7 different sites
- Namespace, admin, and doors centrally
- Graphs show a busy Alice day and a calm Atlas day



## Experience

- Distributed storage great for data taking
  - Somewhat less so for data availability
- Computing on ARC with data staging makes short interventions on storage transparent
- Distributed funding at least as challenging as distributed storage and computing
- Professional HPC sites are good for running clusters and storage, professional developers are good at fixing software issues and making users and ops lives happy





norden

NordForsk

neic

Nordic e-Infrastructure  
Collaboration

# Questions?