



A Distributed Tier-1 for WLCG


*Michael Grønager, PhD
Technical Coordinator, NDGF
CHEP 2007*

Victoria, September the 3rd, 2007

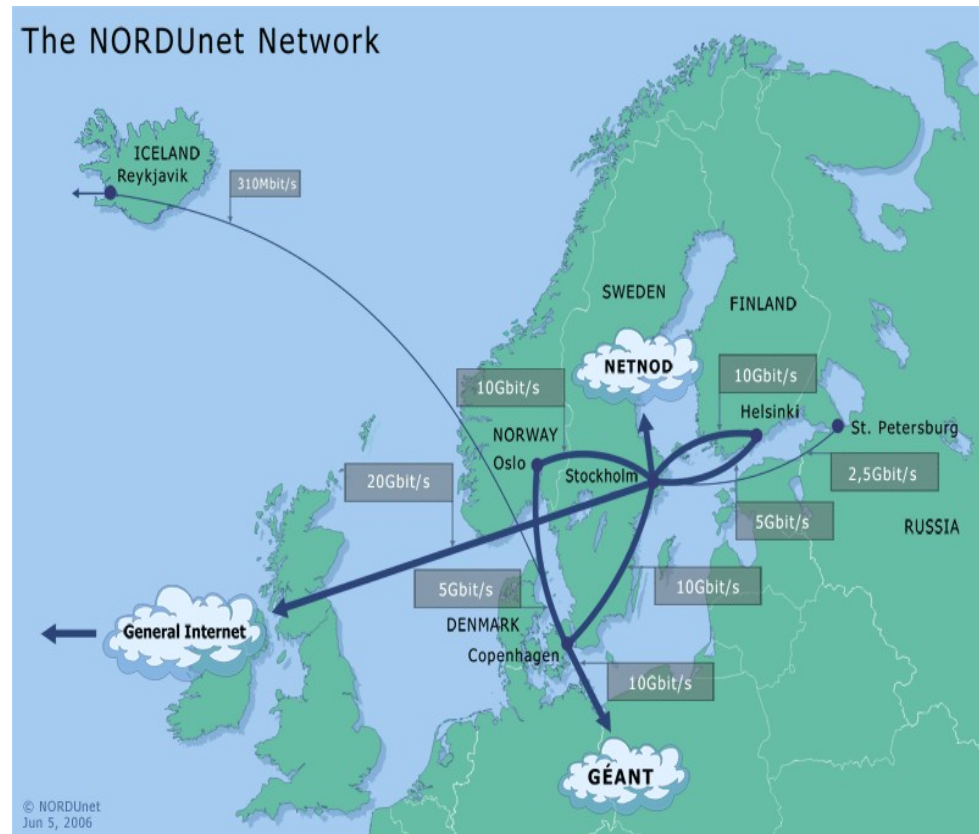
- Background
- Organization / Governance
- Tier-1 Services:
 - Computing
 - Storage
 - ATLAS
 - ALICE
 - Accounting
 - Monitoring
- Operation

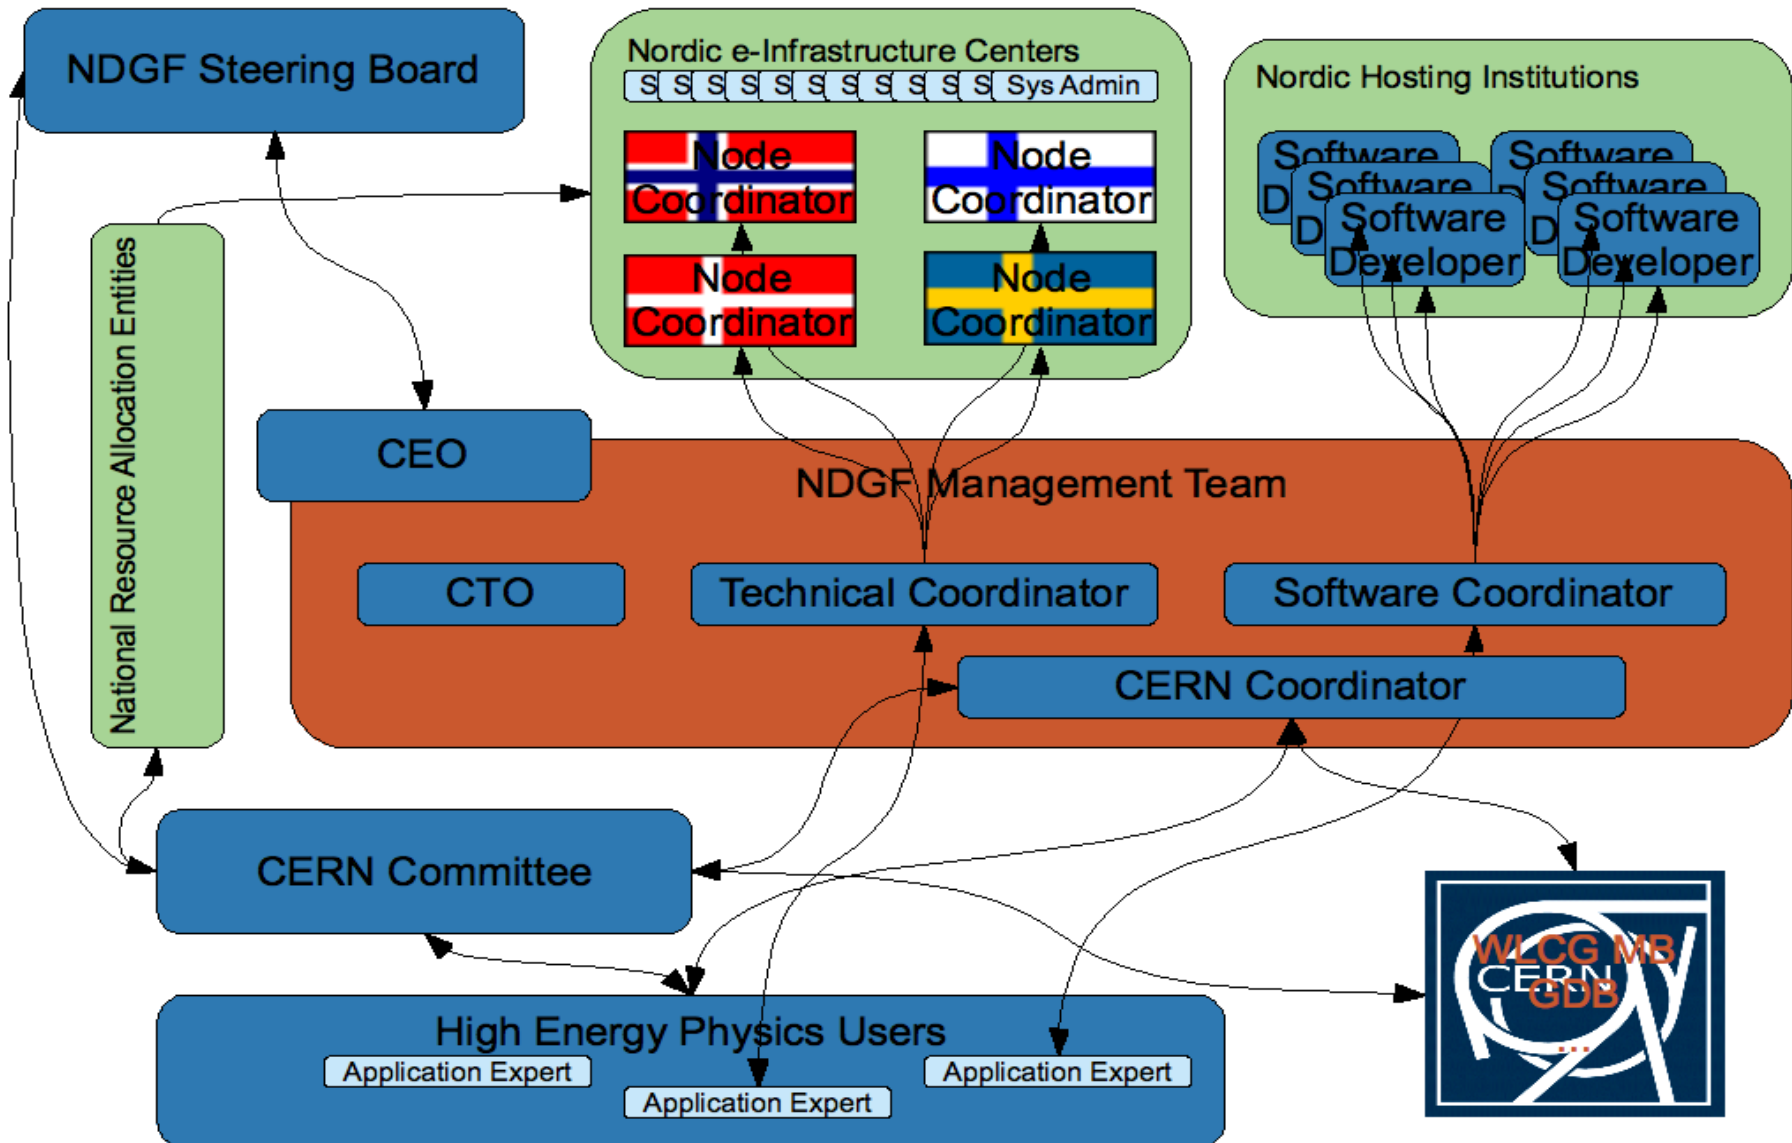
- Nordic Countries constitute together 25Mio People
- No country is bigger than 10Mio People
- Nordic Tier-1 needs to utilize hardware at bigger Nordic compute sites
- Strong Nordic grid tradition: NorduGrid / ARC
 - Deployed at all Nordic compute centers
 - Used heavily also by non-HEP users (>75%)
- Need for a pan-Nordic organization for Tier-1 and possibly other huge inter/Nordic e-Science projects

- A Co-operative Nordic Data and Computing Grid facility
 - Nordic production grid, leveraging national grid resources
 - Common policy framework for Nordic production grid
 - Joint Nordic planning and coordination
 - Operate Nordic storage facility for major projects
 - Co-ordinate & host major e-Science projects (i.e., Nordic WLCG Tier-1)
 - Develop grid middleware and services
- NDGF 2006-2010
 - Funded (2 M.EUR/year) by National Research Councils of the Nordic countries
 - Builds on a history of Nordic grid collaboration




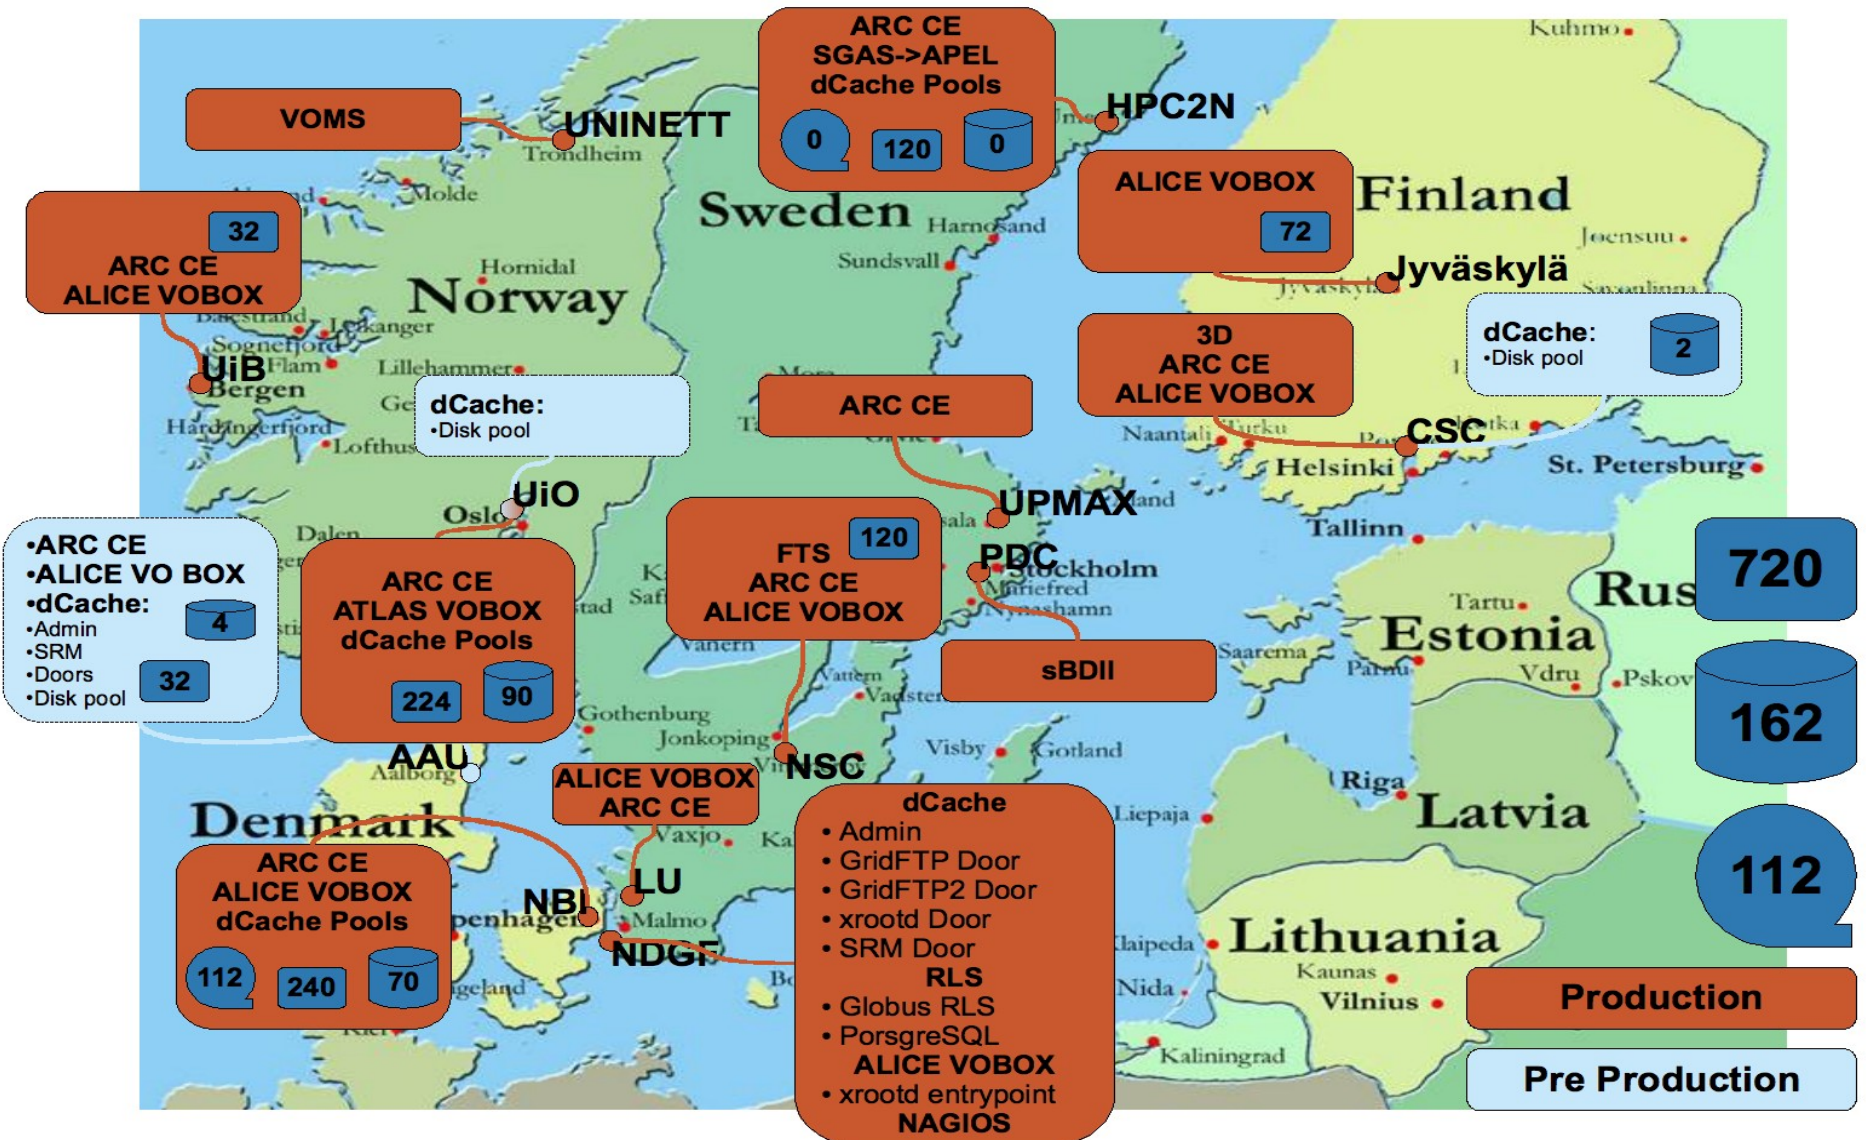
- The Nordic Regional Research and Educational Network (RREN)
- Owned by the 5 Nordic National RENs
- 25 Years of Nordic network collaboration
- Leverage National Initiatives
- Participates in major international efforts
- Represents Nordic NRENS internationally, gateway to the Nordic area



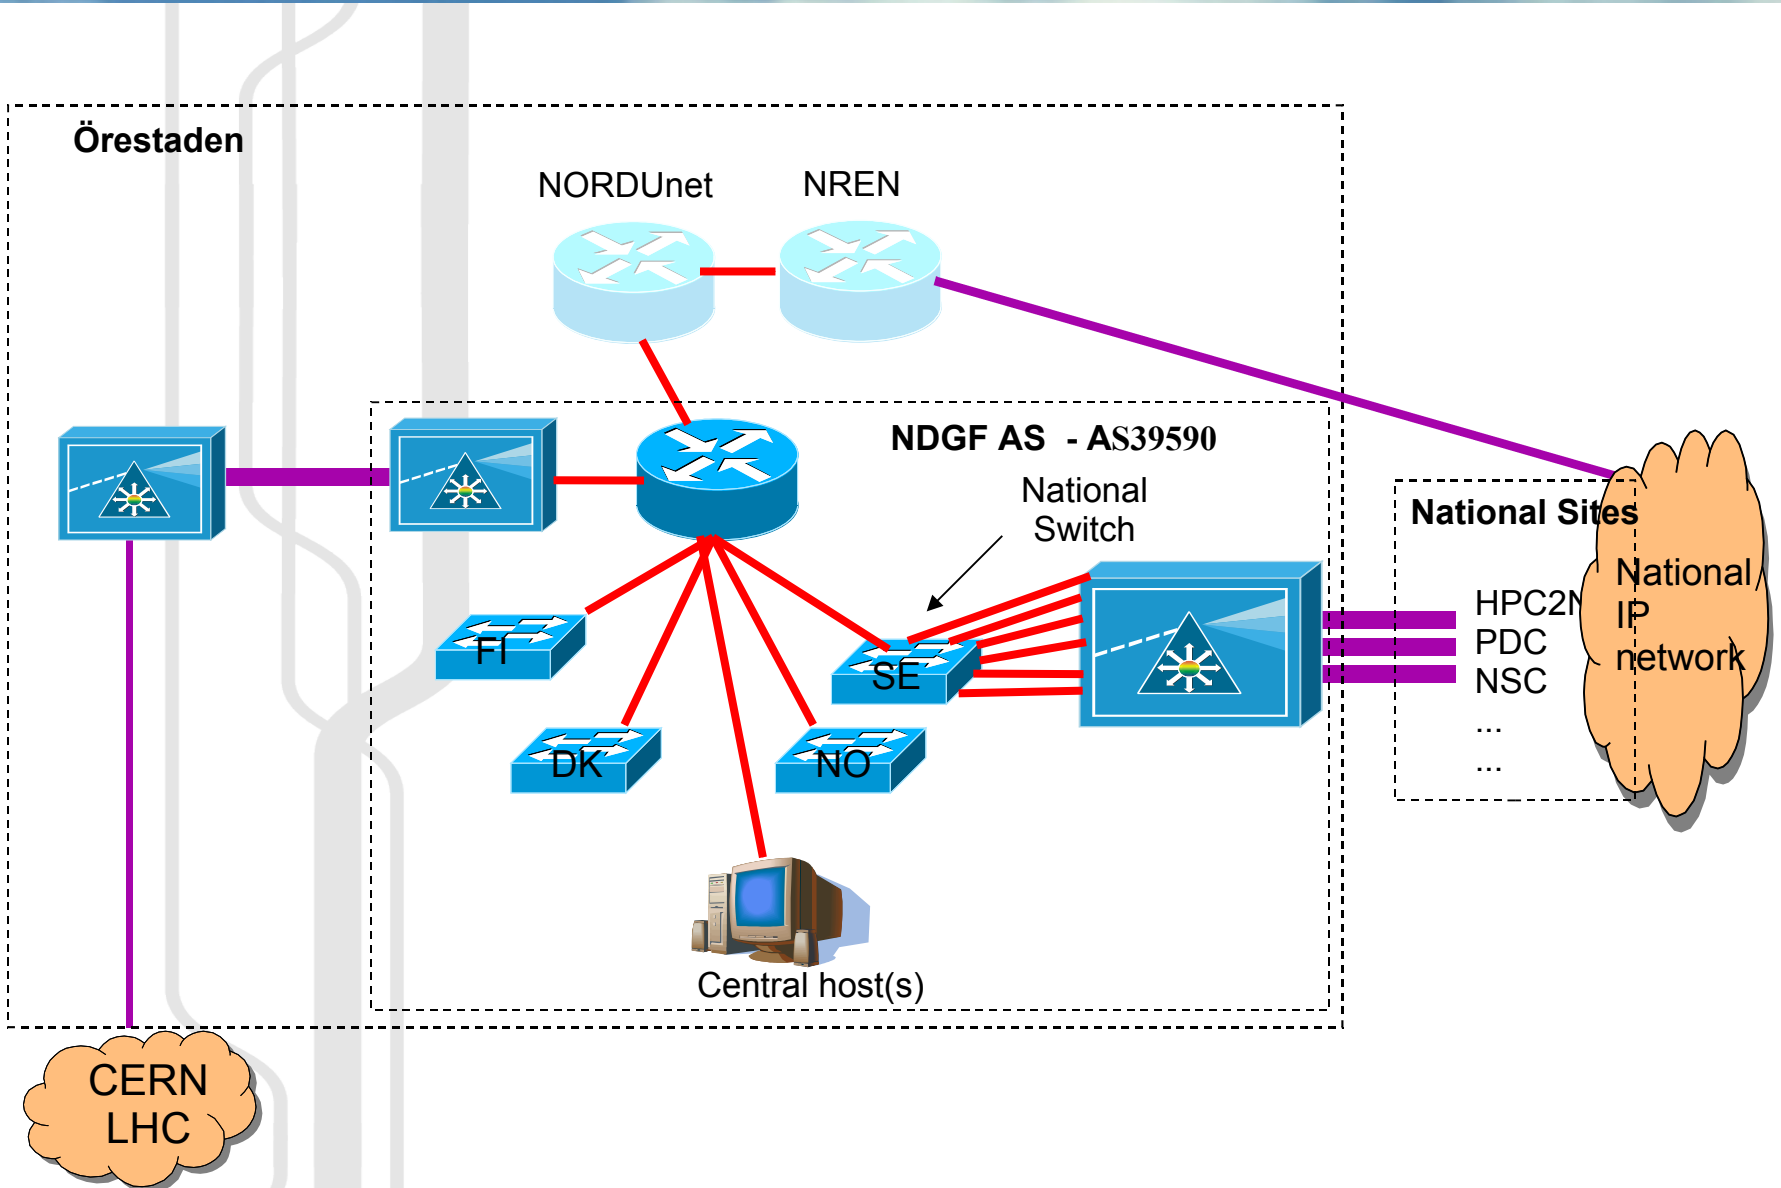


- The 7 biggest Nordic compute centers, dTier-1s, form the NDGF Tier-1
- Resources (Storage and Computing) are scattered
- Services can be centralized
- Advantages in redundancy
- Especially for 24x7 data taking





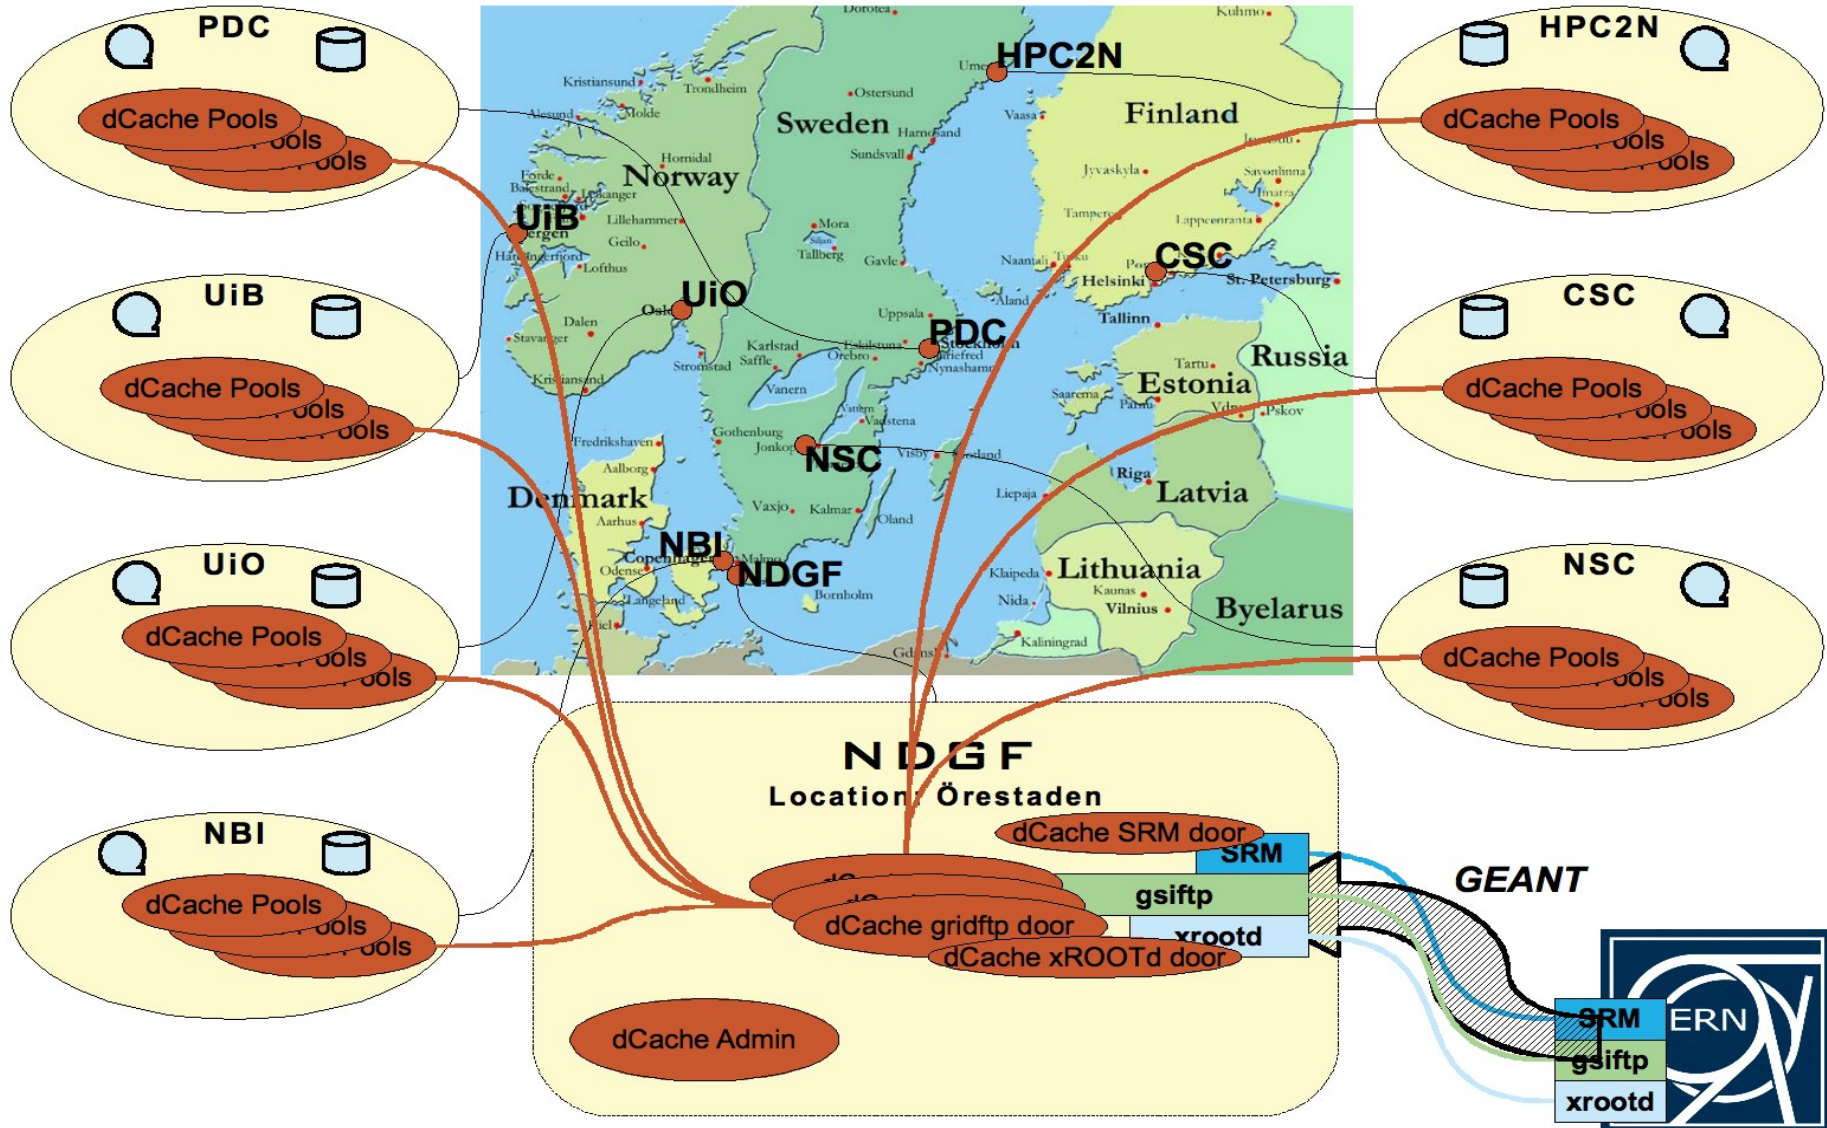
- Today NDGF is connected directly with GEANT 10Gbit fiber to CERN
- Inter-Nordic shared 10Gbit network from NORDUnet
- A Dedicated 10Gbit LAN covering all dTier-1 centers next year




- NorduGrid / ARC middleware for Computing
- Used routinely since 2002 for e.g. ATLAS data challenges
- Deployed at all the dTier-1 sites

Country	Site	CPUs	Load (processes: Grid+local)	Queuing
Australia	Atlas (UniMelb)	26	0+2	0+0
	Charm (UniMelb)	36	0+0 (queue down)	0+0
	Alfred (UniMelb)	20	0+6	2+1
	Resusc (UNU)	10	0+0	0+0
Denmark	Aalborg Grid Gateway	46	38+0	0+0
	Niflheim (DCSC/DTU)	902	0+898	0+17
	Horseshoe (DCSC/SDU)	1192	0+873	0+3
	HEPAX1	1	0+0	0+0
	Morpheus	18	15+0	23+0
Estonia	Theory (DCSC/KU)	112	0+42	0+1
	UT-MCB (Krakow/ka clu)	15	0+0	0+0
	UT-CS Antarctica Clus>	20	0+0	0+0
	CMS on CERN Linux	1	0+0	0+0
	CMS Production server	5	0+0	0+0
	UT DOUG Cluster	2	0+0	0+0
	CMS test cluster	1	0+0	0+0
	EENet cluster	6	0+0	0+0
	UT Physics Cluster	3	3+0	0+0
	CSC Kirppu	1	1+0	6+0
Finland	Mill (Physics)	60	0+15	0+0
	Alpha (HIP)	1	0+0	0+0
	Testbed0 (HIP)	1	0+0	4+1
Germany	FZK cluster	996	83+349	0+0
	LRZ cluster	234	0+230	0+243
Norway	Oslo Temp Cluster	11	0+0	25+0
	Parallab IBM Cluster	58	0+57	0+75
	Bergen Grid Cluster	2	2+0	7+0
	Oslo Grid Cluster	41	9+15	51+0
UiO Grid	100	0+98	0+1	
Slovenia	SIGNET	40	6+31	6+0
Sweden	Bluesmoke (SweGrid,NS>	99	95+0	187+0
	Kosufy farm	60	36+0	0+0
	ISV	4	4+0	14+0
	Hagrid (SweGrid, Uppm>	100	50+0	68+0
	Ingrid (SweGrid,HPC2N)	101	69+0	124+0
	Monolith (NSC)	398	0+342	0+121
	Quark Cluster	7	0+0	0+0
	Beppe (SweGrid PDC KT>	96	92+0	49+0
Switzerland	Sigrid (SweGrid, Luna>	99	49+50	19+25
	Toto7/Whenim64 (Lunar>	192	0+161	0+11
	Bern ATLAS Cluster	8	8+0	12+0
TOTAL	42 sites	5196	570 + 3169	597 + 499

- dCache Installation
- Admin and Door nodes at GEANT endpoint
- Pools at sites
- Very close collaboration with DESY to ensure dCache is suited also for distributed use



- Central Installation:
 - 7 Dell 1950 2xDual Core 2GHz Xeon, 4GB RAM, 2 x 73GB 15k SAS disks (mirrored) (one for spare)
 - 2 x Dell PowerVault MD-1000 direct attached storage enclosures with 7 x 143GB 15k SAS RAID-10 each
- Running:
 - 2 Postgress for PNFS running in HA mode (master-slave) DB on MD-1000
 - 1 PNFS Manager and Pool Manager
 - 1 SRM, location manager, statistics, billing, etc.
 - 1 GridFTP and xrootd door on one machine
 - 1 Monitoring and intrusion detection on one machine



- Central Installation:
 - 7 Dell 1950 2xDual Core 2GHz Xeon, 4GB RAM, 2 x 73GB 15k SAS disks (mirrored) (one for spare)
 - 2 x Dell PowerVault MD-1000 direct attached storage enclosures with 7 x 143GB 15k SAS RAID-10 each
- Running:
 - 2 Postgress for PNFS running in HA mode (master-slave) DB on MD-1000
 - 1 PNFS Manager and Pool Manager
 - 1 SRM, location manager, statistics, billing, etc.
 - 1 GridFTP and xrootd door on one machine
 - 1 Monitoring and intrusion detection on one machine



- See talk on dCache and gridFTP2
- *A Distributed Storage System with dCache*
- *Carson Hall C at 16.50*

- Running FTS2.0
- Patched version of Globus supporting GridFTP2
- Located in Linkjeping:
 - 1 Server for FTS
 - 1 Server for Oracle database
- Channels:
 - STAR-NDGF
 - others...

- Minimal setup located in Helsinki:
 - One dual core dual Xeon box with 4GB of memory
 - no RAC, just one server
 - High availability SAN storage
 - a bit more than one TB of space allocated for data
 - upgrade to 3-5 node RAC in 2008

- ATLAS VOBox (ARC flavor) services fully implemented

- ARC uses Globus RLS

- US-ATLAS-LRC view on the mysql

- Enables outside ATLAS subscription to data stored on old Ses

- and internal through RLS

- See poster #74 Wednesday

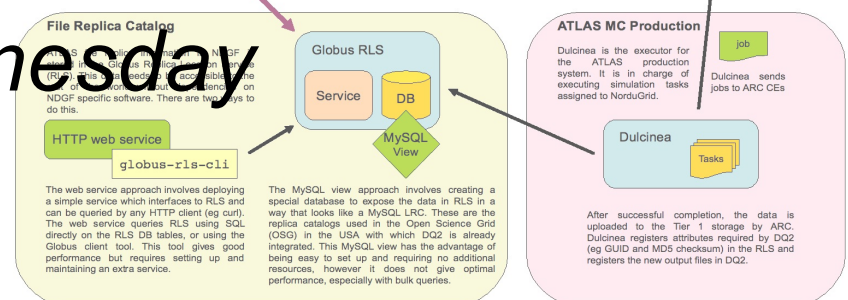
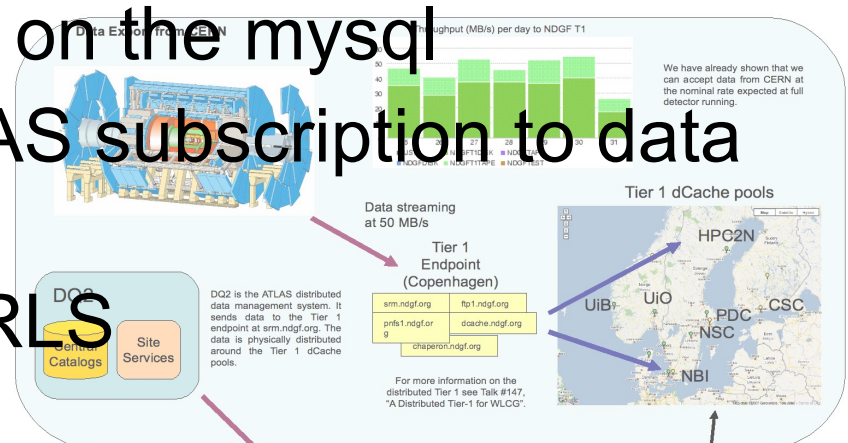


Josva Kleist¹, David Cameron¹, Adrian Taga², Gerd Behrmann¹, Mattias Eilert¹

¹Nordic Data Grid Facility
²University of Oslo, 0316 Oslo, Norway



The Nordic Data Grid Facility (NDGF) consists of Grid resources running ARC middleware in Scandinavia and other countries. These resources serve many virtual organisations and contribute a large fraction of total worldwide resources for the ATLAS experiment, whose data is distributed and managed by the DQ2 software. Managing ATLAS data within NDGF and data distribution between NDGF and other Grids used by ATLAS (the LHC Computing Grid and the Open Science Grid) presents a unique challenge for several reasons. Firstly, the entry point for data, the Tier 1 centre, is physically distributed among heterogeneous resources in several countries and yet must present a single access point for all data stored within the centre. The middleware framework used in NDGF differs significantly from other Grids, specifically in the way that all data movement and registration is performed by services outside the worker node environment. Also, the service used for cataloging the location of data files is different from other Grids but must still be useable by DQ2 and ATLAS users to locate data within NDGF. This poster presents in detail how we solve these issues to allow seamless access worldwide to data within NDGF.



- ALICE VOBox boxes:

- Jyväskylä

- CSC

- NSC

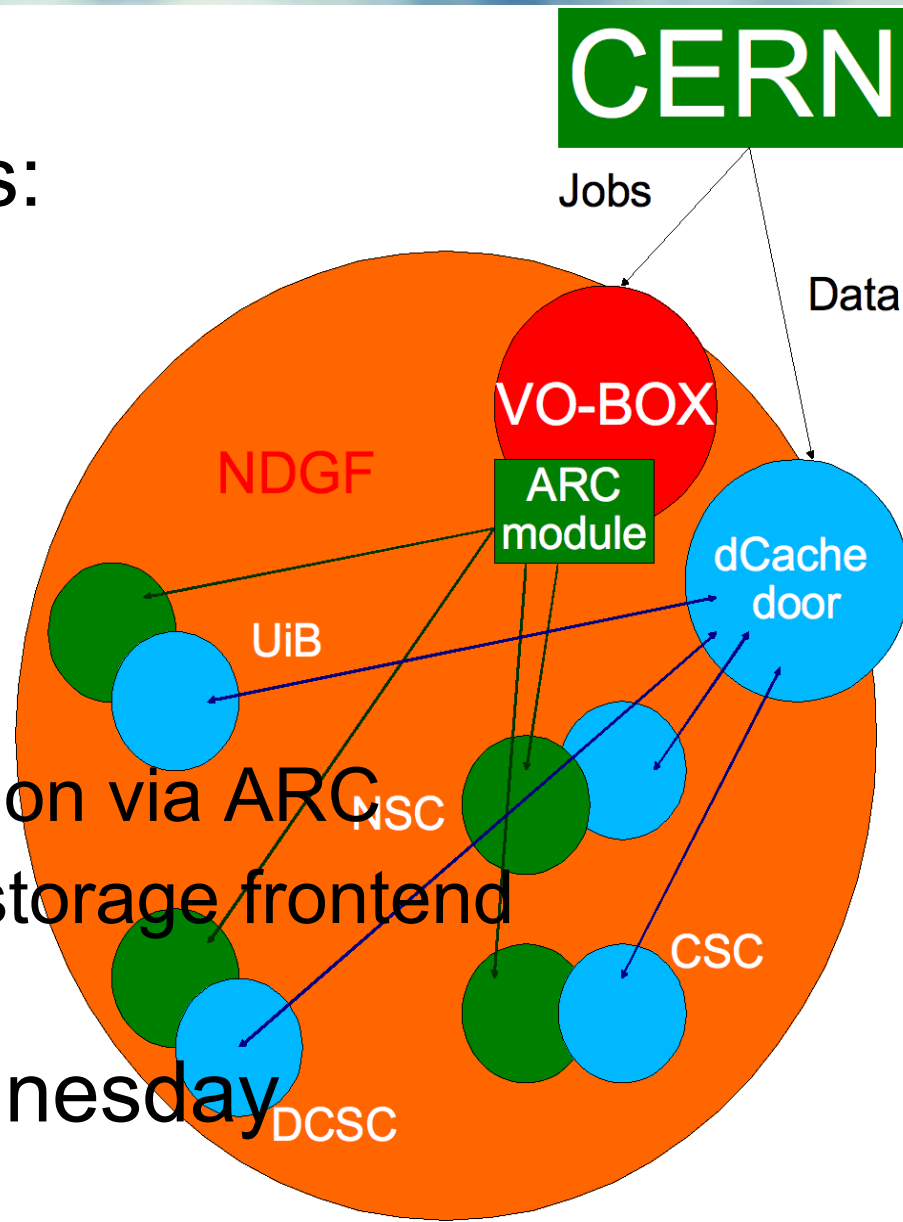
- LUNARC

- DCSC/KU

- UiB – using submission via ARC

- Örestaden – xrootd storage frontend

- See poster #75 Wednesday




- SAM sensors:
 - BDII
 - SE
 - SRM
 - FTS


- ARC-CE

- This is the only different sensor as compared to other sites


Overall Service Availability for site NDGF-T1 : Daily Report



Individual Service Availability for site NDGF-T1 : Daily Report

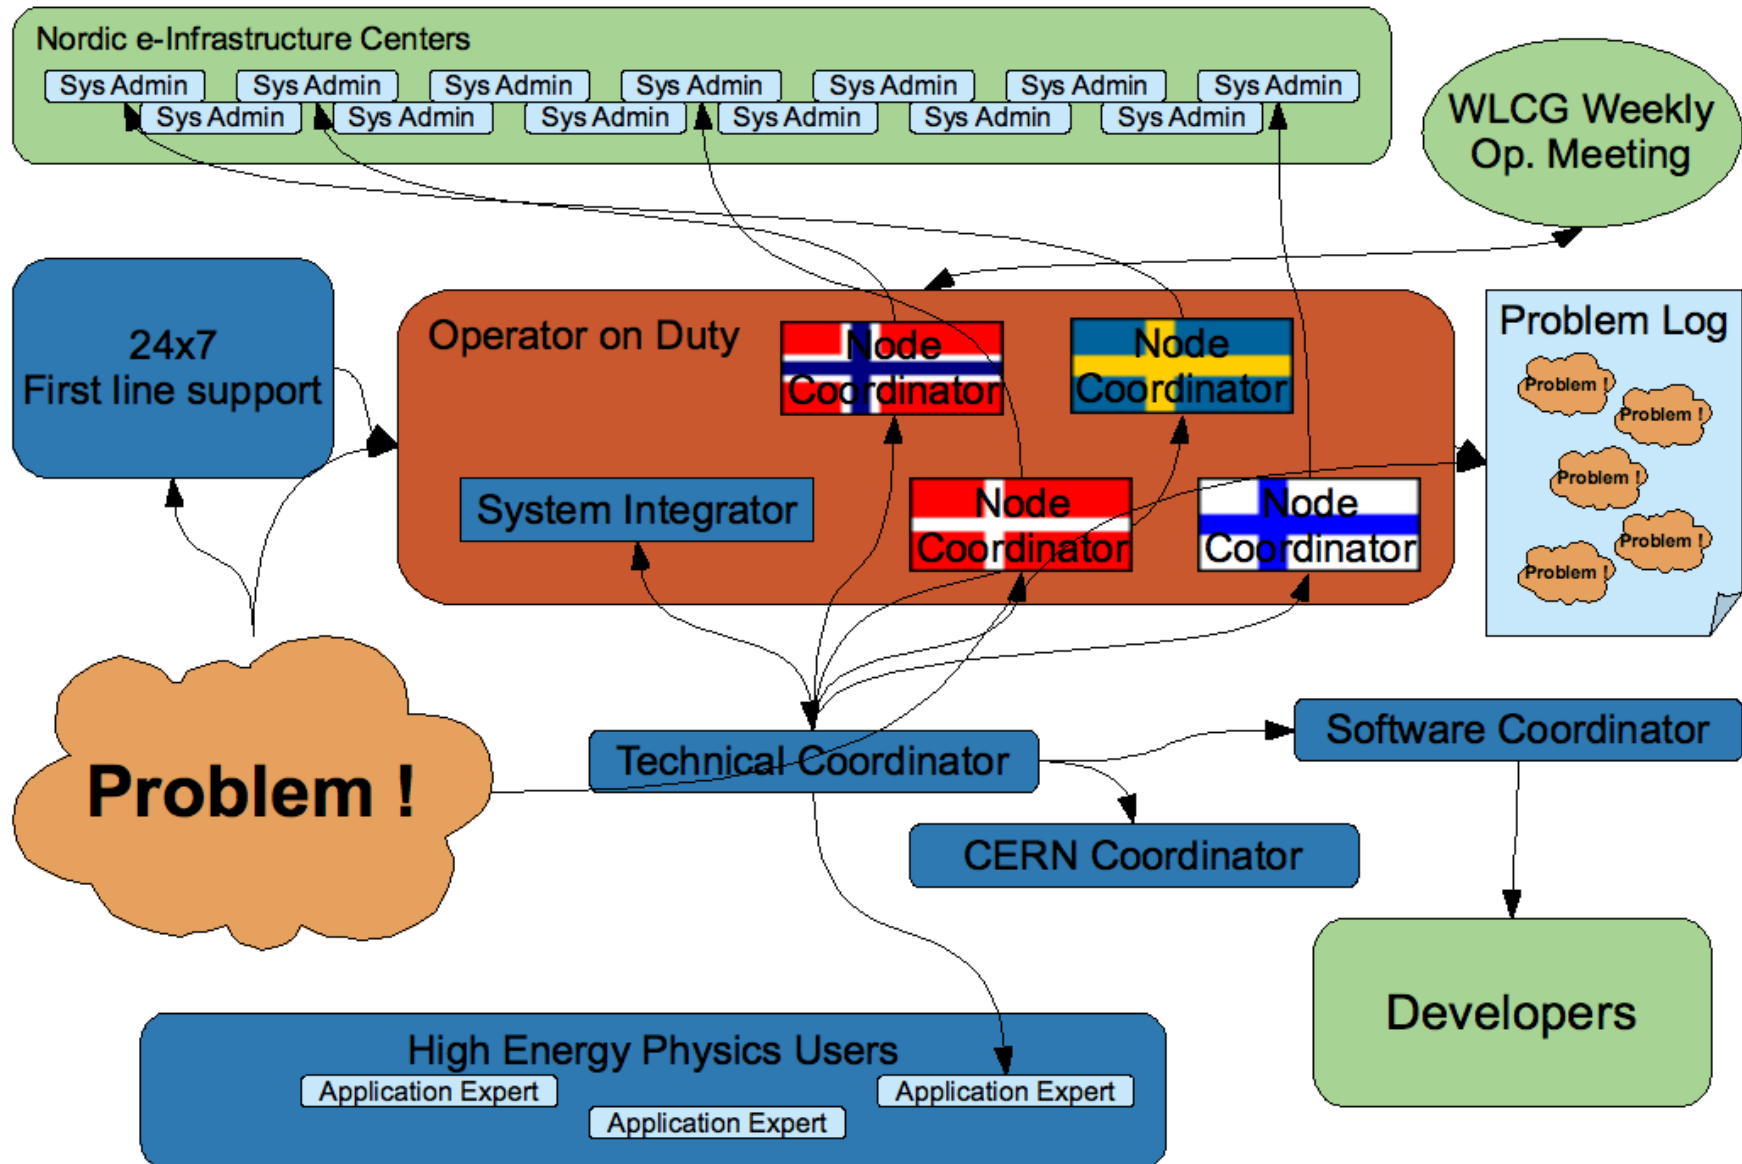


Service Instance Availability for site NDGF-T1 : Daily Report

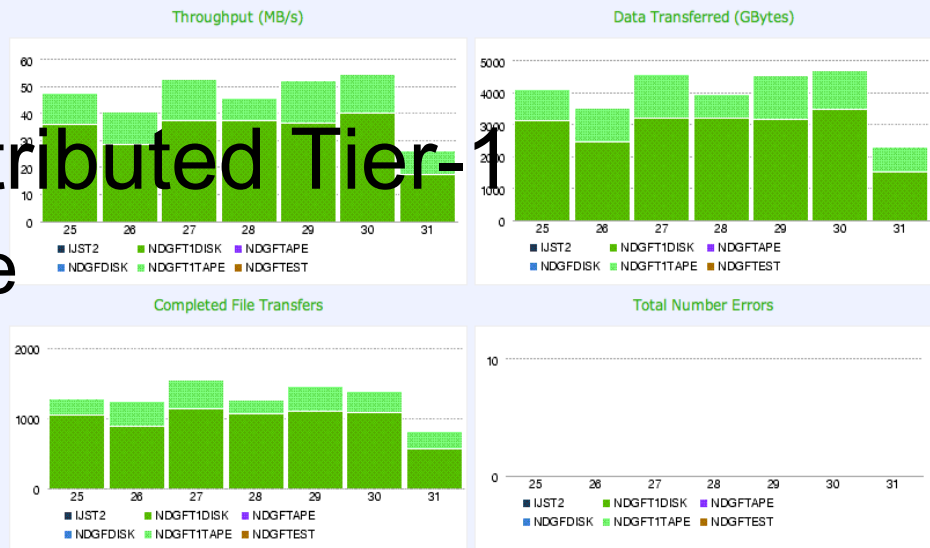


- Sites report using SGAS
 - (SweGrid Accounting System)
- SGAS report translated to APEL
- Injected into the APEL DB


- Functional from September 07
 - some sites already accounted



- We have build a distributed Tier-1
 - dCache – for storage
 - ARC for computing
- Interoperabel with:
 - ALICE
 - ATLAS
 - ARC monitoring and accounting
 - LCG monitoring and accounting
- It works !



Overall Service Availability for site NDGF-T1 : Daily Report



Individual Service Availability for site NDGF-T1 : Daily Report

