




NDGF – A Distributed Computing Centre

Mattias Wadenstein
System Integrator, NDGF
Data Retention Workshop at DESY
Hamburg, 2009-01-26

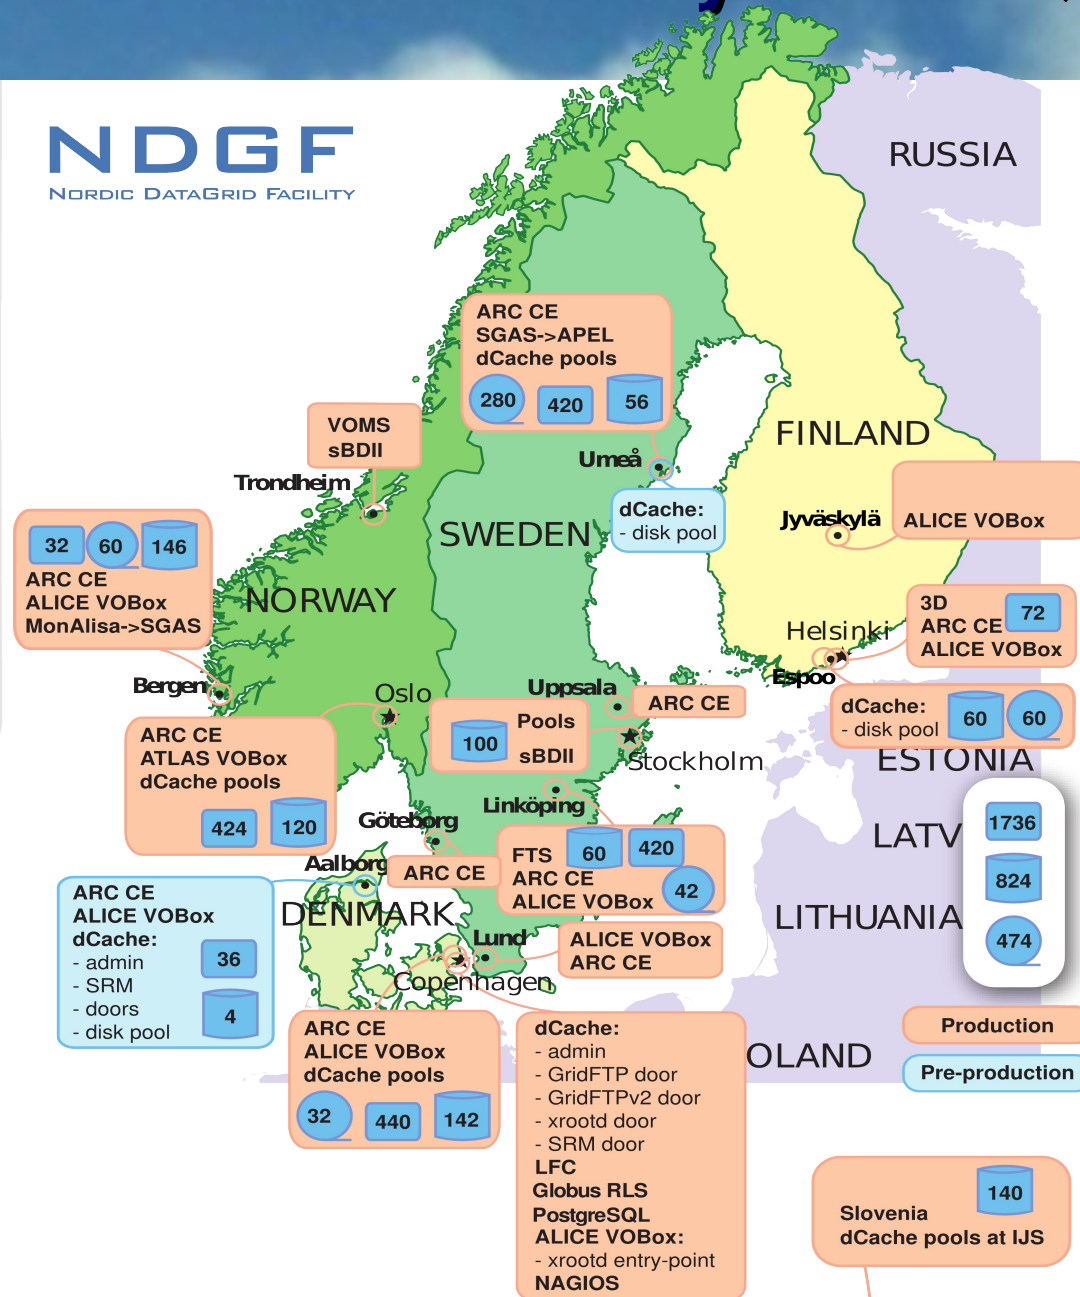
- The NDGF Project
- Status of the infrastructure
 - Resources
 - Services
 - Operation
- The LHC Tier1
- Example site: HPC2N

- 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 eScience projects (i.e., Nordic WLGC Tier-1)
 - Develop grid middleware and services
- NDGF 2006-2010
 - Funded (2 M€/year) by National Research Councils of the Nordic Countries

- NDGF participates in the discussion of future grid infrastructures
 - Participates in the Nordic Task Force who looks to the long term perspective for grid in the nordic countries
 - Elected observers in EGI-PB (European Grid Initiative Policy Board)

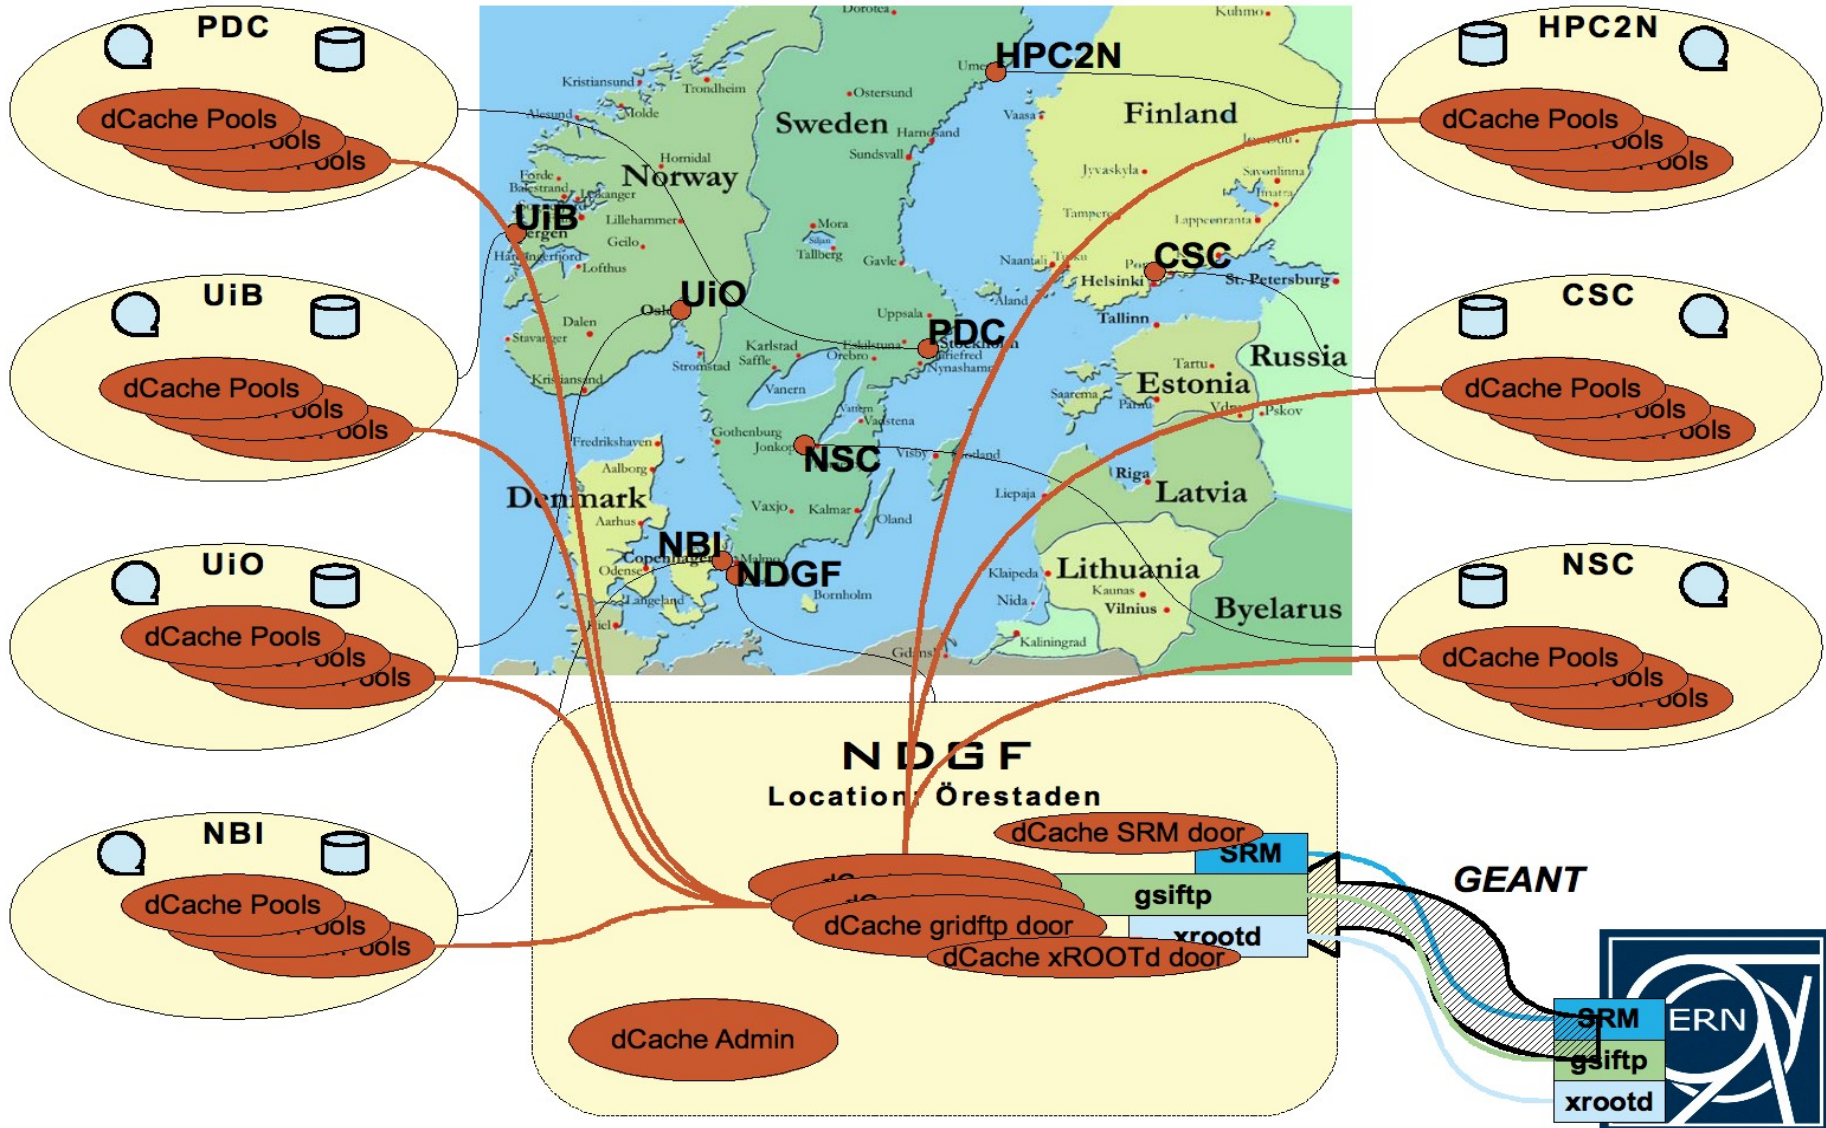


- One of several e-science projects
- Currently by far the biggest at NDGF
- Supports ATLAS and ALICE
 - NDGF also provides technical assistance to the Finnish CMS tier2
- Requires long-term commitment
 - Data available at least until 10 years after LHC shutdown
 - USB 6.0 stick mailed to anyone interested? :)




- 10 Production sites
- 2-4 MSI2k CPU equivalents
- 800 TB of Disk storage
- 500 TB of Tape storage
- Computing running Ubuntu 6.06/8.04, RHEL/CentOS 3,4,5, fedora, gentoo...
- Storage all the above plus Solaris and AIX

- 1st line support – (ready for operation)
 - NORDUnet NOC – 24x7
- 2nd line support – (in operation)
 - Operator on Duty – 8x365
 - EGEE NE-ROC
- 3rd line support – (in operation)
 - NDGF Operation Staff
 - Sys Admins at sites
- Shared tickets with NUNOC



- Central Installation:
 - 7 Dell 1950 2xDual Core 2GHz Xeon, 4GB RAM, 2 x 73GB 15k SAS disks (mirrored) (one spare node)
 - 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
 - 2 LFC servers in failover configuration
 - 1 Monitoring and intrusion detection on one machine



- Central Installation:
 - 7 Dell 1950 2x Dual Core 2GHz Xeon, 4GB RAM, 2 x 73GB 15k SAS disks (mirrored) (one spare node)
 - 2 x Dell PowerVault MD-1000 direct attached storage enclosures with 7 x 145GB 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
 - 2 LFC servers in failover configuration
 - 1 Monitoring and intrusion detection on one machine

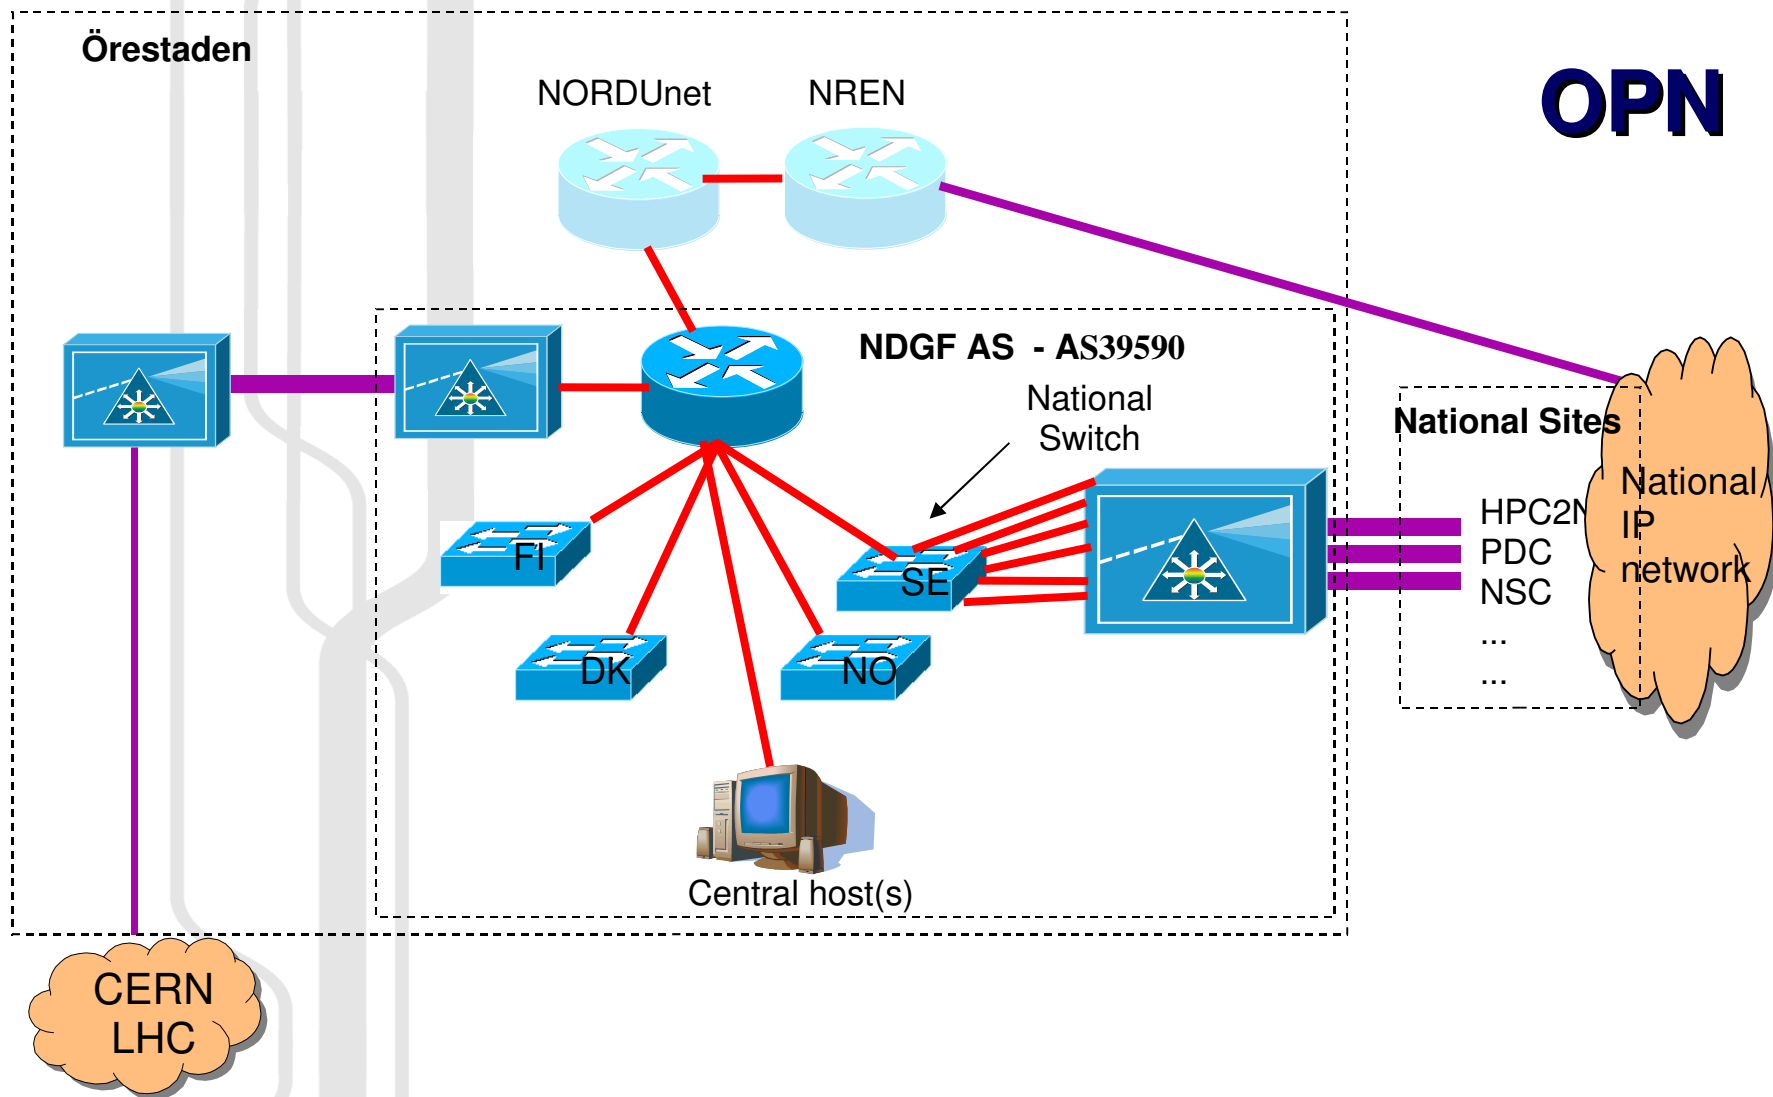


- NorduGrid / ARC middleware for Computing
- Used routinely since 2002 for e.g. ATLAS data challenges
- Deployed at all the dTier-1 sites

Grid Monitor - Microsoft Internet Explorer

Processes: ■ Grid ■ Local

Country	Site	CPUs	Load (processes: Grid+local)	Queueing
Australia	Atlas (UniMelb)	26	0+2	0+0
	Charm (UniMelb)	36	0+0 (queue down)	0+0
	Alfred (UniMelb)	90	0+6	2+1
Denmark	DistLab (DIKU)	10	0+0	0+0
	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
	Theory (DCSC/IKU)	112	0+42	0+1
	VCR (VideoRecorder)	1	1+0 (queue down)	0+0
Estonia	UT IMCB Anakonda clus>	15	3+0	0+0
	UT CS Antarctica Clus>	20	6+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
Finland	CSC Kirppu	1	1+0	6+0
	Mill (Physicum)	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
	Sigrid (SweGrid, Luna>	99	49+50	19+25
Toto7/Whenim64 (Lunar>	192	0+161	0+11	
Switzerland	Bern ATLAS Cluster	8	8+0	12+0
TOTAL	42 sites	5196	570 + 3169	597 + 499



- HPC2N in Umeå, Sweden
 - provides computational resources for swedish academic researches
 - National allocation
 - Mostly MPI jobs
- Also participates in various grid projects
 - Swegrid, NDGF, EGEE, GIRD, ETC


- compute resource
 - 5376-core IBM blade system with infiniband for general usage
 - 432-core supermicro twin nodes dedicated for grid
 - 380-core opteron & myrinet cluster



- 100 TB GPFS
- 1 TB AFS
 - Software, home, etc
- IBM 3584
 - 1000 tape slots
 - LTO4
- LHC tier1 disk
 - 56TB running
 - 200 TB ordered



- New machine room
 - strict hot/cold aisle
 - 28 racks
 - 25 kW front to back aircooling per rack
 - total ~400kW
 - no redundant power
 - Only for computing, storage&servers in old machine room
 - no raised floor
 - 3.5m high



- Per rack:
- 1 x 32A 400V 3-phase
- Or 3 x 32A 230V 1-phase
- Separate cabling ladders for:
- HPC interconnect (infiniband, myrinet, etc)
- Ethernet
- Power





Questions?