



NDGF Site Report

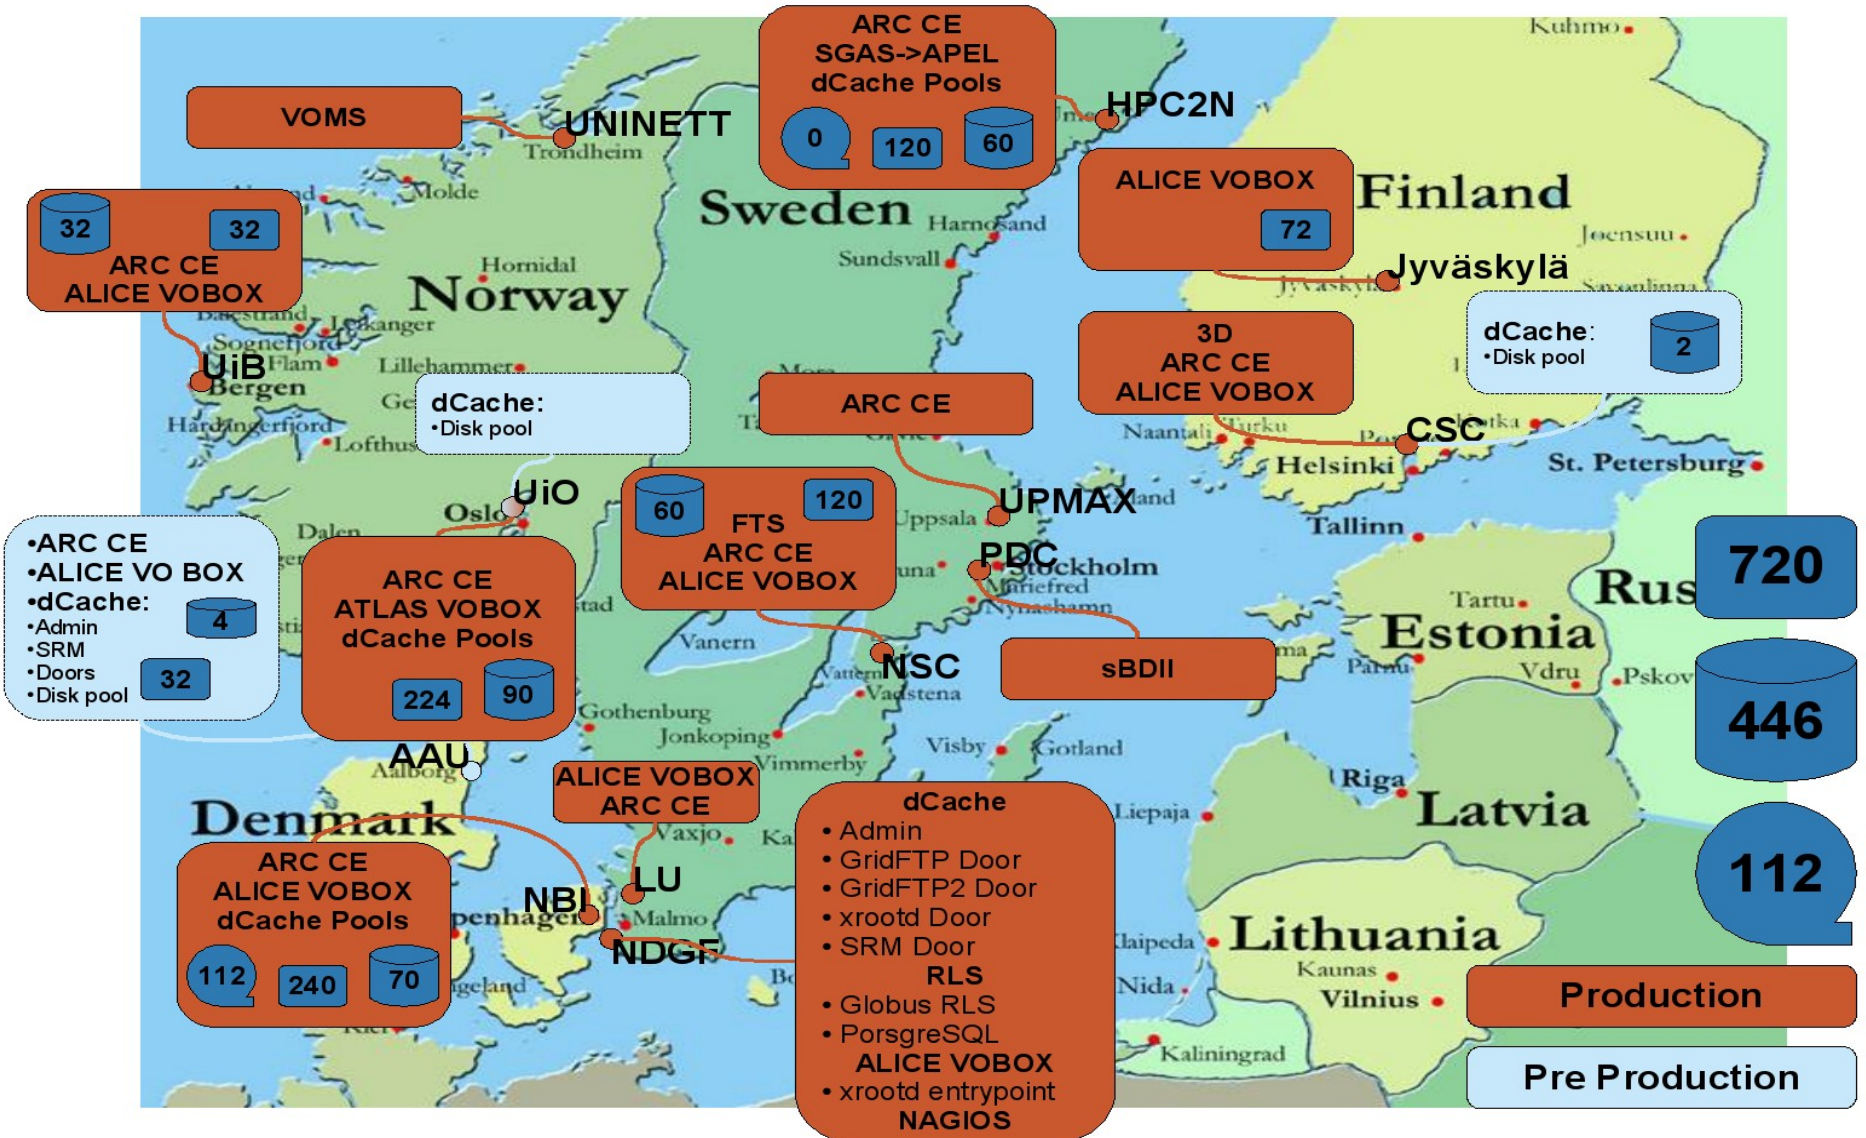
Mattias Wadenstein
System Integrator, NDGF
HEPiX fall 2007 at GSC
St. Louis, 2007-11-06

- Looking back
- Status of the infrastructure
 - People
 - Resources
 - Services
 - Operation
- e-Science Projects
- Looking ahead
- Example of one site's hardware (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 WLGCC Tier-1)
 - Develop grid middleware and services
- NDGF 2006-2010
 - Funded (2 M€/year) by National Research Councils of the Nordic Countries


- One year ago
 - Management group just hired
 - Initial contact to important partners
 - No operation staff
 - No development staff
- Major challenges
 - Hire people
 - Acquire resources
 - Build a WLCG Tier-1
 - Do it distributed
 - On heterogeneous resources
 - Startup new e-Science projects

- The WLCG Tier-1 Challenge
 - No Nordic contact point towards CERN
 - No coordination of resources
 - No coordination of efforts
- The Nordic Tier-1 was behind with everything
- A lot of skepticism from CERN, EGEE, Nordic community
- Also a need for starting up new projects
- And to be a european player




- Graduate head count ramp up over the last year
 - All positions full since 2007Q2
- Operation Team (Technical Coord.)
 - 1 System integrator
 - 4 Node Coordinators
- Development Team (Middleware Coord.)
 - 4 Infrastructure developers
 - 2 Project developers
 - 1 New Project developer
- CERN Coordinator

- 8 Production sites
- 720 kSI2k CPU equivalents
- 304 TB of Disk storage
- 112 TB of Tape storage
- Running Ubuntu 6.06, RHEL/CentOS 3,4,5, fedora core 1...
- 3 Pre Production sites
- 64 kSI2k CPU equivalents
- 4 TB of Disk storage

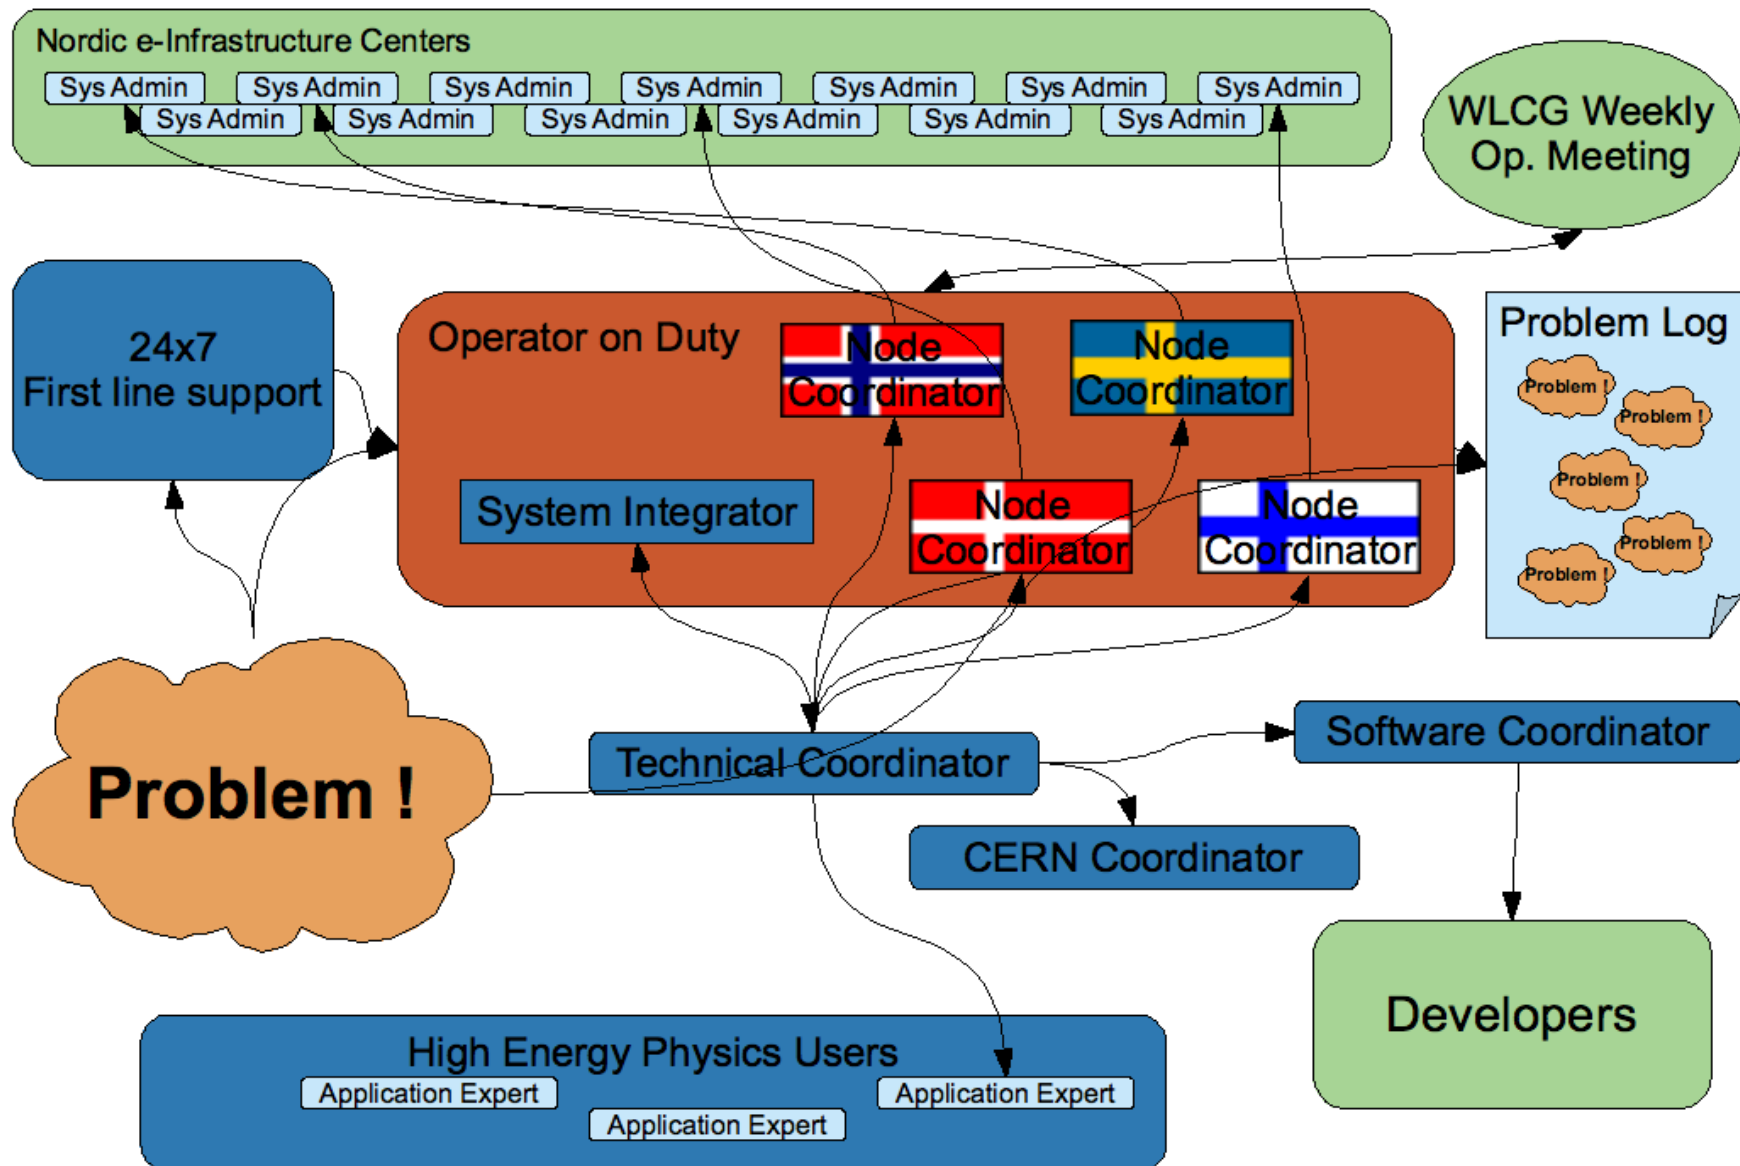


- Central Installation:
 - 7 Dell 1950 2xDual Core 2GHz Xeon, 4GB RAM, 2 x 73GB 15k SAS disks (mirrored) (one forspare)
 - 2 x Dell PowerVault MD-1000 direct attached storage enclosures with 7 x 143GB 15k SAS RAID-10 each
- Running:
 - 2 Postgress for PNFS running in HA mode (master-slave) DB on MD-1000
 - 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 forspare)
 - 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





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

- CERN
 - Build up of Tier-1: Completed
 - Running the Tier-1: Ongoing
 - Integrating Tier-2s: Ongoing
- BIO (A Nordic Biogrid)
 - Under development
- CO2 (CO2 sequestration)
 - Under development

The Infrastructure: Computing


- 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 (DCSCIKU)	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

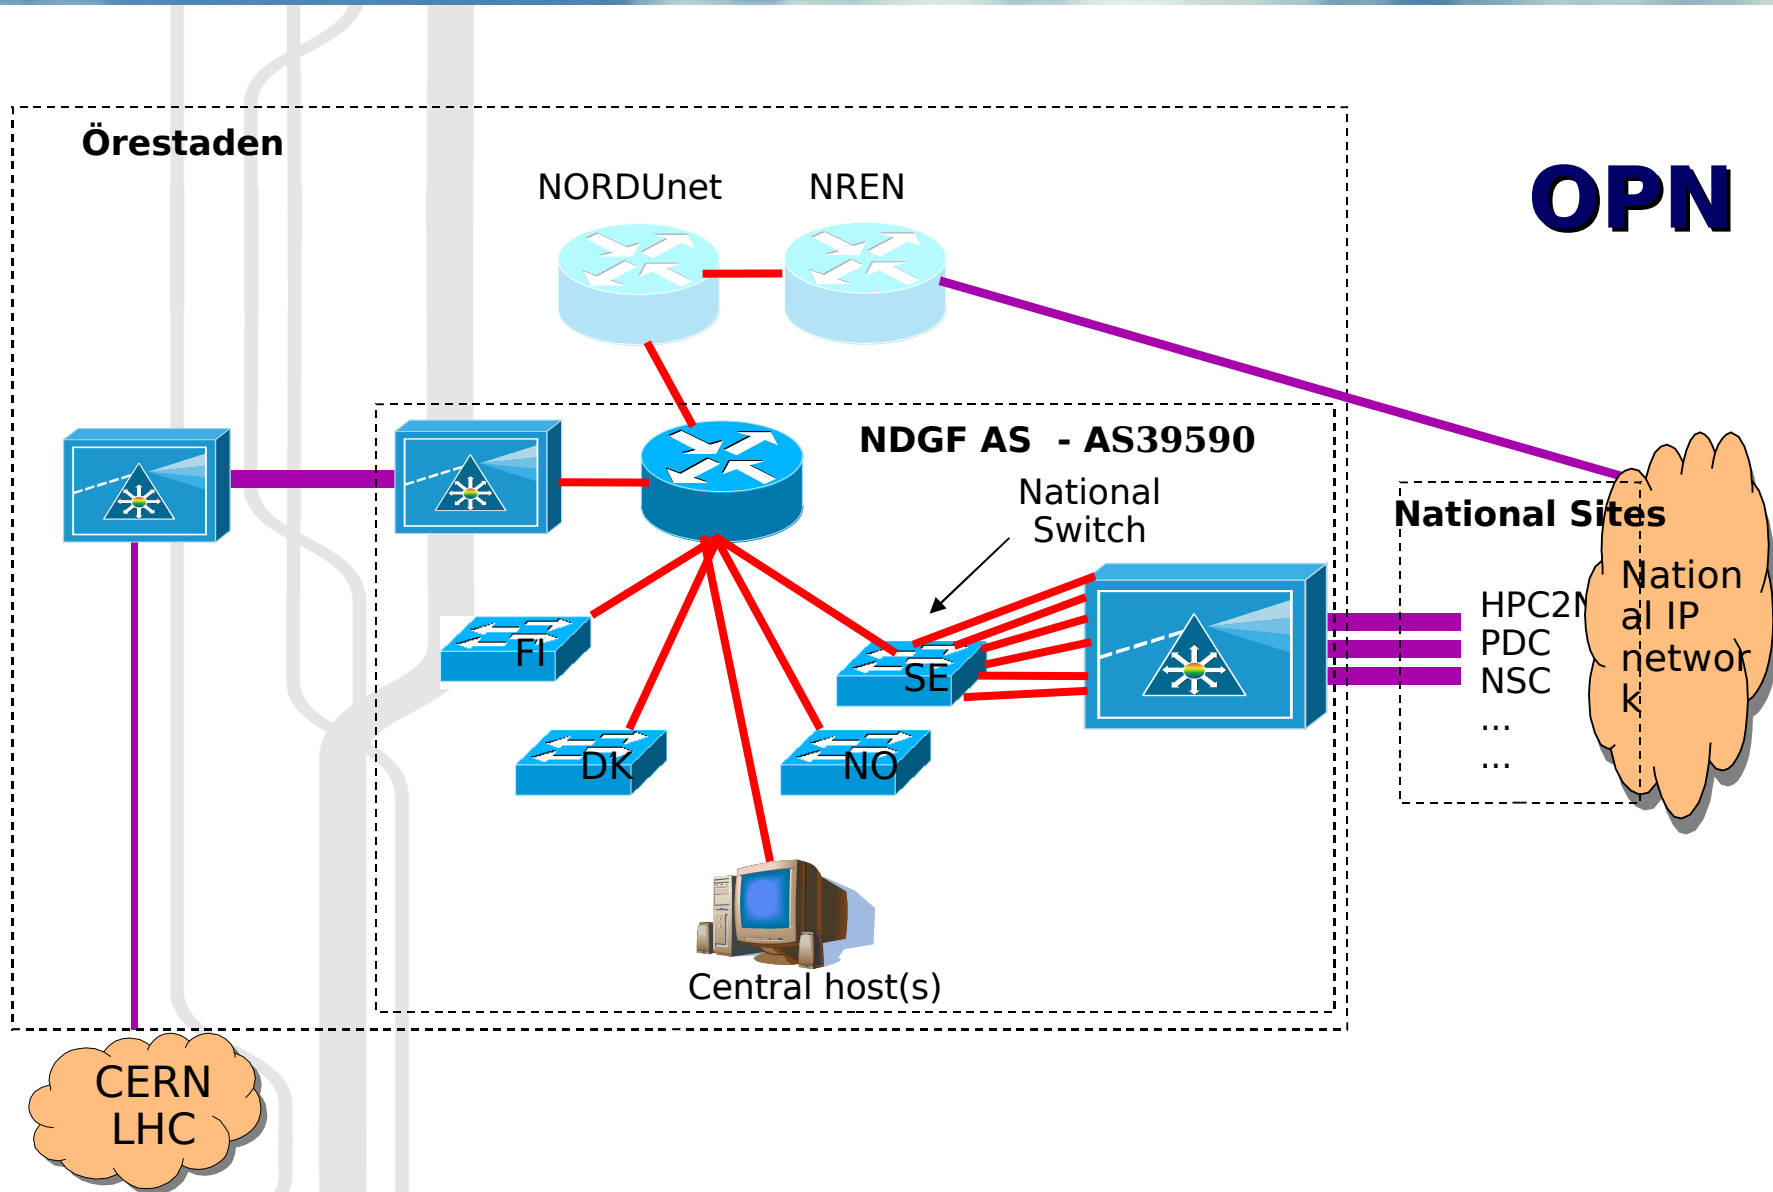
- Today NDGF is connected directly with GEANT 10Gbit fiber to CERN
- Inter-Nordic shared 10Gbit network from NORDUnet
- A Dedicated 10Gbit LAN covering all dTier-1 centers next year



- 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




The Infrastructure: Networking



- National academic HPC site
 - Typically 2-4 100+ node clusters with shared usage (heavy MPI usage)
 - Time on clusters granted nationally or locally
 - Also participates in grid initiatives, like NDGF
 - Expanding infrastructure (power&cooling)

- Current NDGF resources
 - 8 hp dl320s storage services – 60TB disk
 - 100 old computing nodes (single p4)
- Planned NDGF resources this year
 - ~40 new computing nodes (dual quad-core)
 - A few tens TB more disk
 - A few tens TB tape

- 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



- Current status
- Moving in march
- 400kW should last about 2-4 years
- Machine room expansion at many other NDGF sites

