



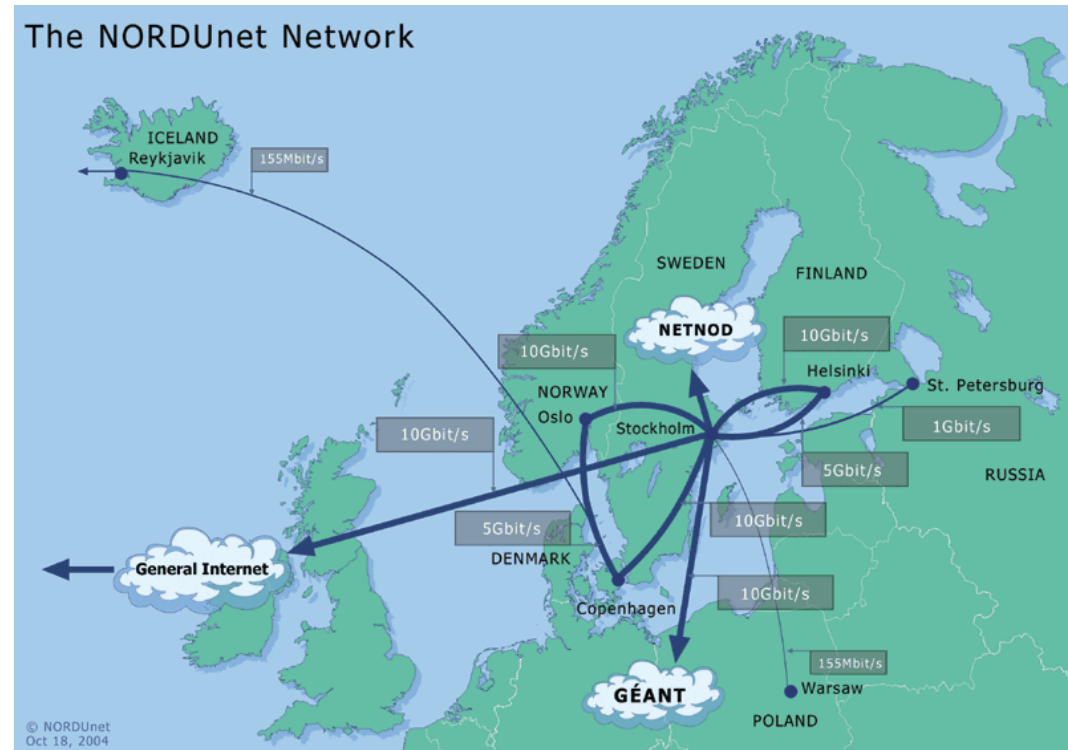
NDGF – a Joint Nordic Production Grid

Lars Fischer

7th Annual Workshop on Linux Clusters for Super Computing
Linköping, 18 October 2006

- NORDUnet – The Nordic Network collaboration
- NDGF – A Collaborative Grid Production Facility
- NDGF – organization, role, and structure
- ARC Middleware and the NorduGrid Consortia
- An Example: The Nordic WLCG Tier-1
- Summary

- The Nordic Regional Research and Education Network (REN)
- Owned by the 5 Nordic national RENs
- 25 years of Nordic network collaboration
- Leverage National Initiatives
- Participates in major international efforts
- Represent Nordic NRENs internationally, gateway to Nordic area



- 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
 - Strategic planning ongoing.

- "...to establish a Nordic Data Grid Facility and to involve Nordic countries in European and global co-operation in data sharing in a variety of fields."
- To *coordinate* and *facilitate* the creation of a Nordic e-Infrastructure sharing platform
- To enable Nordic researchers to participate in major international projects
- To optimize and standardize use of resources
- To optimize Nordic participation in international projects

- NorduGrid project (2001-2003)
 - Middleware: Architecture and First Implementation
 - Grid test-bed – in operation since August 2002
- Nordic grid middleware: ARC
 - Deployment of ARC at computing centers and clusters in Scandinavia
 - Successful participation in LHC Service Challenges
 - Forming of the NorduGrid middleware consortia (2005)
- NDGF Pilot Project
 - Two-year trial 2004-2005
 - Deployment of a pilot production facility
- NDGF 2006-2010

- Integration with NORDUnet
 - Shared management structure
 - Shared administration and office
 - Separate board of directors
 - Separate funding
- Mix of centralized and distributed organization
 - Central management and administration
 - Central key technical and project coordination functions
 - Central support systems and tools
 - Distributed production grid management (w/partners)
 - Distributed node management (w/partners)
 - Distributed software development (w/partners)

- Organization
 - CEO and CTO shared with NORDUnet
 - Technical Coordinator
 - Software Coordinator
 - Node Coordinators (one per country)
 - E-Science Project Coordinators
 - Software Developers (6 FTE)
 - About 15 people in all
- Governance
 - Steering Board w/national representation
 - Advisory Board per supported e-Science Project

- A Nordic Production Grid leveraging national resources
 - Extending the Network with Compute Services
 - Create a single production grid facility from existing resources (supercomputers, clusters, storage) in participating countries
 - A production grid: Processing, Storage, AAI, Services
- Middleware Development for Collaborative e-Science
 - Deliver the middleware required by e-Science projects
 - Leverage existing middleware initiatives (i.e., ARC, KnowARC, NORDUnet3, ...)
 - Hardening, usability, from proof-of-concept to production code
 - Development of facilities not currently available

- Provide a single point of contact for e-Science initiatives
 - A place to go for researchers initiating projects
 - Support projects in high-energy physics, nanotech, earth science, astronomy, medicine, etc.
- Services for user groups and applications
 - Portals for user groups
 - Gridifying major applications
- Host and Coordinate major e- Science projects
 - Coordinate technical implementation
 - Coordinate collaboration
 - Represent project internationally
 - Example: Nordic Tier-1 for the CERN LHC

- Represent the Nordic Grid community internationally
 - Allows the Nordic Grid community to speak with one voice
 - Creates visibility for Nordic technology and solutions
 - Creates a single Nordic representation in major e-Science projects
- Single Point of Entry for international partners
 - To reach the Nordic Grid community, talk to us
 - Point of contact for e-Science projects, middleware development, and grid facility deployment
- Leverage NORDUnet international contacts and activities
 - Joint participation in conferences and workshops
 - Joint representation of Nordic networking and grid views

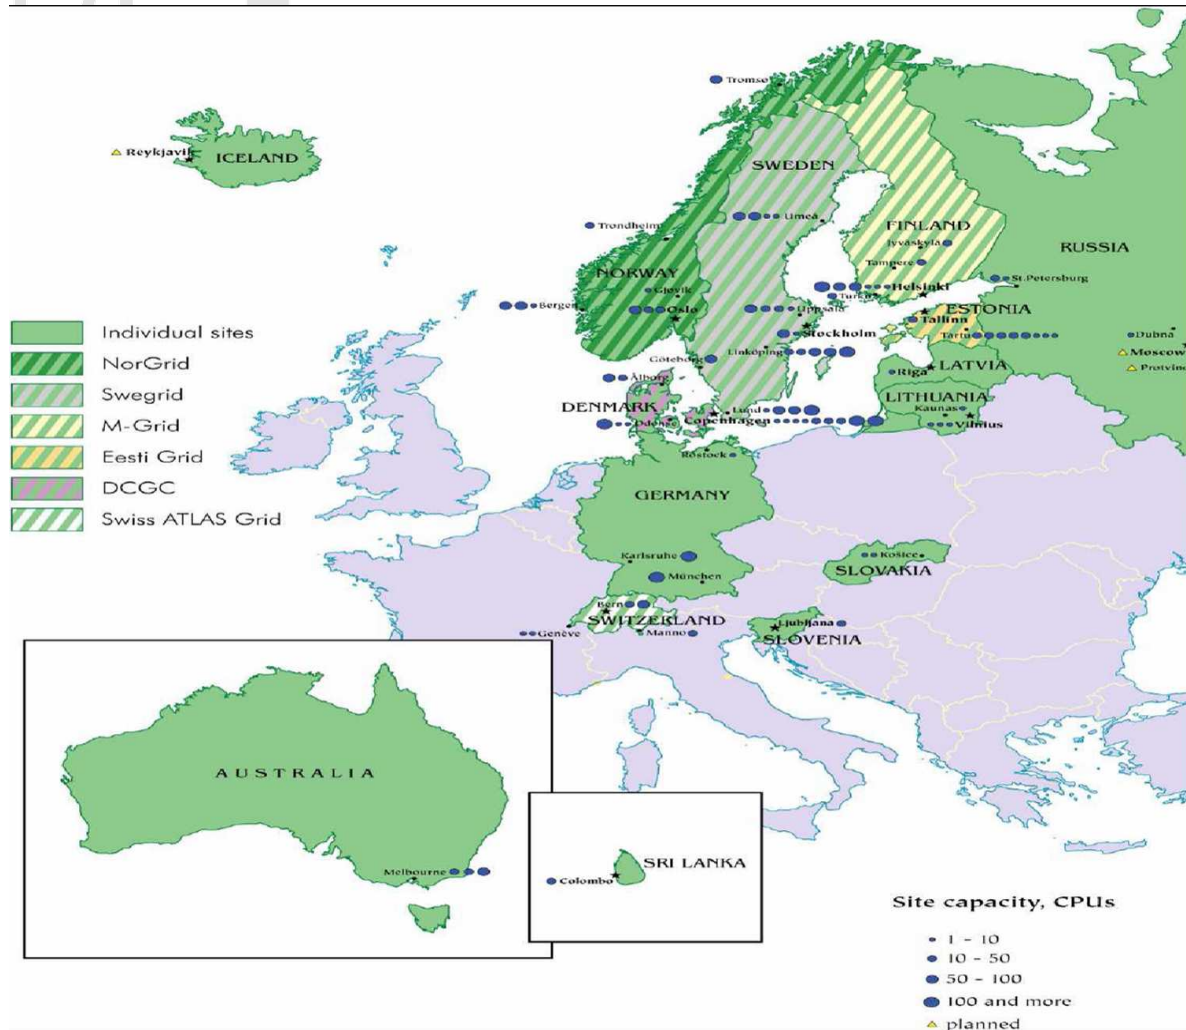
- NDGF is a Facilitator and an Integrator, not a resource owner
 - NDGF does not own a network
 - NDGF does not own computing resources
 - NDGF does not own the middleware platform
 - NDGF provides human resources
- Resources are owned and managed by partners
- NDGF Partnerships
 - NORDUnet
 - National Grid Projects (DCSC, SweGrid, NorGrid,)
 - National Computing resources
 - NorduGrid middleware consortia

- ARC – the Advanced Resource Connector
- Open Source Grid Middleware platform
 - Available for free download
 - Source code available for modification and incorporation in new projects
- An Open Process for Collaboration
 - Open for international participation
 - Open for contribution of code and ideas
 - Open for deployment
 - "Rough Consensus and Working Code"
- Controlled by the NorduGrid middleware consortia
- Think of ARC as Linux, NDGF as a Linux Distribution

- ARC Features
 - ARC has no single point of failure
 - ARC is non-intrusive (does not require changes to existing computing resources)
 - ARC is easy to install and deploy
- LCG Interoperability
 - The ARC middleware interoperates with the LCG middleware
- Co-existence
 - The ARC middleware can co- exist with other types of middleware

- Open Source development process
 - Core group of developers (5)
 - University researchers & developers, student projects
 - Extensions for specific projects
- A number of specific, funded development projects
 - KnowARC (EU-funded ARC development, 3 MEUR)
 - Nordunet3: New and Innovative Services for NorduGrid (Nordic funded ARC development, 1 MEUR)
- Nordic and International partners
 - ARC development process is open for contributions from the international Grid community
 - The NorduGrid consortia welcomes collaboration with researchers and developers worldwide

- NorduGrid collaboration is to deliver a robust, scalable, portable and fully featured solution for a global computational and data Grid system.
- Goals of the NorduGrid consortia:
 - Develop, coordinate, and support the ARC middleware
 - Define strategic directions for development of the ARC middleware following latest tendencies in the Grid technologies.
 - Promote ARC middleware development and deployment.
 - Contribute to Grid standardization
- Consortia governed by Memorandum of Understanding
 - Currently about 20 partners
 - The consortia is open to new partners, non-binding



- 60 Sites
- 13 countries
- 6000 Processors
- 60 TB storage
- 1600 Grid users
- HEP, Bio, Chemistry, Material Sciences, ...

Grid Monitor

2004-05-06 CEST 02:15:09

Processes: ■ Grid ■ Local

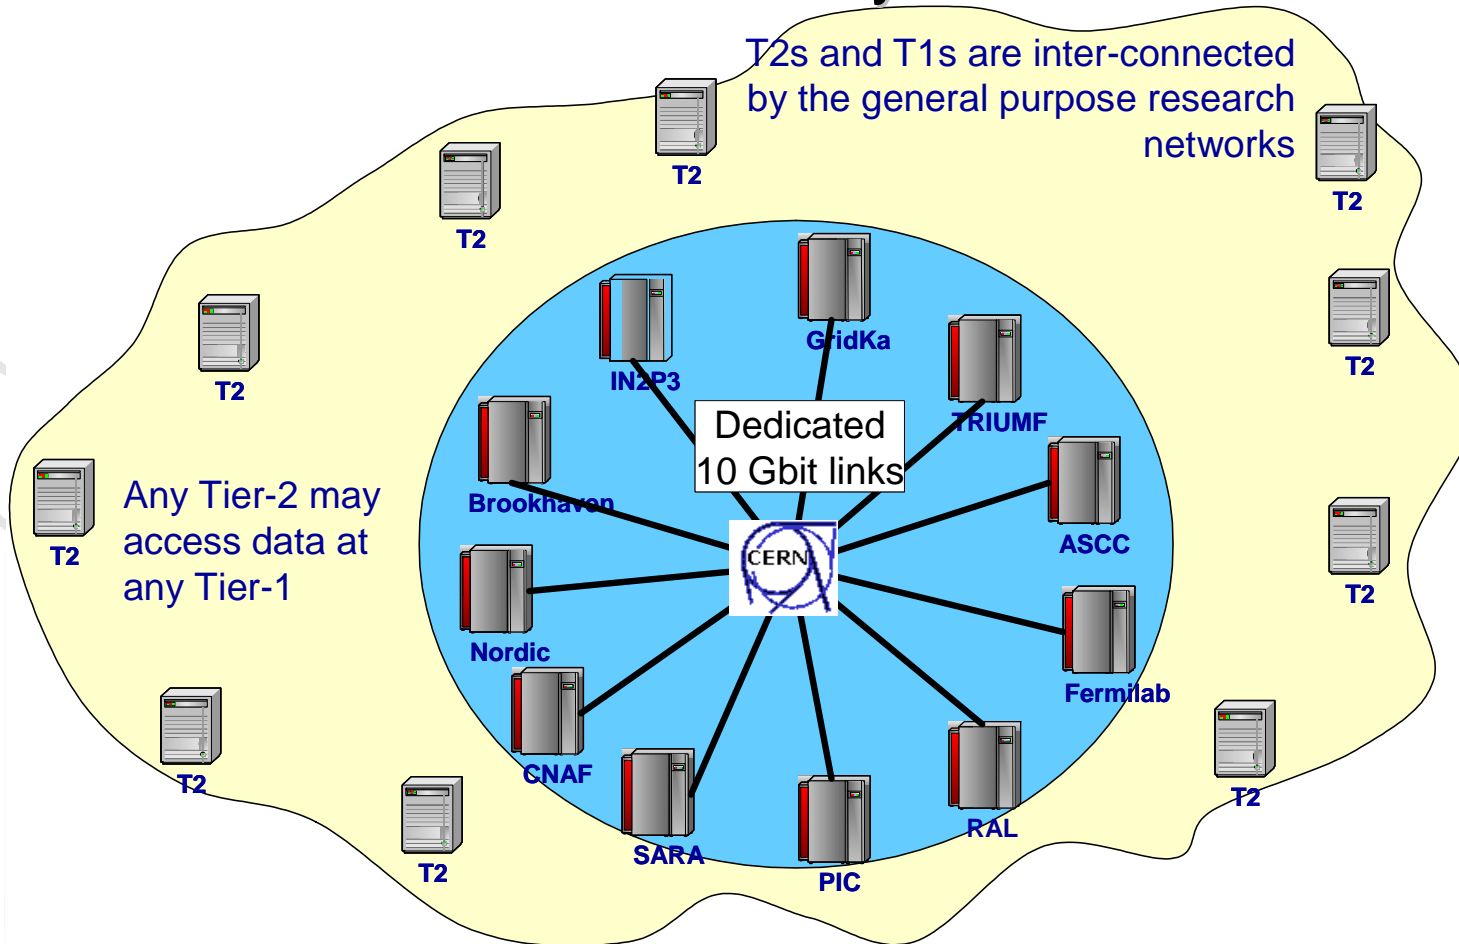
Site	CPU's	Load (processes: Grid+local)	Queueing
Australia			
ATLAS (UniMelb)	30	0+0	0
DistLab (DIKU)	9	0+0	0
Benedict (AAU/DCGC)	32	0+0	0
Horseshoe (DCSC/SDU)	561	0+497	40
Denmark			
NBI GRID	4	3+0	0
HEPAX1	1	0+0	0
Morpheus	18	0+0	0
Theory (DCSC/KU)	104	0+49	1
VCR (VideoRecorder)	1	1+0	0
Estonia			
CMS on CERN Linux	1	0+0	0
CMS test cluster	1	1+0	0
Finland			
CSC Kirppu	1	0+0	0
Hirmu Cluster (HIP)	16	0+0	0
Alpha (HIP)	1	0+0	0
Norway			
Parallab IBM Cluster	58	0+58	65
Bergen Grid Cluster	4	2+0	0
Oslo Grid Cluster	36	0+0	0
Gjovik Grid Cluster	2	0+0	0
Slovakia			
UPJS GRID	1	0+0	0
Slovenia			
SIGNET	42	0+18	0
Sweden			
Bluesmoke (SweGrid, NS>	100	62+38	93
Kosufy farm	66	30+0	0
Grendel	14	0+0	5
ISV	4	2+0	0
Hagrid (SweGrid, Uppma>	94	0+0	0
Hive (SweGrid, UNICC)	99	0+66	0
Ingrid (SweGrid, HPC2N)	101	101+0	330
Ingvar (NSC)	31	0+27	2
Monolith (NSC)	394	0+344	222
Quark Cluster	7	2+0	0
Seth (HPC2N)	202	0+133	36
Beppe (SweGrid PDC KT>	92	87+4	238
Sigrld (SweGrid, Luna>	99	0+76	3
Switzerland			
HIP CH	1	0+0	0
TOTAL	34 sites	2227 291 + 1302	

- ARC is Open Source and Free Software
 - Download the middleware
 - Install on a cluster
 - Optionally register a NorduGrid indexing databases (GIIS) to collaborate with other ARC sites
- NorduGrid website: <http://www.nordugrid.org>
 - Middleware download
 - Documentation
 - Presentations
- Join the ARC mailing list



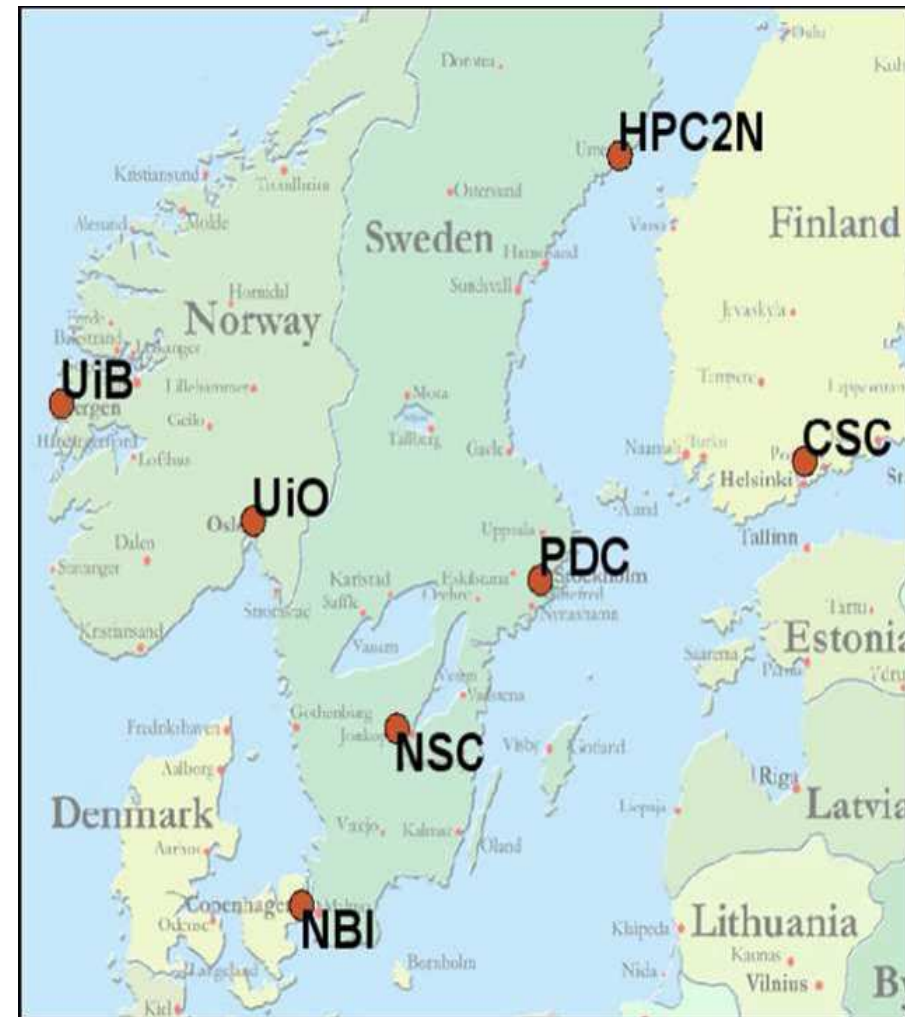
- Will require vast amounts of storage and computation
- Distributed on the Worldwide LHC Computing Grid (WLCG)

T0/T1/T2 Interconnectivity

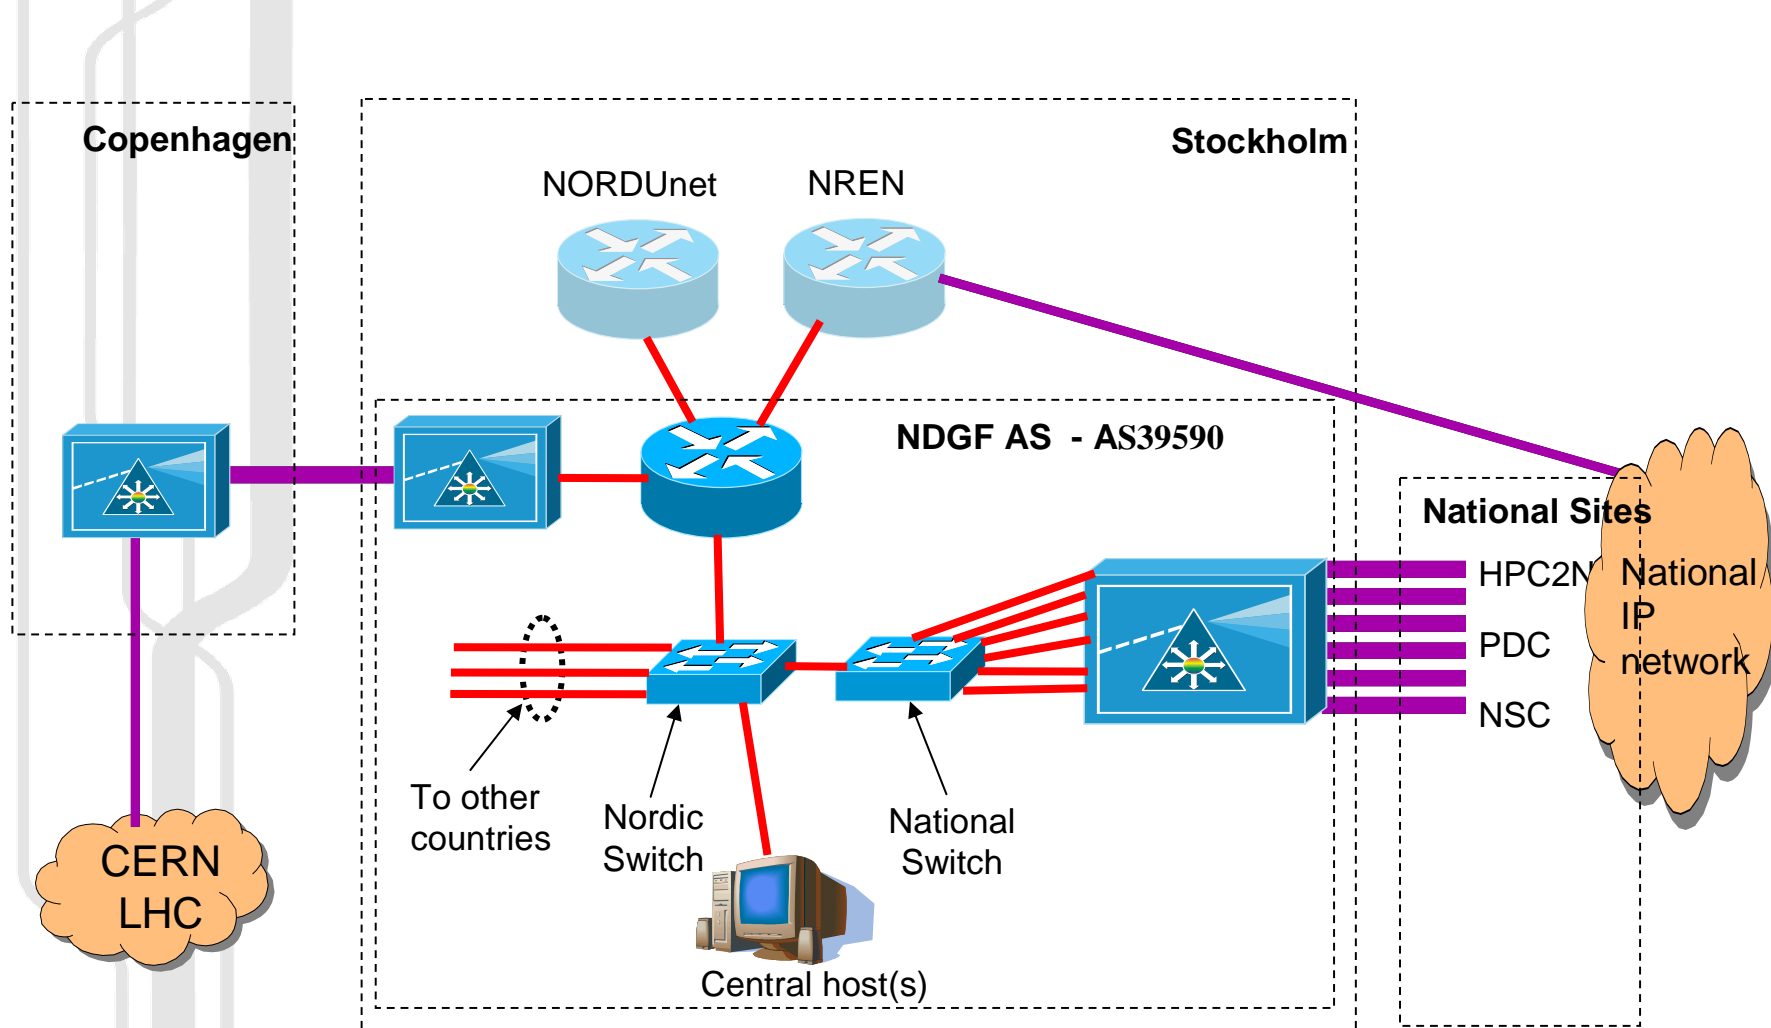


- Nordic Tier-1 for CERN LHC Experiments
 - One Nordic country cannot host a Tier-1 alone
 - ALICE (4 countries), ATLAS (DK, NO, SE), CMS ()
- NDGF will
 - Host the Nordic LHC Tier-1
 - Co-ordinate network, storage, and computing resources
 - Co-ordinate towards CERN and LHC project partners
- NDGF creates one technical facility to host the Tier-1
 - NDGF Tier-1 Network, connecting to CERN and to participating sites
 - NDGF Tier-1 computing and storage infrastructure, employing national grid resources and national supercomputing centers

- Appears as a Single Site
- Has one interface towards CERN (One SRM endpoint)
- The Storage is distributed
- The Computing Resources are distributed – like most other Tier-1's
- Has Storage and Computing Resources attached to a "longreach" LAN
- Most resources run ARC
- LCG-ARC interoperability, ALICE ARC-VOBOX, ATLAS ARC-DDM



- Connect all participating sites
 - Appear as a single LAN inside the NDGF Tier-1
 - Shared access to computing and storage resources
 - Creates one virtual site from participating sites
- Provides a application-specific IP network
 - Interface to LHC OPN at GEANT2 PoP in Copenhagen
 - Tier-1 – Tier-1 links from Copenhagen, using CBF
 - Peering with LHC OPN, NORDUnet IP, Nordic NREN IP.
 - AS39590, announce NDGF IP-range
- Outside access to NDGF Tier-1 resources through IP peering (i.e., from Tier-2's).



- NDGF is a Computational and Grid Production Facility for 4 Nordic countries
- NDGF supports Nordic e-Science Projects
- NDGF is the interface to Grid activities in the Nordic countries.
 - Single Point of Entry for collaboration, middleware development, and deployment, and for e-Science projects
 - Represents the Nordic Grid community internationally
- NDGF coordinates activities and leverages existing resources - NDGF does not own resources or middleware

Questions?

<http://www.ndgf.org>

<http://www.nordugrid.org>

<http://www.nordu.net>

Lars Fischer (CTO, NDGF), lars@nordu.net