

Nordic Data Grid Facility

Lars Fischer, NORDUnet

Chinese-Nordic Workshop
Trondheim, 19-21 April 2006

Overview

- NDGF - A Collaborative Grid Production Facility
- ARC Middleware: An Open Grid platform, and an Open Process for Participation, Contribution, and Use
- NDGF is the interface to Grid activities in the Nordic countries.
 - Single Point of Entry for collaboration, middleware development, and deployment, and e-Science projects
 - Represents the Nordic Grid community internationally
- NDGF coordinates activities - NDGF do not own resources or middleware

What is the NDGF?

- A Co-operative Nordic 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 grid projects (i.e., Nordic LHC Tier-1)
 - Develop grid middleware
- 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.

History of Nordic Grid Collaboration

- Collaboration initiated with Nordunet2 programme (2001-2003)
 - "Nordic Testbed for Wide Area Computing and Data Handling"
 - Development of middleware: Architecture and First Implementation
 - Grid test-bed - in operation since August 2002
- Nordic grid middleware: ARC (Advanced Resource Locator)
 - Deployment of ARC at computing centres 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, based on national resources.
- NDGF 2006-2010

Organization of the NDGF

- 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)

NDGF Technical Platform

- 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
 - Provide grid project storage - a current challenge
- Middleware Development for Collaborative e-Science projects
 - Deliver the middleware required for a production grid for e-Science
 - Leverage existing middleware development initiatives (i.e., ARC)
 - Hardening, usability, from proof-of-concept to production code
 - Development of facilities not currently available (i.e., storage sharing)

Supporting Nordic e-Science

- 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

International Collaboration

- 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 Partnerships

- 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
- Resources are owned and managed by partners
- NDGF Partnerships
 - NORDUnet
 - National Grid Projects (DCSC, SweGrid, NorGrid,)
 - National Computing resources
 - NorduGrid middleware consortia

ARC – an Open Grid Middleware

- 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 Overview

- **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, designed for the CERN LHC project

- **Co-existence**

- The ARC middleware can co-exist with other types of middleware

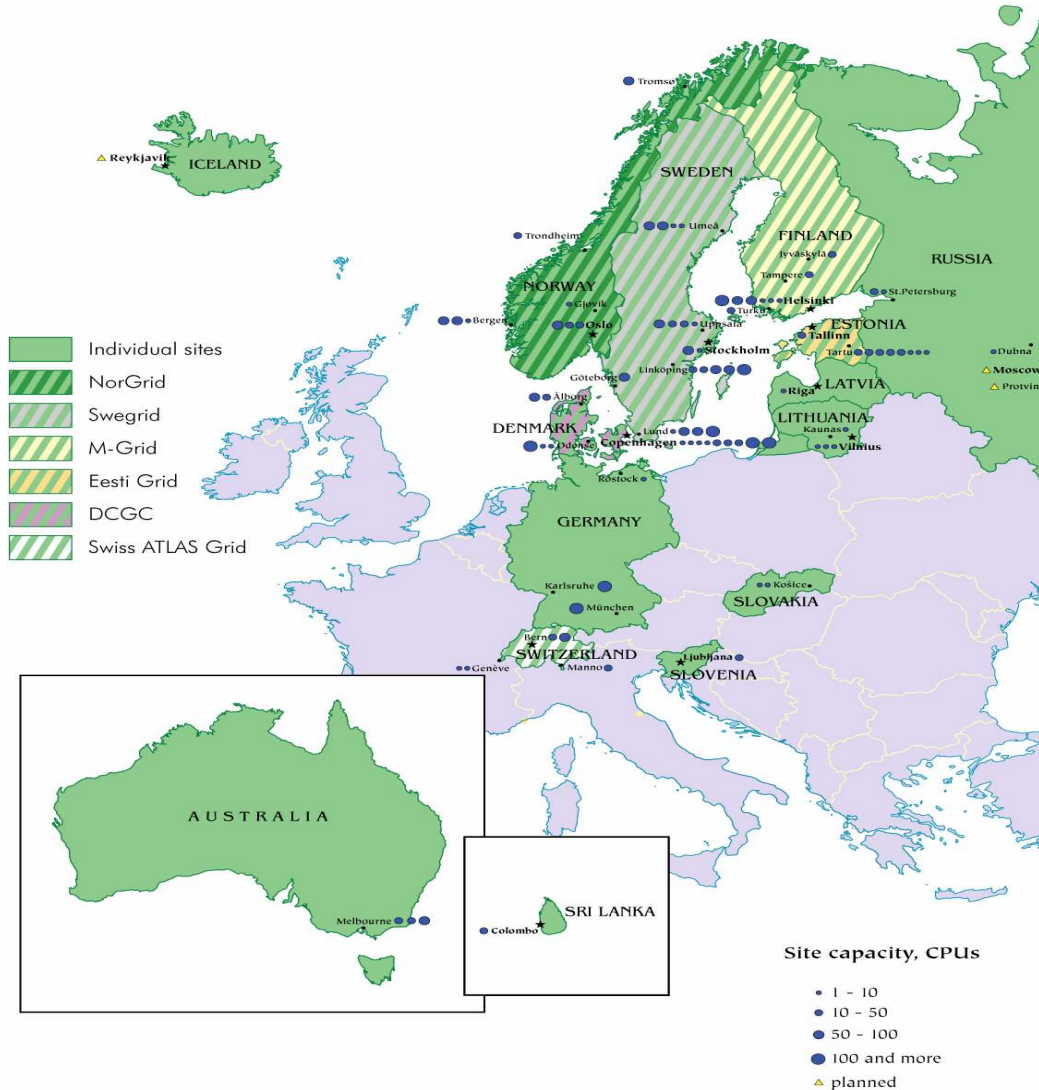
ARC Development

- **Open Source development process**
 - Much like Linux and other Open Source projects
 - Core group of developers (5)
 - University researchers & developers
 - Student projects
 - Extensions for specific projects
- **A number of collaborative projects**
 - KnowARC (EU-funded ARC development, 6 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 Consortia

- Based on Memorandum of Understanding between initial partners
 - The universities in Helsinki (Finland) , Lund (Sweden), Uppsala (Sweden), Oslo (Norway), Copenhagen (Denmark)
 - Currently about 20 partners in the consortia
 - The NorduGrid Consortia is open to new partners
- 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 and support the ARC middleware.
 - Coordinate contributions to the ARC code.
 - Define strategic directions for development of the ARC middleware following latest tendencies in the Grid technologies.
 - Promote ARC middleware solutions in such areas as Grid development, deployment and usage.
 - Contribute to development of Grid standards, e.g. via Global Grid Forum.

ARC Deployment



- 60 Sites
- 6000 Processors
- 60 TB storage
- 1600 Grid users

ARC Grid Monitor

Grid Monitor

2004-05-06 CEST 02:15:09



Processes: ■ Grid ■ Local

	Site	CPUs	Load (processes: Grid+local)	Queueing
	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
	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
	CMS on CERN Linux	1	0+0	0
	CMS test cluster	1	1+0	0
	CSC Kirppu	1	0+0	0
	Hirmu Cluster (HIP)	16	0+0	0
	Alpha (HIP)	1	0+0	0
	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
	UPJS GRID	1	0+0	0
	SIGNET	42	0+10	0
	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
	Sigrid (SweGrid, Luna>	99	0+76	3
	HIP CH	1	0+0	0
TOTAL		34 sites	2227 291 + 1302	

Getting Started with ARC


- 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

Case: NDGF LHC Tier-1

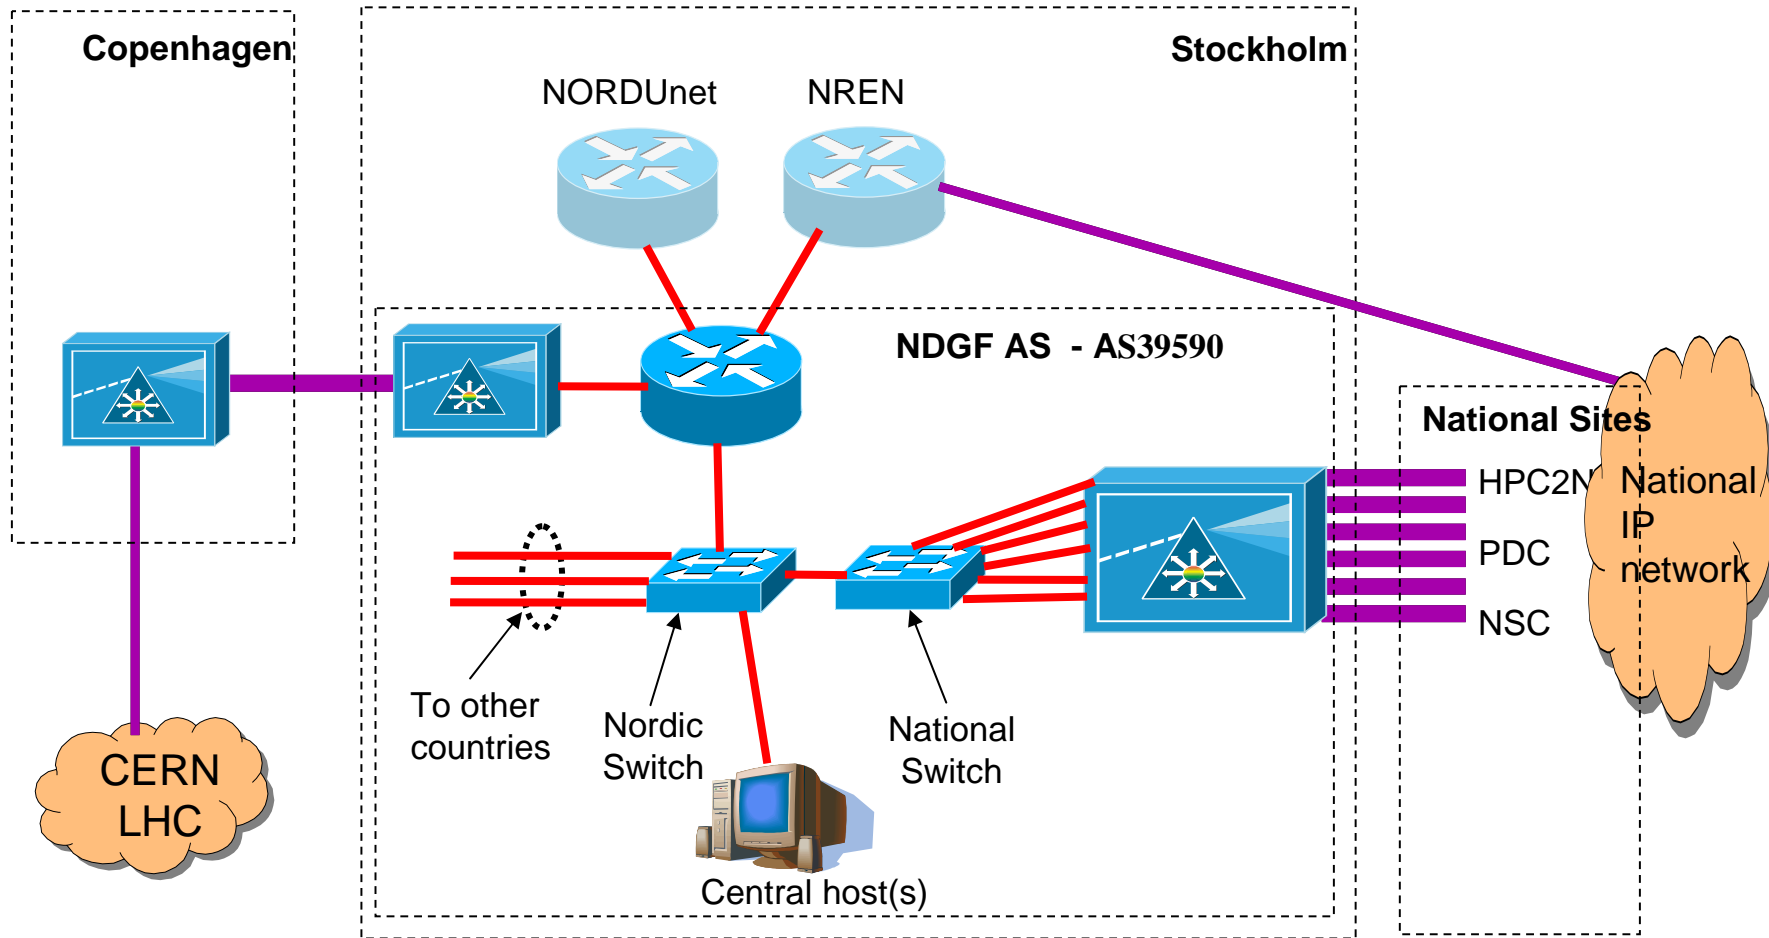
- Nordic Tier-1 for CERN LHC Experiments
 - None of the Nordic countries have the resources to host a Tier-1
 - The Nordic HEP community would like to have Tier-1
- NDGF will
 - Host the Nordic LHC Tier-1
 - Co-ordinate network, storage, and computing resources
 - Co-ordinate towards CERN and LHC project partners
- NDGF creates *one* technical facility to host the Tier-1
 - NDGF Tier-1 Network, connecting to CERN (via GEANT2) and to participating sites
 - NDGF Tier-1 computing infrastructure, employing national grid resources
 - NDGF Tier-1 storage infrastructure, deployed at national supercomputing centres

NDGF Tier-1 Sites

- 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 "long-reach" LAN
- Most resources run ARC
- LCG-ARC interoperability, ALICE ARC-VOBOX, ATLAS ARC-DDM



NDGF Tier-1 Optical Private Network



NDGF Tier-1 OPN - details

- Star Network, centre in Stockholm
- Interface to LHC OPN at GEANT2 PoP in Copenhagen
- Possible Tier-1 - Tier-1 links from Copenhagen, using CBF
- Will provide one IP network, with peering with LHC OPN, NORDUnet IP, Nordic NREN IP.
 - Will use AS39590 (ownership: ORG-NDGF1-RIPE)
 - Will announce NDGF IP-range
- Outside access to NDGF Tier-1 resources through IP peering
- Inside NDGF Tier-1 will appear as one LAN
 - Connect all participating sites
 - Provide inter-Tier-1 access to computing and storage resources
 - Will create one virtual site from participating national sites
- Layered LAN approach, one switch per country (where needed)

Thank You

Questions?

<http://www.ndgf.org>

<http://www.nordugrid.org>

<http://www.nordu.net>

Lars Fischer, lars@nordu.net