

Griding the Nordic Supercomputer Infrastructure

Anders Ynnerman

Director Swedish National Infrastructure for Computing
Linköping University
Sweden



Outline of presentation



- Status in the Nordic Countries
- Some GRID efforts in the Nordic Region
 - SweGrid testbed for production
 - NorduGrid - ARC
 - North European Grid
 - Nordic DataGrid Facility
- Identifying potential problems for Nordic Grid collaborations
- Path forward
- BEEGEE

DENMARK



HPC-infrastructure

- HPC program initiated 2001
 - Danish Center for Scientific Computing (DCSC)
 - Distributed program
- Location of facilities based on scientific evaluation of participating groups not on computer centers or infrastructure
 - Århus
 - Odense
 - DTU
 - Copenhagen University
- Budget
 - 2.1 MEuros (2001)
 - 1.6 MEuros (2002)
 - 1.5 MEuros (2003)
 - 1.4 MEuros (2004)

GRID-initiatives

- Copenhagen University is a partner in NorduGrid
- Danish Center for Grid Computing (DCGC) has just been established
 - Virtual center located at 5 sites
 - Main location at the Niels Bohr Institute, Copenhagen University
 - 0.75 MEuros/year
 - Operations and support of a Danish Grid
 - Research agenda with 10 senior researchers and 5 PhD students
 - Strong ties to NorduGrid and Nordic DataGrid Facility
- Participation in EGEE through NBI

NORWAY



HPC-infrastructure

- NOTUR – The Norwegian High-Performance Computational Infrastructure
 - Trondheim
 - Oslo
 - Bergen
 - Tromsø
- Budget
 - 168 MEuros (2000-2003)
 - 50% from reserach councils
 - Statoil, DNMI
 - In kind contributions
- New program similar to the Swedish program is being initated and placed under the NREN (UNINETT)

GRID-initiatives

- Metacenter goal – To establish a uniform and as seamless as possible, access to HPC resources
- Parallab at Bergen is participating in the Nordic Grid Consortium with PDC and CSC
- Parallab is contributing to EGEE security team
- Bergen and Oslo are partners in NorduGrid
- NORGRID project has been initiated

FINLAND



HPC-infrastructure

- Center for Scientific Computing (CSC) is the (main) center in Finland
- The largest HPC center in Scandinavia
- Runs services for academic users all over Finland and FMI
- Budget
 - Directly from the ministry of education
- Smaller university HPC-systems (clusters) are beginning to appear

GRID-initiatives

- CSC is one of the founding partners in the Nordic Grid Consortium
- CSC is contributing resources to NorduGrid
- CSC is a partner in DEISA
- Helsinki Institute of Physics is running several GRID projects and is contributing to the EGEE security team
- A Finnish Grid dedicated to Condensed Matter Physics is being built

SWEDEN



HPC-infrastructure

- Swedish National Infrastructure for Computing (SNIC) formed during 2003
- 6 participating centers
 - Umeå
 - Uppsala
 - Stockholm
 - Linköping
 - Göteborg
 - Lund
- National resource allocation (SNAC)
- Budget
 - 4.5 MEuros/year (Government)
 - 4.0 MEuros/year private foundations

GRID-initiatives

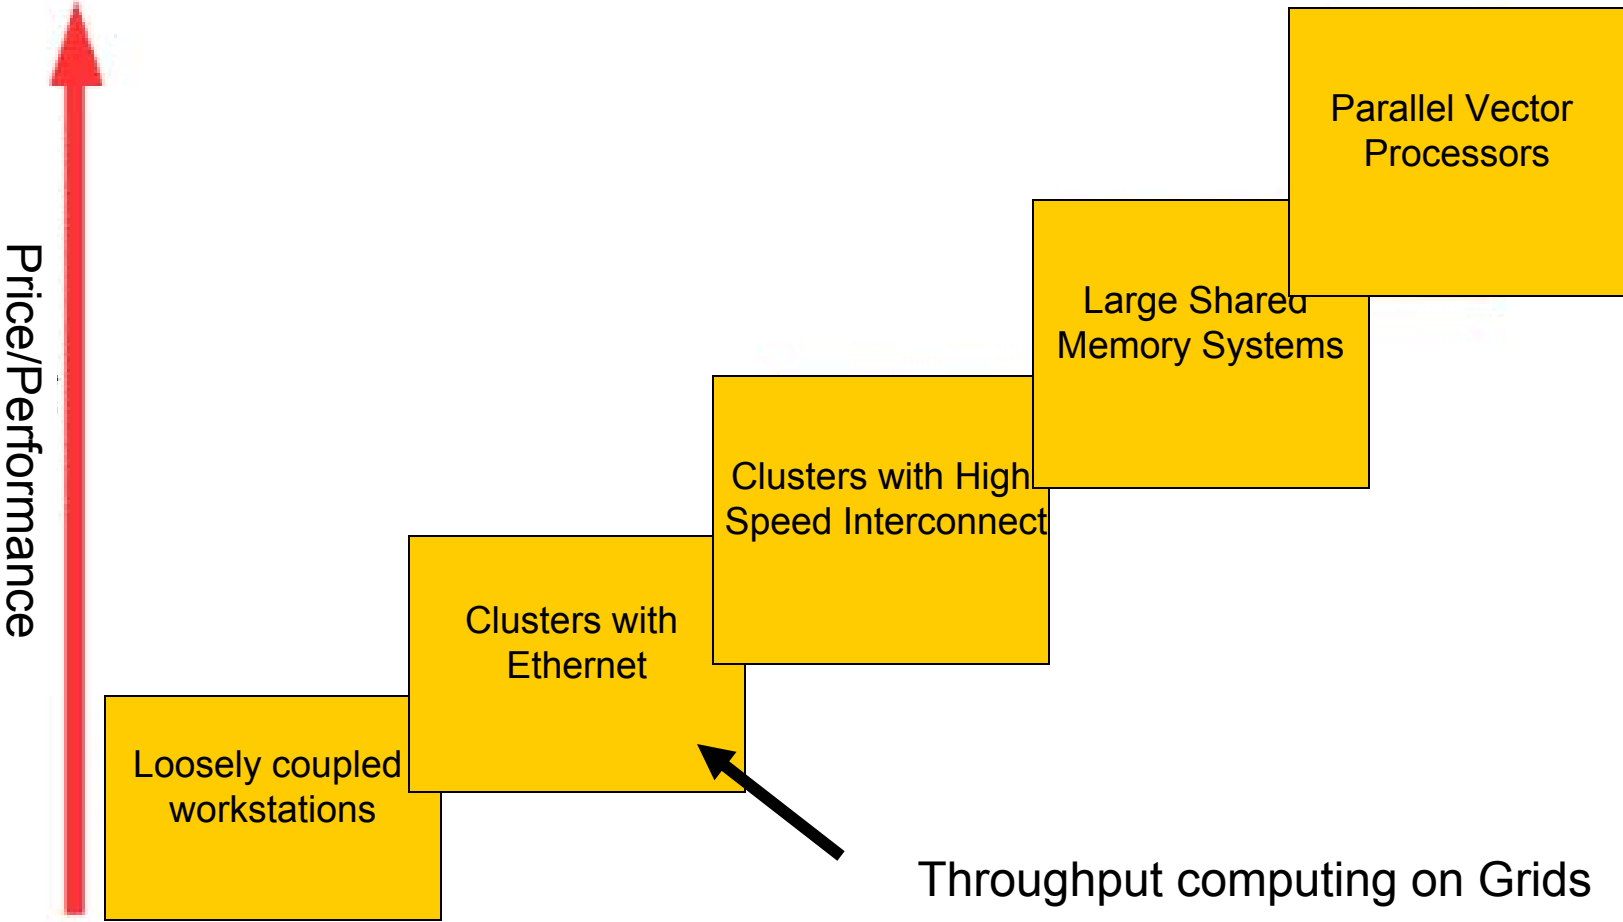
- Lund and Uppsala are partners in NorduGrid
- Largest resource contributions to NorduGrid are SNIC clusters at Umeå and Linköping
- SweGrid production Grid (600 CPUs) over 6 sites 120 Tb storage
- SNIC is hosting the EGEE Nordic Regional Operations Center
- Stockholm (PDC) is coordinating the EGEE security activity
- PDC is a partner in the European Grid Support Center

A Typical (Nordic) HPC Center



- Receives funding from a national agency (Research Council)
- Serves a large user community spanning “all” fields of science
- Users are allocated time through a (national) application procedure
- Operates several large scale computer services
- Is involved in HPC R&D efforts
- Accounting statistics and scientific reports show the results generated
- Sees Grid technology as the future

Grids in the context of HPC centers






SweGrid production testbed



- The first step towards HPC center Gridification
- Initiative from
 - All HPC-centers in Sweden
 - IT-researchers wanting to research Grid technology
 - Users
 - Life Science
 - Earth Sciences
 - Space & Astro Physics
 - High energy physics
- PC-clusters with large storage capacity
- Build for GRID production
- Participation in international collaborations
 - LCG
 - EGEE
 - NorduGrid
 - ...



SweGrid production test bed



- Total budget 3.6 MEuro
- 6 GRID nodes
- 600 CPUs
 - IA-32, 1 processor/server
 - 875P with 800 MHz FSB and dual memory busses
 - 2.8 GHz Intel P4
 - 2 Gbyte
 - Gigabit Ethernet
- 12 TByte temporary storage
 - FibreChannel for bandwidth
 - 14 x 146 GByte 10000 rpm
- 120 TByte nearline storage
 - 60TByte disk
 - 60 TByte tape (To be installed)
- 1 Gigabit direct connection to SUNET (10 Gbps)




SUNET connectivity




GigaSunet 10 Gbit/s

Typical POP at Univ.




“The snowman topology”






Persistent storage on SweGrid?

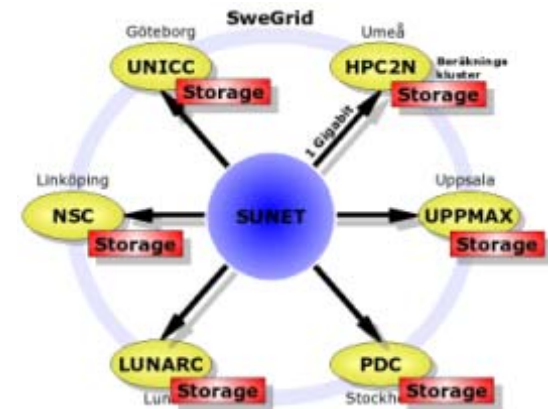


1

Size
Administration



2



3

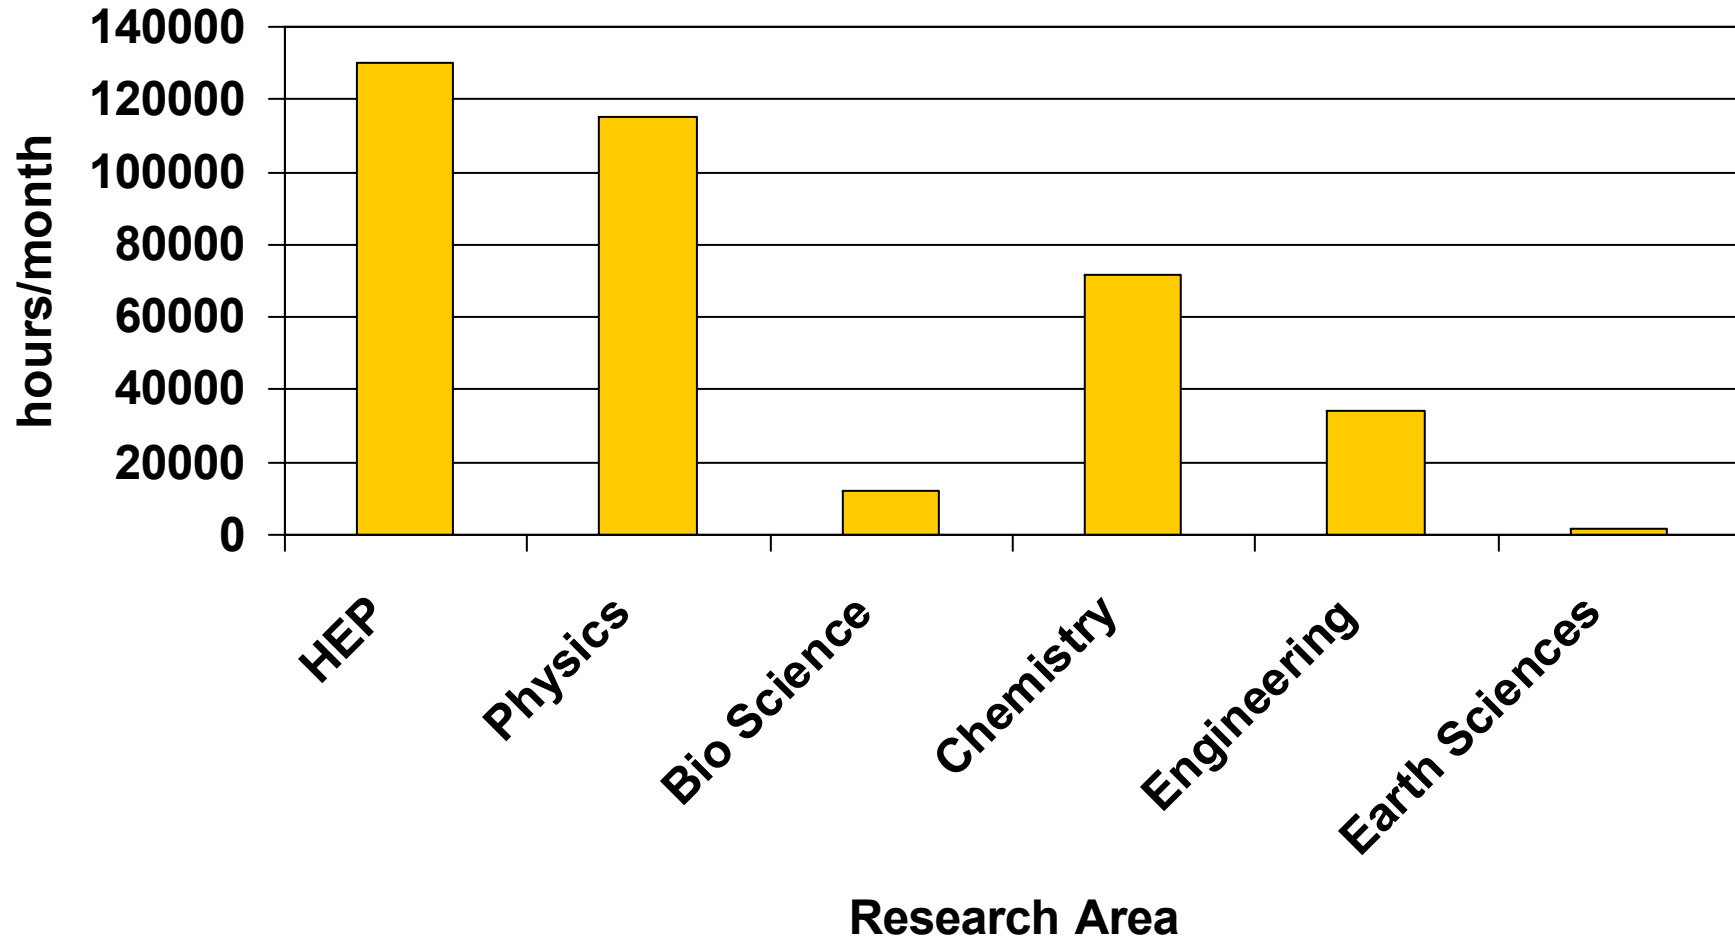
Bandwidth
Availability

SweGrid status



- All nodes installed during January 2004
- Extensive use of the resources already
 - Local batch queues
 - GRID queues through the NorduGrid middleware - ARC
 - 60 users
- 1/3 of SweGrid is dedicated to HEP (200 CPUs)
- Contributing to Atlas Data Challenge 2
 - As a partner in NorduGrid
- Currently deploying LCG-2
 - Investigating compatibility between ARC and LCG-2
- Forms the core of the Northern EGEE ROC
- Accounting is still missing ...

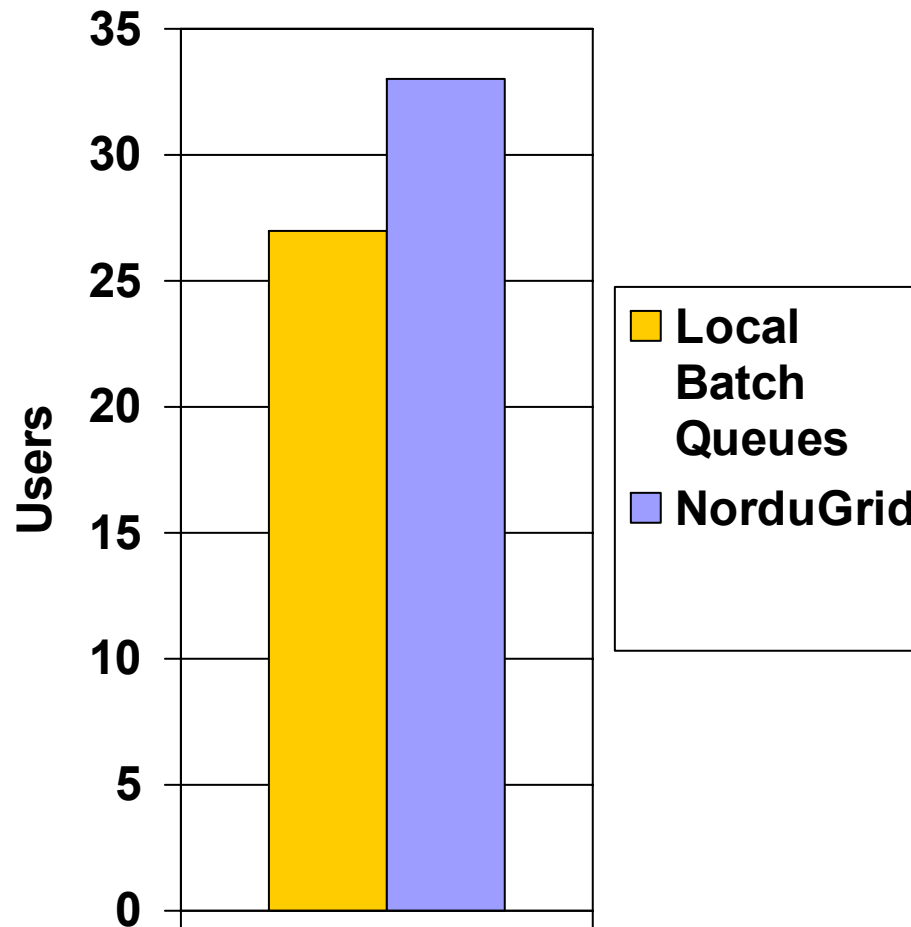
The first users of SweGrid



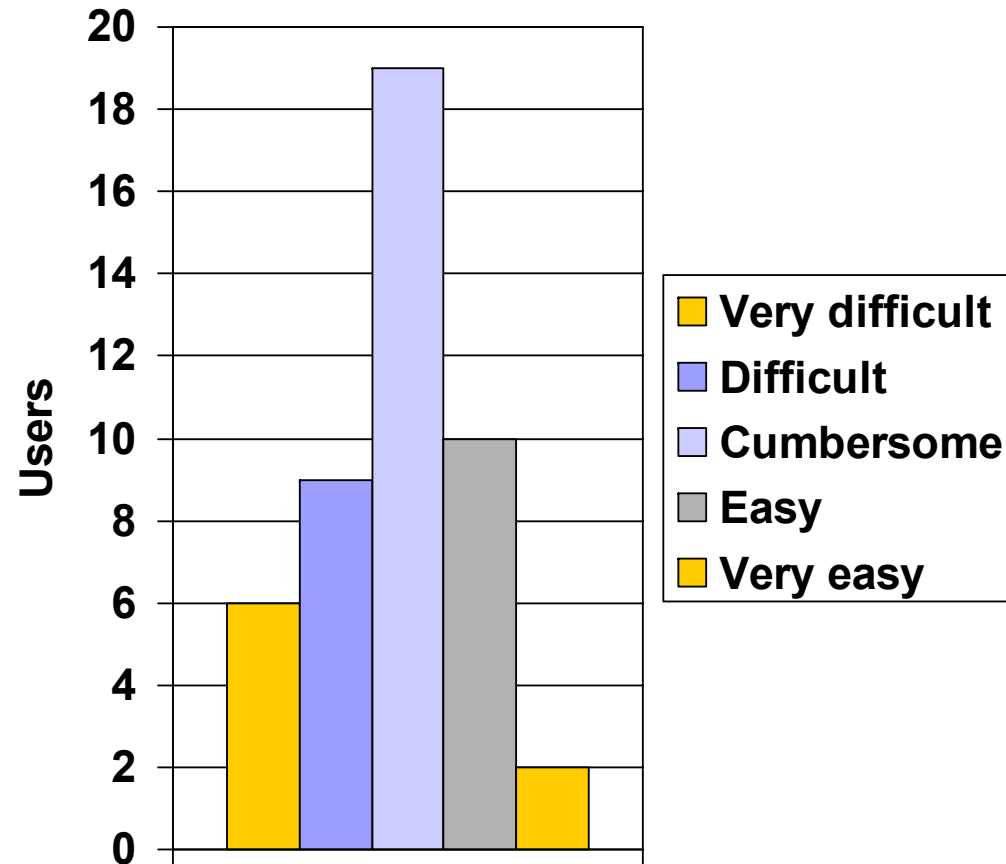


What did they think?


How have you accessed SweGrid Resources?




How have you found porting your applications to SweGrid to be?



What is your overall impression of use of SweGrid resources?



Do you think all supercomputing resources should be available on a Grid?



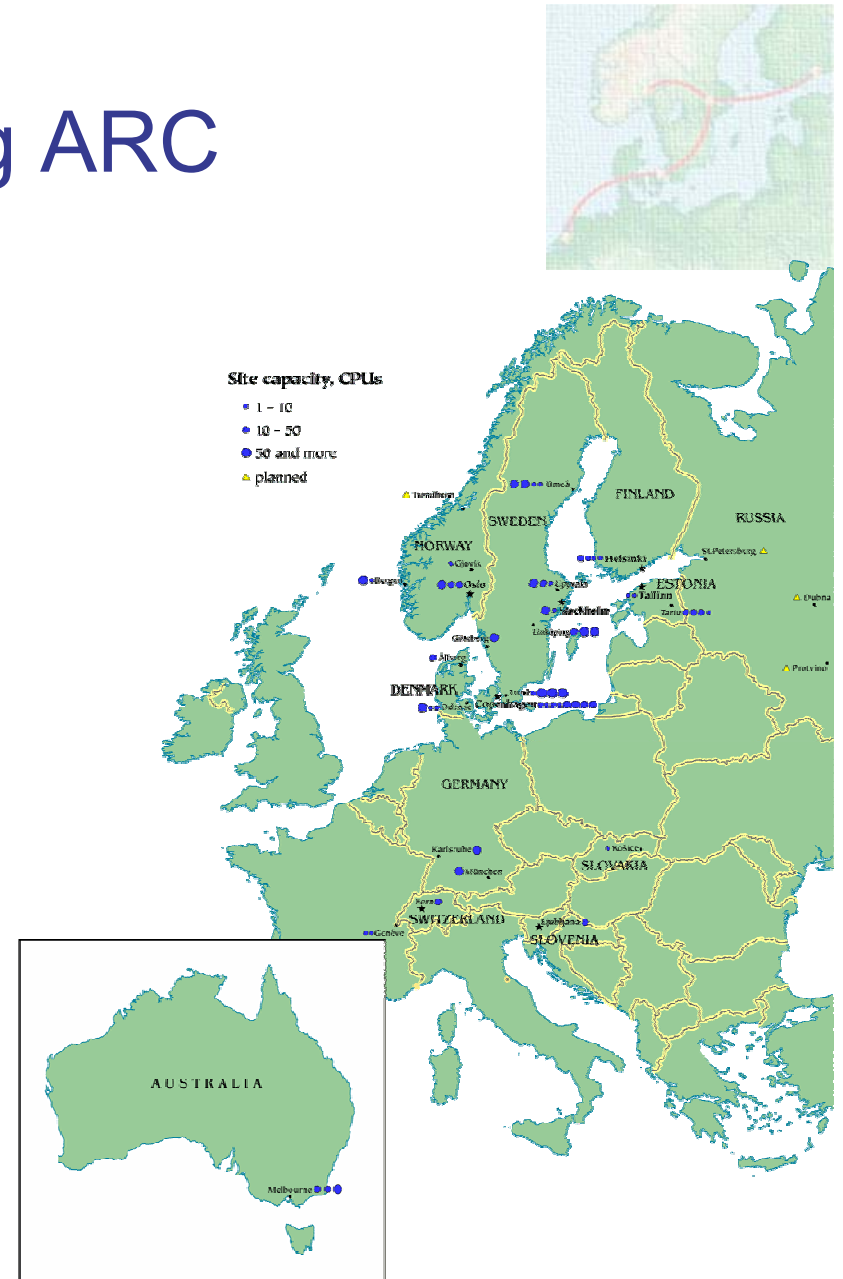
The NorduGrid project



- Started in January 2001 & funded by NorduNet-2
 - Initial goal: to deploy DataGrid middleware to run “ATLAS Data Challenge”
- NorduGrid essentials
 - Built on GT-2
 - Replaces some Globus core services and introduces some new services
 - Grid-manager, Gridftp, User interface & Broker, information model, Monitoring
 - Middleware named ARC
- Track record
 - Used in the ATLAS DC tests in May 2002
 - Contributed 30% of the total resources to ATLAS DC II
- Continuation
 - Could be included in the framework of the ”Nordic Data Grid Facility”
 - Close co-operation with EGEE/LCG

Resources running ARC

- Currently available resources:
 - 10 countries, 40+ sites, ~4000 CPUs, ~30 TB storage
 - 4 dedicated test clusters (3-4 CPUs)
 - SweGrid
 - Few university production-class facilities (20 to 60 CPUs)
 - Three world-class clusters in Sweden and Denmark, listed in Top500
- Other resources come and go
 - Canada, Japan – test set-ups
 - CERN, Russia – clients
 - Australia
 - Estonia
 - It's open, anybody can join or part
- People:
 - the “core” team grew to 7 persons
 - local sysadmins are only called up when users need an upgrade



A NorduGrid snapshot

2004-09-24 CEST 15:57:02

Processes: ■ Grid ■ Local



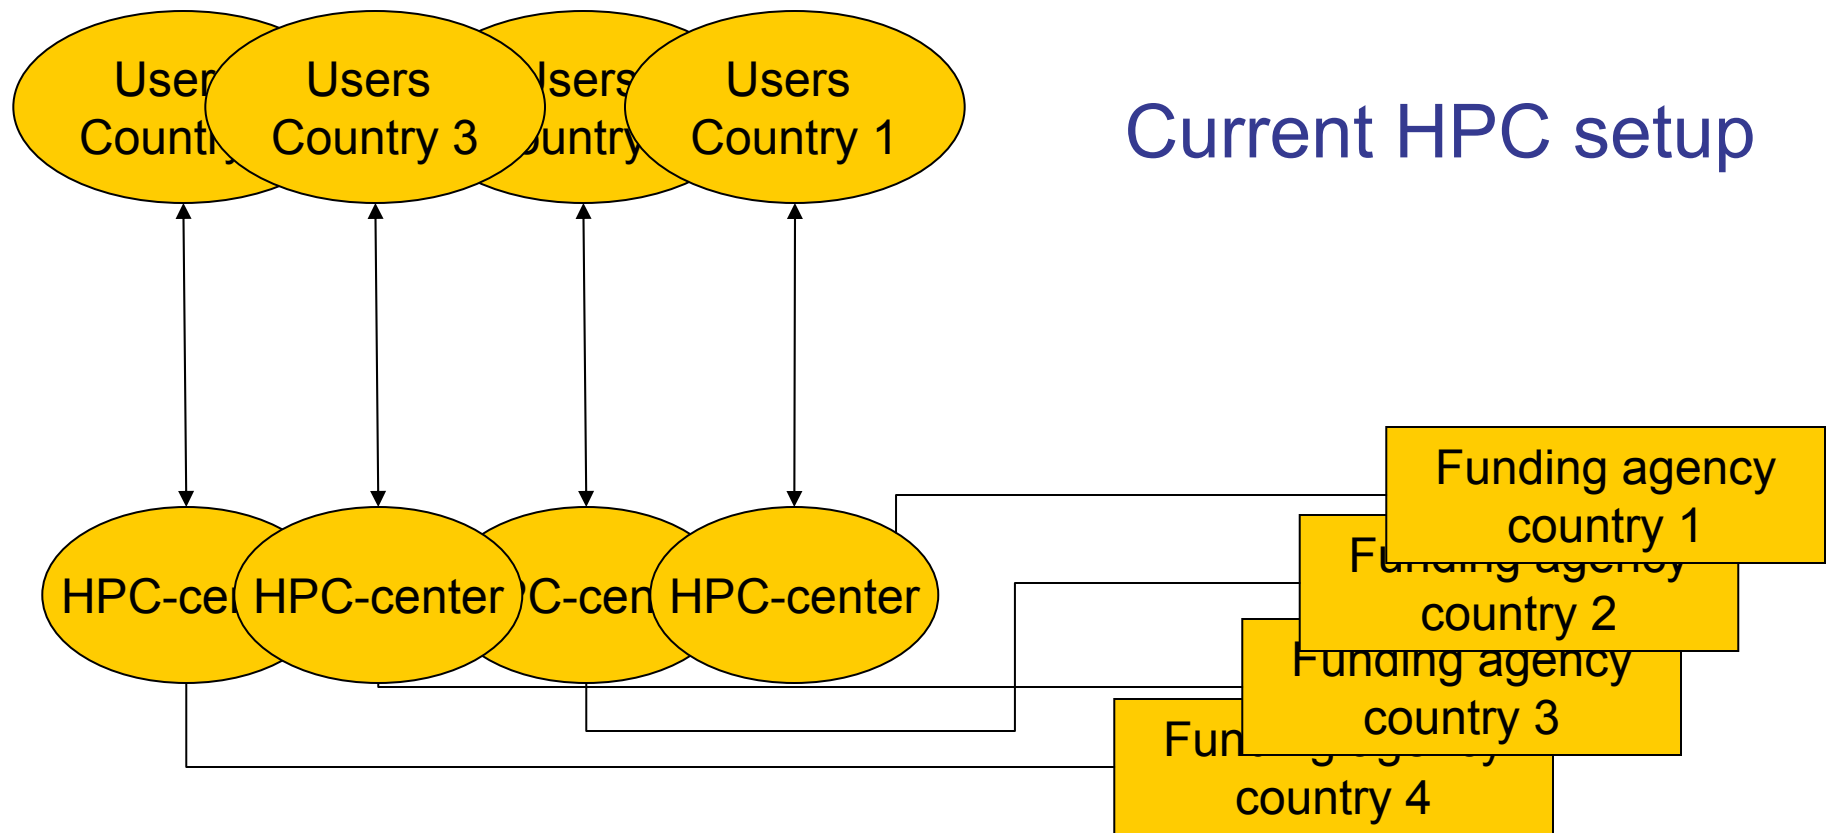
Country	Site	CPU's	Load (processes: Grid+local)	Queueing
	Atlas (UniMelb)	26	8+2	0+0
	Charm (UniMelb)	19	8+8	0+0
	Alfred (UniMelb)	88	8+27	8+1
	DistLab (DIKU)	10	8+8	0+0
	Aalborg Grid Gateway	46	46+8	27+0
	Niflheim (DCSC/DTU)	425	8+423	0+10
	Horseshoe (DCSC/SDU)	1195	9+883	0+0
	HEPAX1	1	8+8	0+0
	Morpheus	18	17+8	8+0
	Theory (DCSC/KU)	112	8+71	0+1
	VCR (VideoRecorder)	1	1+8 (queue down)	0+0
	UT IMCB Anakonda clus>	15	6+8	0+0
	UT CS Antarctica Clus>	20	12+8	0+0
	CMS on CERN Linux	1	8+8	0+0
	CMS Production server	5	8+8	0+0
	UT DOUG Cluster	2	8+8	0+0
	CMS test cluster	1	8+8	0+0
	EENet cluster	6	2+8	0+0
	UT Physics Cluster	3	3+8	0+0
	CSC Kirppu	1	8+8	0+0
	Mill (Physicum)	64	8+6	0+0
	Alpha (HIP)	1	8+8	0+0
	Testbed0 (HIP)	1	8+8	4+1
	FZK cluster	886	246+322	0+0
	LRZ cluster	234	8+231	34+346
	Oslo Temp Cluster	13	8+8	5+0
	Parallab IBM Cluster	58	8+49	9+77
	Oslo Grid Cluster	43	16+8	16+0
	UiO Grid	105	51+53	1+50
	UPJS GRID	1	8+8	0+0
	SIGNET	51	43+8	14+0
	Bluesmoke (Swegrid,NS>	98	53+8	490+0
	Kosufy farm	60	58+8	81+0
	ISV	4	4+8	2+0
	Hagrid (SweGrid, Uppm>	100	95+8	84+0
	Ingrid (SweGrid,HPC2N)	101	96+8	165+0
	Monolith (NSC)	400	8+318 (queue down)	0+131
	Quark Cluster	7	3+8	3+0
	Beppe (SweGrid PDC KT>	96	92+2	230+0
	Sigrd (SweGrid, Luna>	99	49+58	7+35
Toto7/Whenim64 (Lunar>	191	8+156	0+27	
	Bern ATLAS Cluster	8	8+8	10+0
TOTAL	42 sites	4616	918 + 2610	1198 + 679

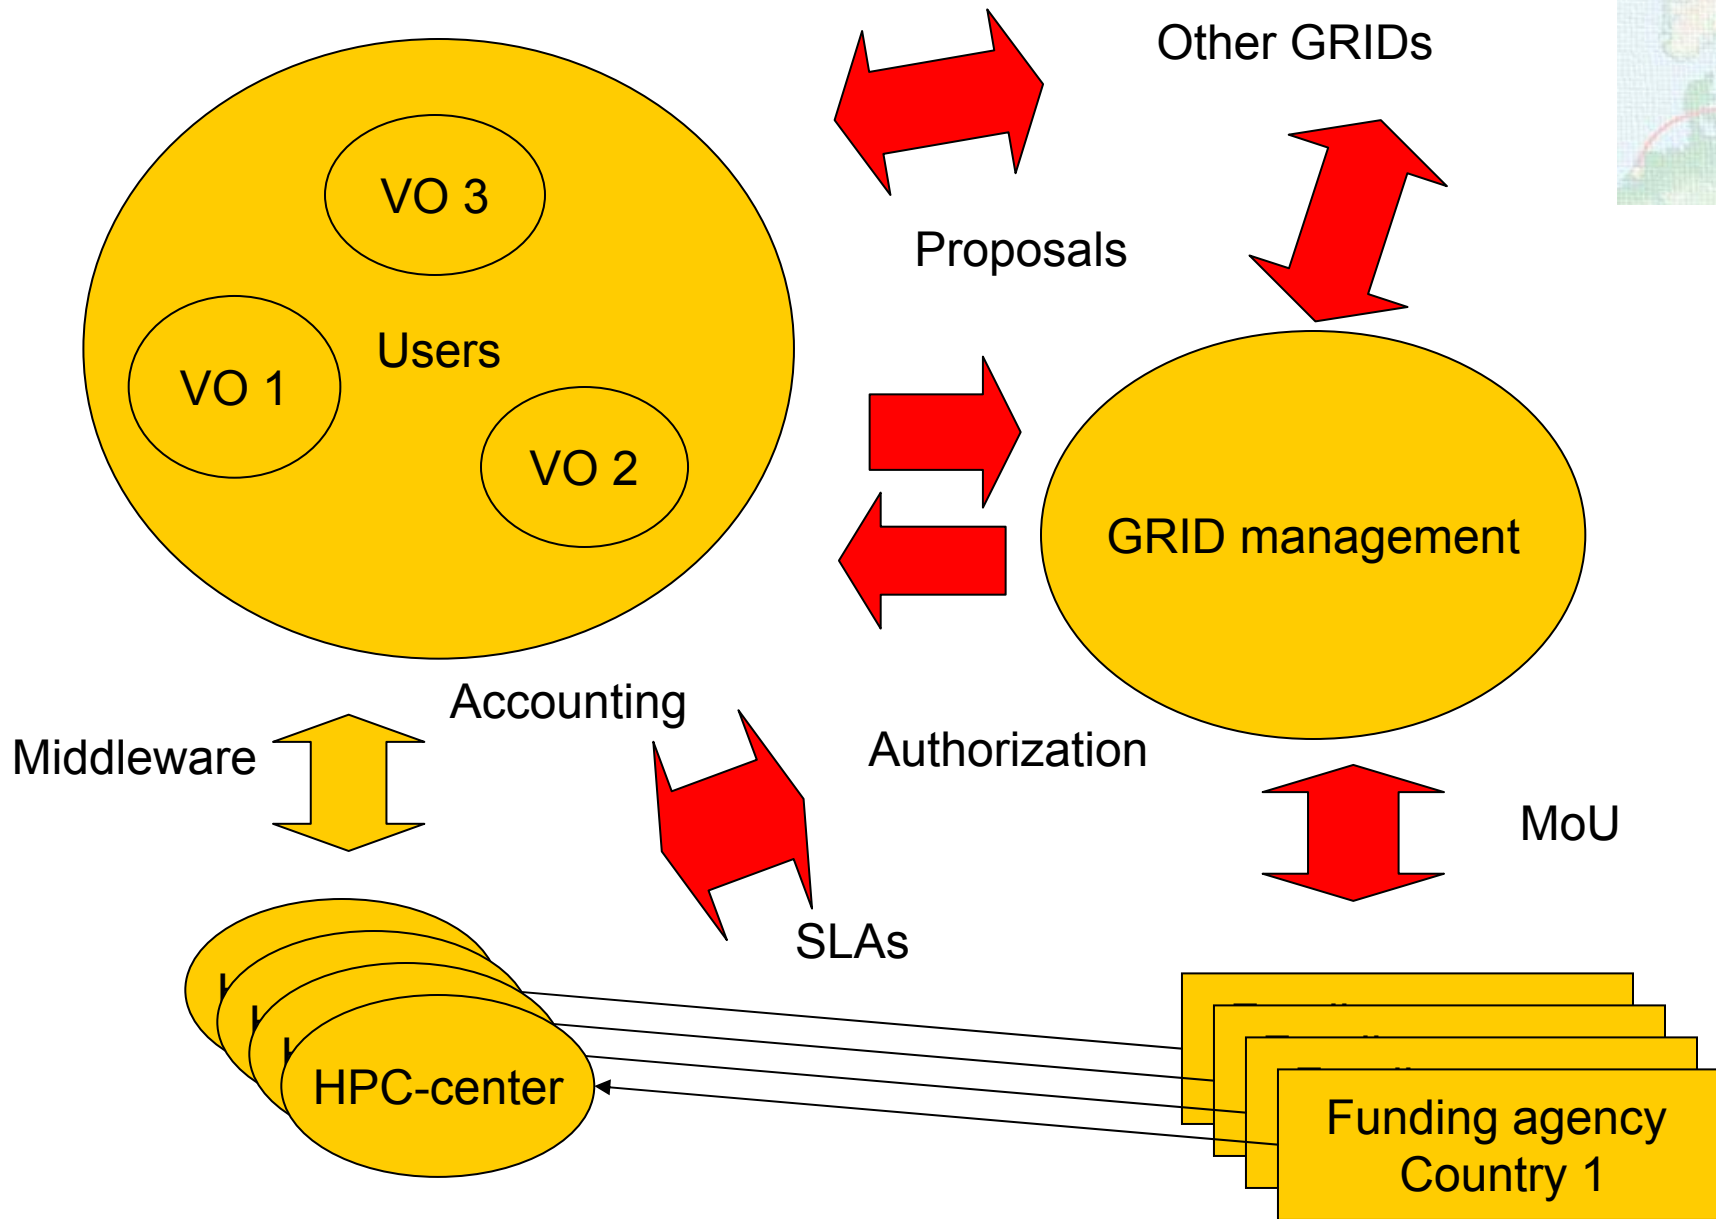
Reflections on NorduGrid



- Bottom up project driven by an application motivated group of talented people
- Middleware adaptation and development has followed a flexible and minimally invasive approach
- Nordic HPC centers are currently “connecting” large resources since it is good PR for the centers
- As soon as NorduGrid usage of these resources increases they will be disconnected. There is no such thing as free cycles!
- Motivation of resource allocations is missing – no Authorization
 - NorduGrid lacks an approved procedure for resource allocation to VOs and individual user groups based on scientific reviews of proposals

Challenges for Nordic Grids



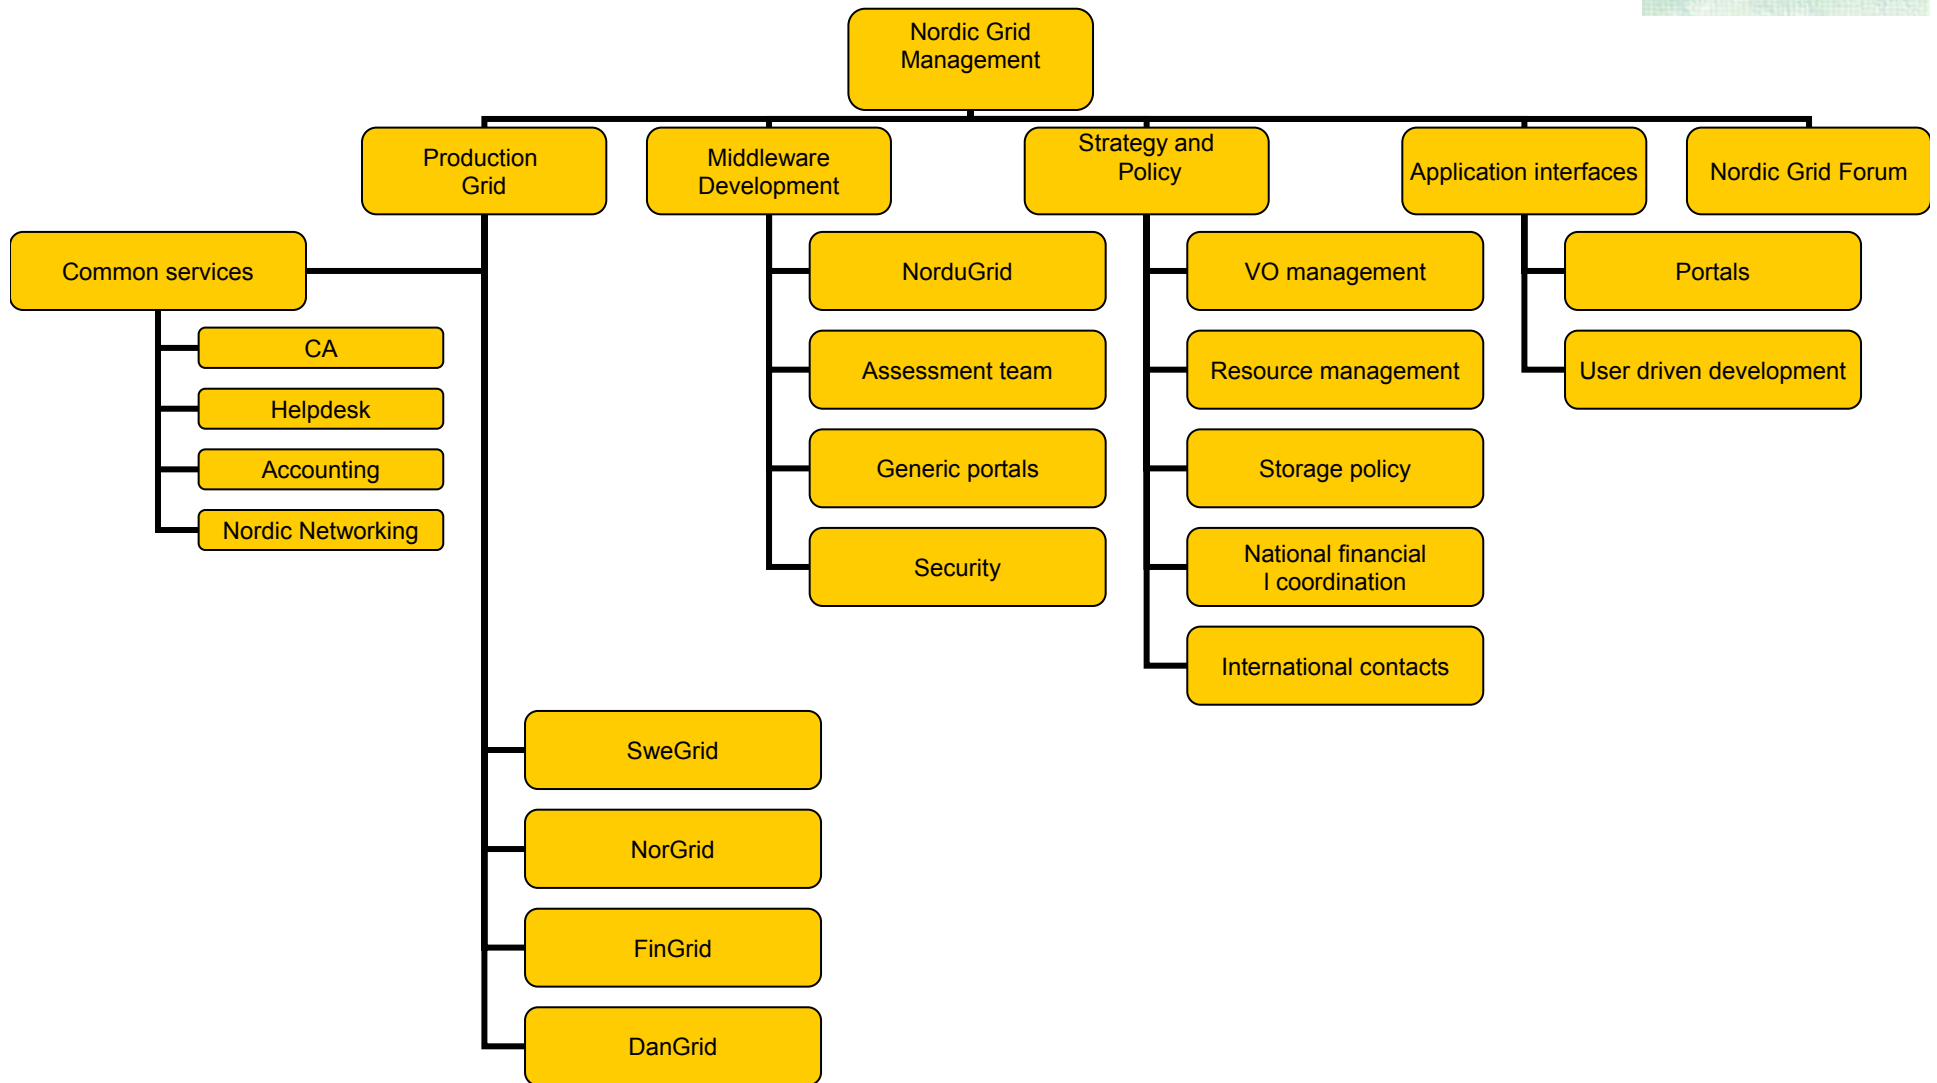


Path Forward - The Nordic Grid Facility



- Most likely:
 - NGF will be built as a Grid of Grids with a large centralized storage facility
 - Common services and projects will be identified
 - Will serve as a testbed for Grid Economic models
 - NorduGrid software (ARC) development CAN continue in the Framework of NGF
- In this way the Nordic countries will have one interface to the outside world to the benefit of:
 - LCG
 - EGEE
 - Could form the basis for a Nordic T1 centre
- Proposal will be put to NOS-N end of 2004
- Decision expected early 2005

Nordic Grid Facility- Organization



BEEGEE



- Baltic Extension of Enabling Grids for EsienceE
- An opportunity to make use of:
 - A well defined region
 - Several different Grid initiatives in the region
 - Merge or collaborate on similar projects
 - Middleware
 - National storage
 - National Authentication
 - Create a unified interface to other Grid projects
 - Try out economic models for Grid computing
 - Obtain significant EU-funding for Grid development and dissemination

