NDGF approach to serve the non-LHC communities needs: Material Science VO

Ivan Degtyarenko, D.Sc. (Tech)

NDGF

CSC – IT Center for Scientific Computing
Espoo, Finland

Because of LHC

the European Grid exists because of LHC, it has been designed in order to serve LHC needs

NDGF: BIO and CO2 CPU time consumption < 1%

EGEE: of 190 active VOs, 42 are HEP and consume over 90% of the CPU and storage resources

a lot of data and short time serial jobs

Source: NDGF 2009 Quaterly reports, EGEE accounting portal

Non-LHC communities

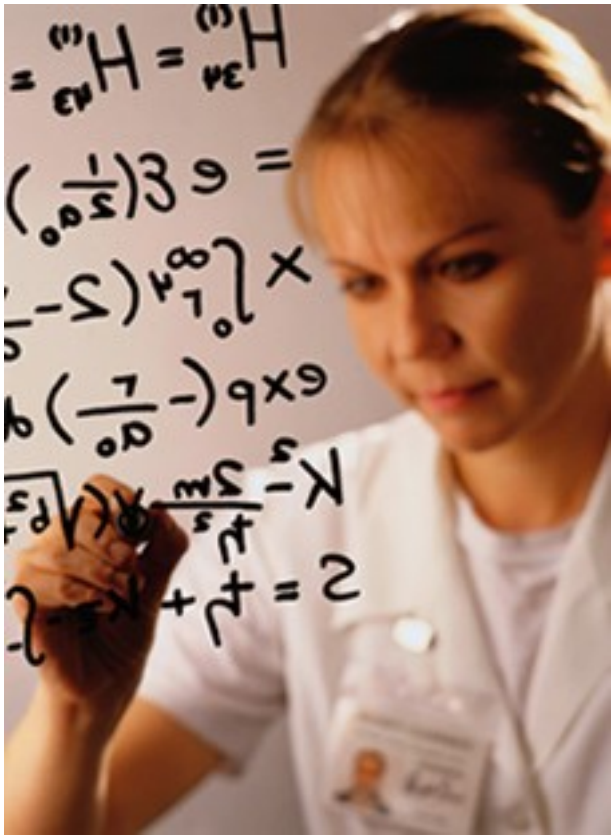

Who are them? What are they running? Where? What are their CPU and storage resource needs?

Why the vast majority of them do not care about grid interfaces?

Material Science VO

The joint project between NDGF and CSC on extending the Finnish Material Science VO to a Nordic countries level.

- support of the non-LHC communities
- not limited to any particular field of science
- become a strong actor at the European scale

Initial project goals

extend the VO resources, user experience and collaboration by attracting new computer centers and labs within Nordics

easily expandable / shrinkable system based on the solid statistical numbers and users preferences

- a number of popular applications support over all sites
- MPI support and impact on ARC development towards MPI environment support

Who?

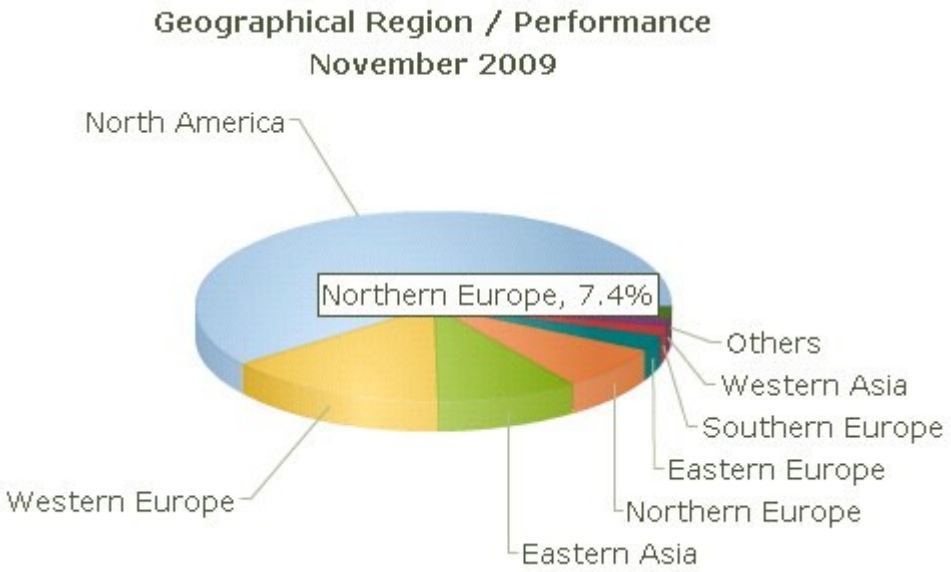

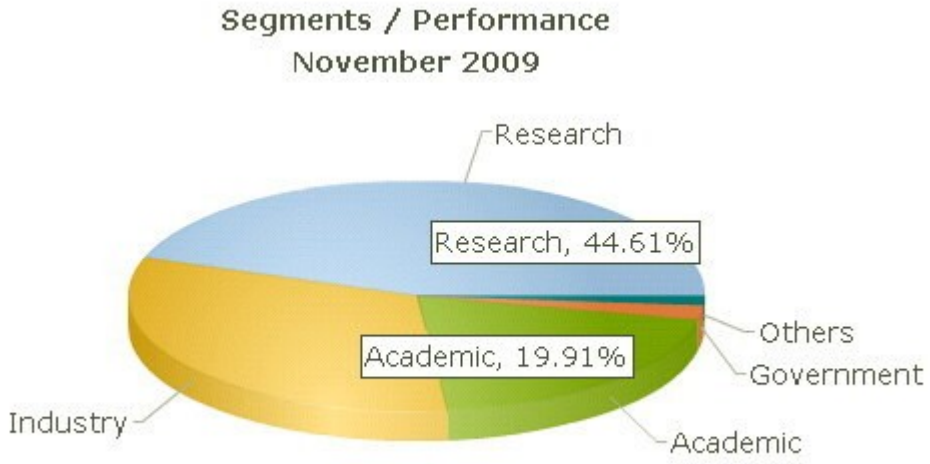

Who are them?

computational scientists at universities, polytechnics and research institutions

by fields of science: bio-sciences (including drug design and structural analyzes), physics and chemistry (and corresponding nanosciences), astrophysics, fluid dynamics, economics, medical, environmental, social and earth sciences, mathematics and computer science, language research ... others where scientific computing has influenced

What?

NUMBER of JOBS

total	6213569	
serial	4753250	(**76.50%**)
parallel	1460319	(23.50%)
lam	56640	(3.88% of pe jobs)
mpich	888456	(60.84%)
mpich2	51152	(3.50%)
openmpi	349598	(23.94%)
mvapich	79519	(5.45%)
threaded	31385	(2.15%)

WALLTIME (hours)

total	48920078	
serial	9535418	(19.49%)
parallel	39384660	(**80.51%**)
lam	2616030	(6.64% of pe jobs)
mpich	11557678	(29.35%)
mpich2	889372	(2.26%)
openmpi	11481871	(29.15%)
mvapich	12234506	(31.06%)
threaded	100057	(0.25%)

Statistics: Finnish M-grid clusters

Use of CPU cores in parallel

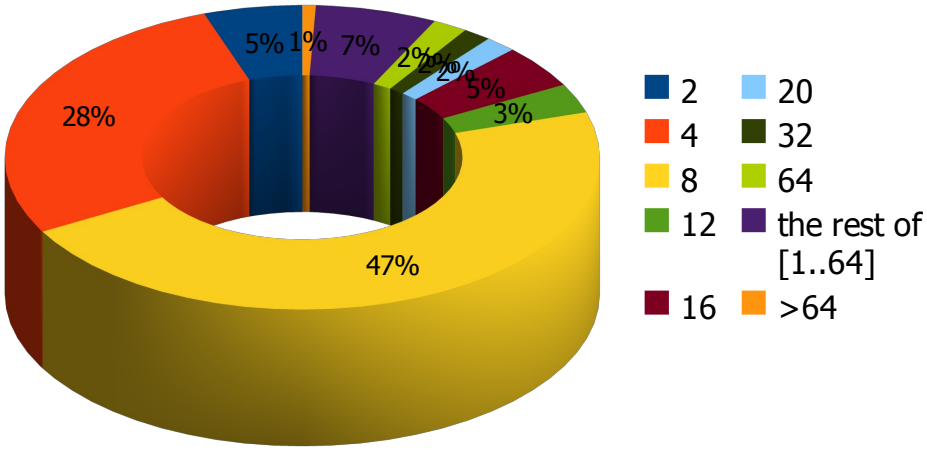

Walltime distribution for parallel jobs

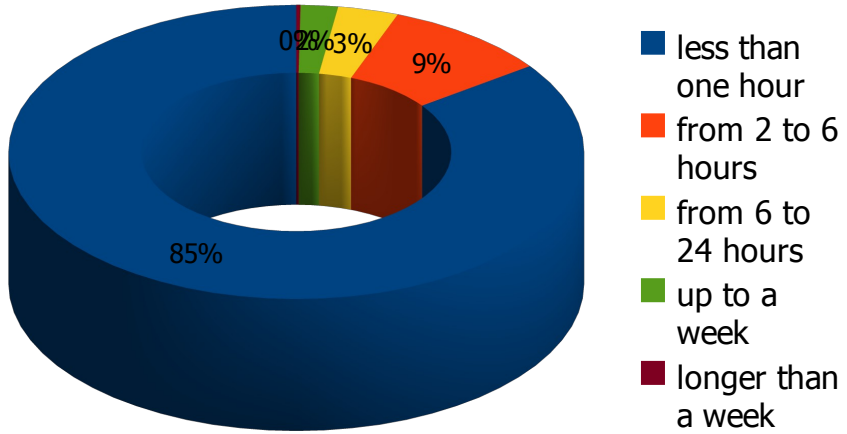

Walltime distribution for serial jobs

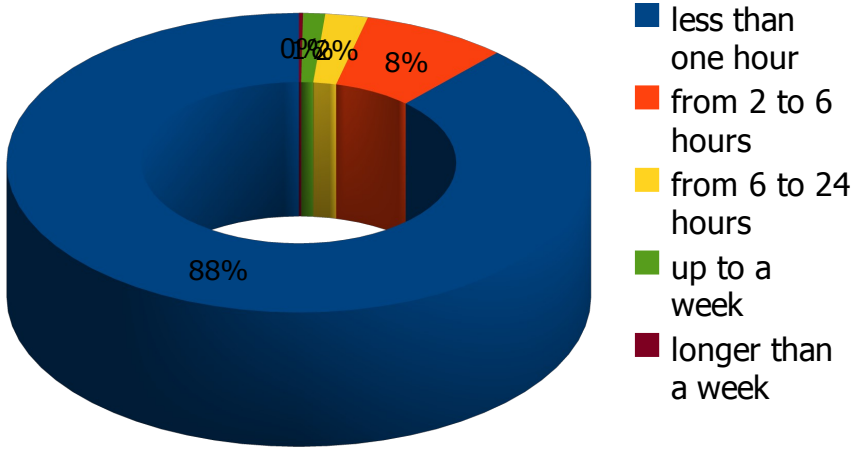

Statistics: cluster at CSC

Number of jobs distribution

total jobs number run on 1, 2, 4 ... 512 CPU cores

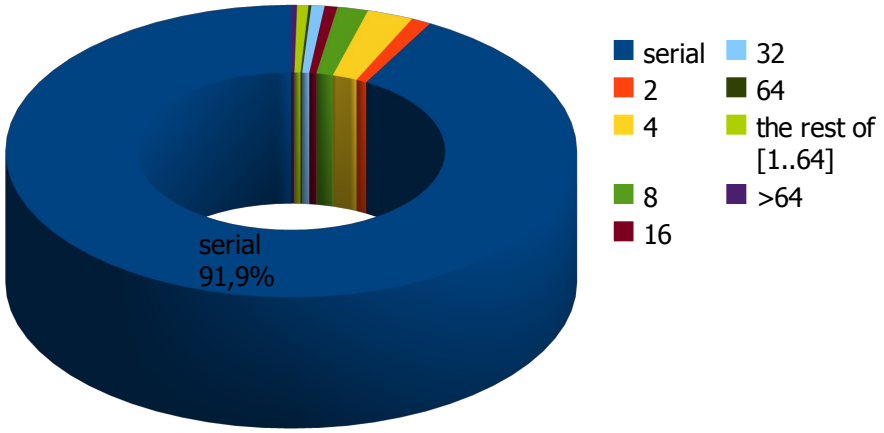

CPU time consumtion distribution

total CPU time consumed by 1, 2, 4 ... 512 CPU core jobs

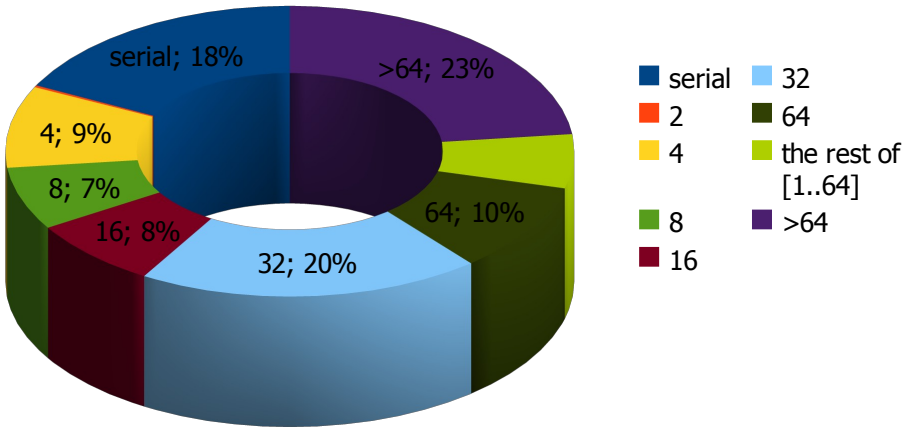

parallel jobs walltime distribution

job length statistics

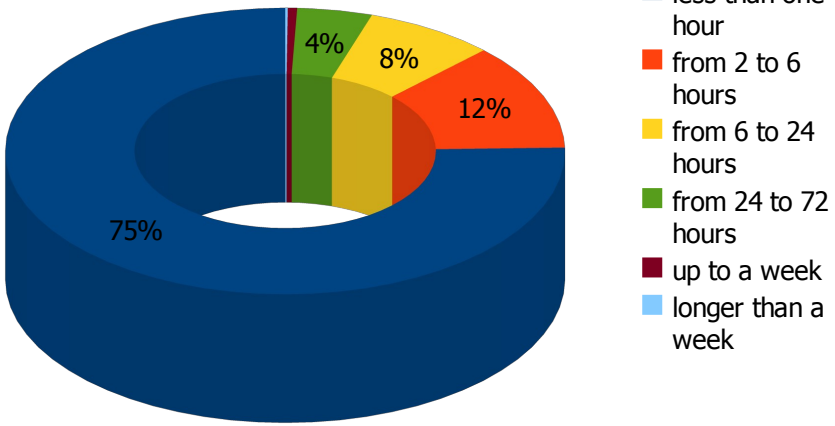

serial jobs walltime distribution

job length statistics

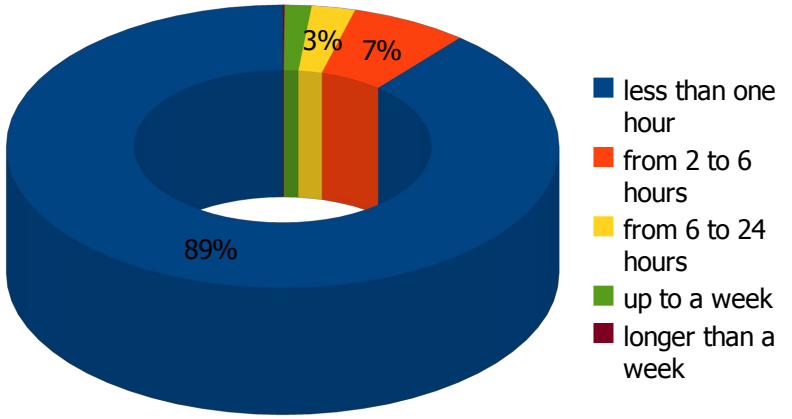

Runtime needs summary

majority of the jobs are:

- no longer than 24 hours
- serial (in terms of job number)
- parallel (in terms of CPU time)

Type of applications

a.out (70%, another 30% are pre-installed apps)

Top list of pre-installed apps (Finland): GPAW, NAMD, Gromacs, TURBOMOLE, Gaussian, CPMD, ADF, BLAST, Elmer, Molpro, Dmol, NWChem

(Norway): WRF, VASP, ROMS, CCSM, EnKF, Dalton, Gaussian, ADF

Others?

Licensing: an obstacle or opportunity?

MPI must be taken seriously

ability to run and compile MPI easily

OpenMPI: the default recommended flavor

ability to request the varying number of slots

ability to request logical CPUs within

one physical CPU only, or one WN

available memory per logical CPU

interconnecting choice

Others (?)

Canonical use case

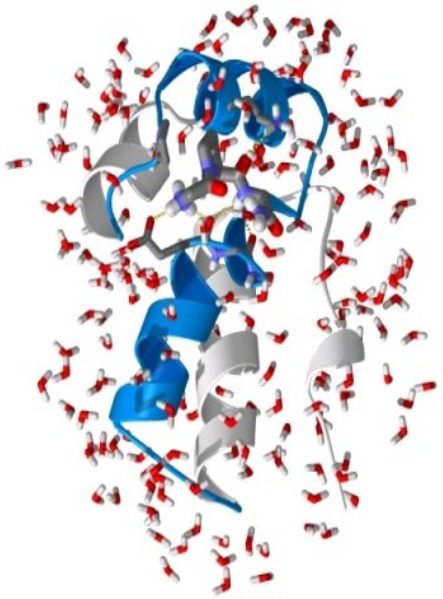

Grid type of job

compile/install app (optional), prepare the application input files, submit the job, fetch the results, analyze&visualize

canonical use case: Gromacs

description: classical molecular dynamics package, widely used by biophysicists and other bio-scientists

input: a few files, ~megabytes

output: no limits, typically a few gigabytes

execution: may take a few hours or weeks, in serial, parallel (typically), as an array jobs

analyzes&visualization: VMD, PyMol, gOpenMol

General observations

batch system users are the target auditory

make sure they know about grid interface opportunity

the community of computational scientists must be a driven force, look for the critical mass

the concrete user, and not a user

HPC and not HTC