




IT4Innovations

Centrum excellence

Vít Vondrák



EVROPSKÁ UNIE
EVROPSKÝ FOND PRO REGIONÁLNÍ ROZVOJ
INVESTICE DO VAŠÍ BUDOUCNOSTI



Introduction

Project objective: to establish a centre focused on research in the field of IT with emphasis on development of HPC

3 basic project levels:

1. Resources management – creating a functioning research centre
2. Research programmes implementation
3. Acquisition of research infrastructure – particularly supercomputing facilities



Introduction II

Implementation period for the grant-covered part of the project:

July 2011 to December 2015

EU+MEYS grant: over 75 million EUR

Principal stages:

7/11 – start of implementation


9/11 – start of research programmes

III.Q/12 – acquisition of “small cluster” – container solution

I.Q/13 – “small cluster” ready for utilization by external academic institutions



2014 – installation of “big cluster” in the new building

2015 – full operation of IT4I centre



Current HPC activities

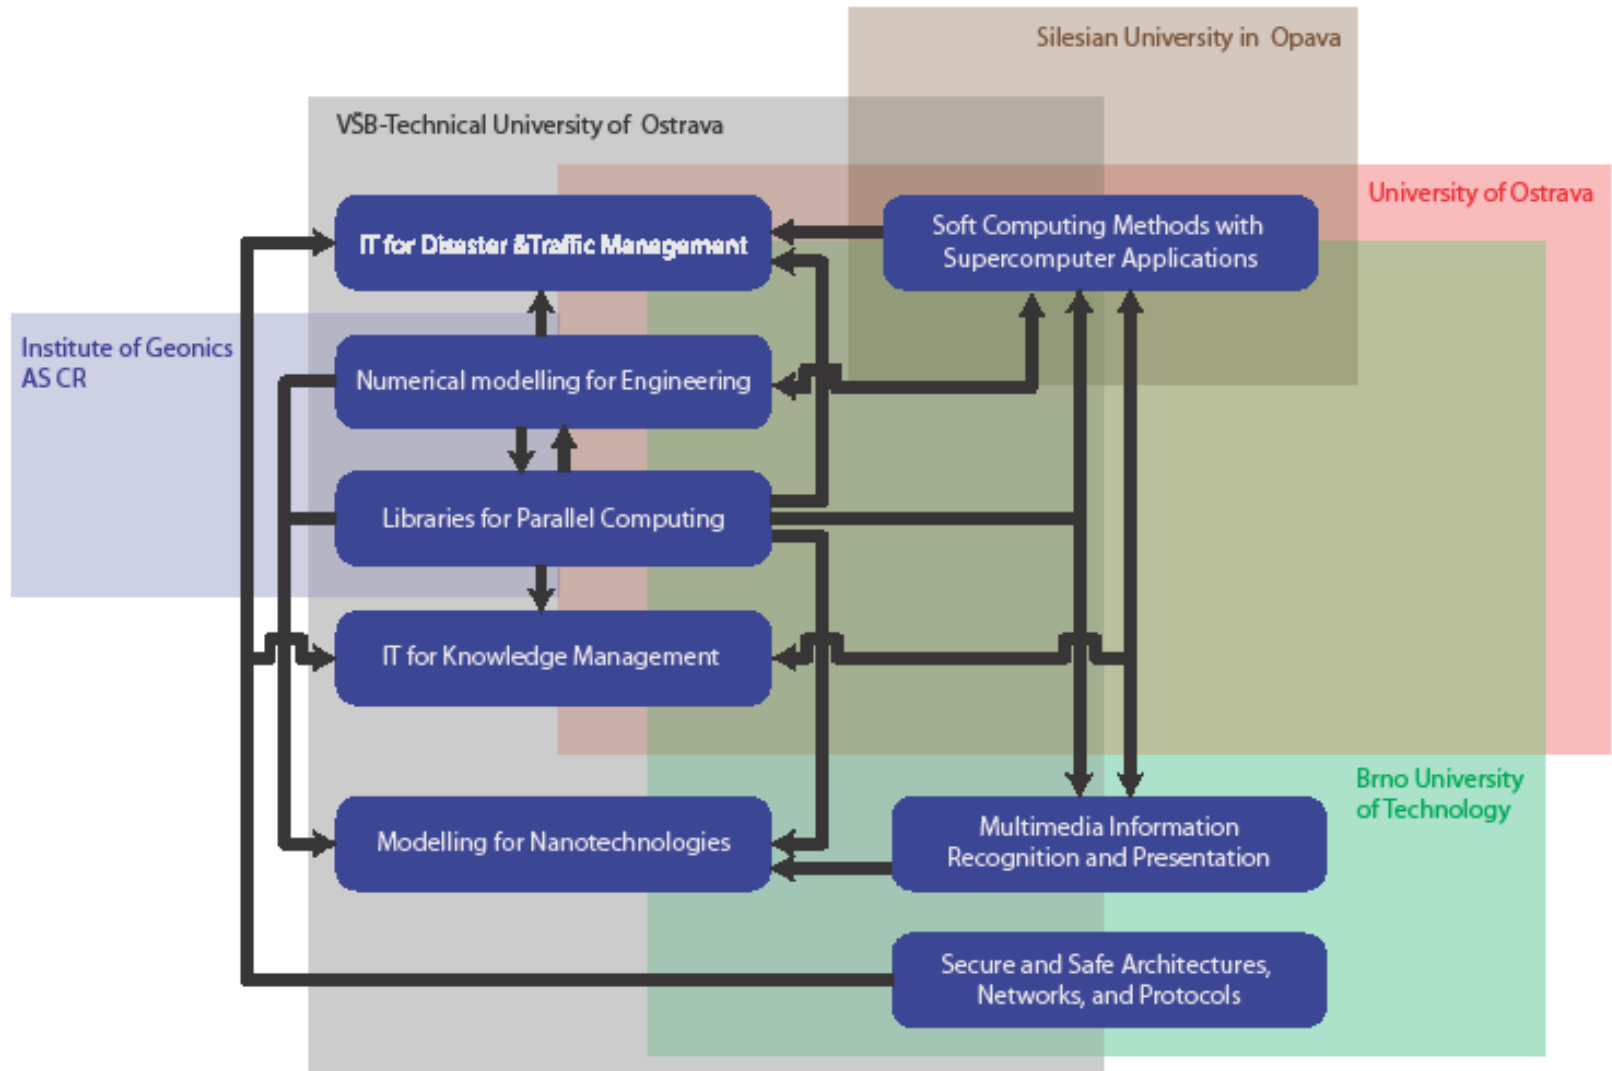
- Disaster management
 - Floods and pollution modelling
 - Traffic simulations
- Computational mathematics and mechanics
 - Large scale nonlinear FE models
 - Multiscale and multiphysic problems
 - Optimal design problems
 - Uncertainty modelling
- Computational physics and chemistry
 - Development of new pharmaceutical forms
 - Nanocomposites and nanothreads
- Computational problems in informatics
 - Information retrieval
 - Multimedia information recognition and presentation
 - Knowledge base





IT4Innovations
Centrum excellence

Research programmes



Research Areas and Programmes

Information Technologies for People - IT4People

- RP1 IT for Disaster and Traffic Management – VSB - TUO
- RP7 Multimedia Information Recognition and Presentation – BUT

Supercomputing for Industry - SC4Industry

- RP2 Numerical Modelling for Engineering – VSB-TUO and IG AS
- RP3 Libraries for Parallel Computing – VSB - TUO
- RP4 Modelling for Nanotechnologies – VSB-TUO

Theory for Information Technologies - Theory4IT

- RP5 IT for Knowledge Management – VSB - TUO
- RP6 Soft Computing Methods with Supercomputer Applications – UO
- RP8 Secure and Safe Architectures, Networks, and Protocols – BUT

All the partners - VSB-TUO (Technical University of Ostrava), BUT (Brno University of Technology), OU (University of Ostrava), SU (Silesian University) and IG AS (Institute of Geonics) – cooperates within these programmes.



panasas Data + Performance + Private Cloud = Panasas Intersect360 RESEARCH DOWNLOAD REPORT

- PROJECT
- LISTS
- STATISTICS
- RESOURCES
- NEWS
- CONTACT
- SUBMISSIONS
- LINKS
- HOME

Home > Lists > November 2011

TOP500 List - November 2011 (1-100)

R_{max} and R_{peak} values are in TFlops. For more details about other fields, check the [TOP500 description](#).

Power data in KW for entire system

next

Rank	Site	Computer/Year Vendor	Cores	R _{max}	R _{peak}	Power
1	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIx 2.0GHz, Tofu interconnect / 2011 Fujitsu	705024	10510.00	11280.38	12659.9
2	National Supercomputing Center in Tianjin China	NUDT YH MPP, Xeon X5670 6C 2.93 GHz, NVIDIA 2050 / 2010 NUDT	186368	2566.00	4701.00	4040.0
3	DOE/SC/Oak Ridge National Laboratory United States	Cray XT5-HE Opteron 6-core 2.6 GHz / 2009 Cray Inc.	224162	1759.00	2331.00	6950.0
4	National Supercomputing Centre in Shenzhen (NSCS) China	Dawning TC3600 Blade System, Xeon X5650 6C 2.66GHz, Infiniband QDR, NVIDIA 2050 / 2010 Dawning	120640	1271.00	2984.30	2580.0
5	GSIC Center, Tokyo Institute of Technology Japan	HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows / 2010 NEC/HP	73278	1192.00	2287.63	1398.6

SUPERMICR GPU Blade
20 GPUs in 7U
intel inside Xeon

FUJITSU PRIMEHPC FX10

To se mi líbí f Toto se líbí 425 lidem.

ISC INTERNATIONAL SUPERCOMPUTING CONFERENCE
June 17 - 21, 2012 Hamburg, Germany

APPRO
Based on the future Intel® Xeon® Processor E5 Family

Want to



IT4Innovations

Centrum excellence

130	Naval Oceanographic Office - NAVO DSRC United States	Cray XT5 QC 2.4 GHz / 2011 Cray Inc.	12720	96.55	122.11	
131	EDF R&D France	Blue Gene/P Solution / 2008 IBM	32768	95.45	111.41	252.0
132	IT Service Provider Germany	Cluster Platform 3000 BL2x220, E54xx 3.0 Ghz, Infiniband / 2009 HP	10240	94.74	122.88	
133	Gaming Company United States	Cluster Platform 3000 BL460c G7, Xeon X5660 2.80 GHz, 10G Ethernet / 2011 HP	11712	94.31	131.17	
134	Geoscience United States	iDataPlex DX360M3, Xeon X5650 6C 2.66 GHz, Infiniband, NVIDIA 2090 / 2011 IBM	7392	93.65	244.89	108.2
135	Geoscience United States	iDataPlex DX360M3, Xeon X5650 6C 2.66 GHz, Infiniband, NVIDIA 2090 / 2011 IBM	7392	93.65	244.89	108.2



Research infrastructure: HW 2015

“big cluster“

- arch. x86, min. 32768 cores, 4GB RAM / core, IB QDR (fat-tree topology)

SMP/NUMA

- arch. x86, 512 cores, min. 2TB RAM, dedicated scratch storage

storage capacity

- ca. 2.4 PB storage (5/20/75 % cap. SAS/SATA/tape – HSM)

visualization centre

- arch. x86, ca. 10 graph. stations with GPU, stereoscopic projection, pre- and post- processing of tasks

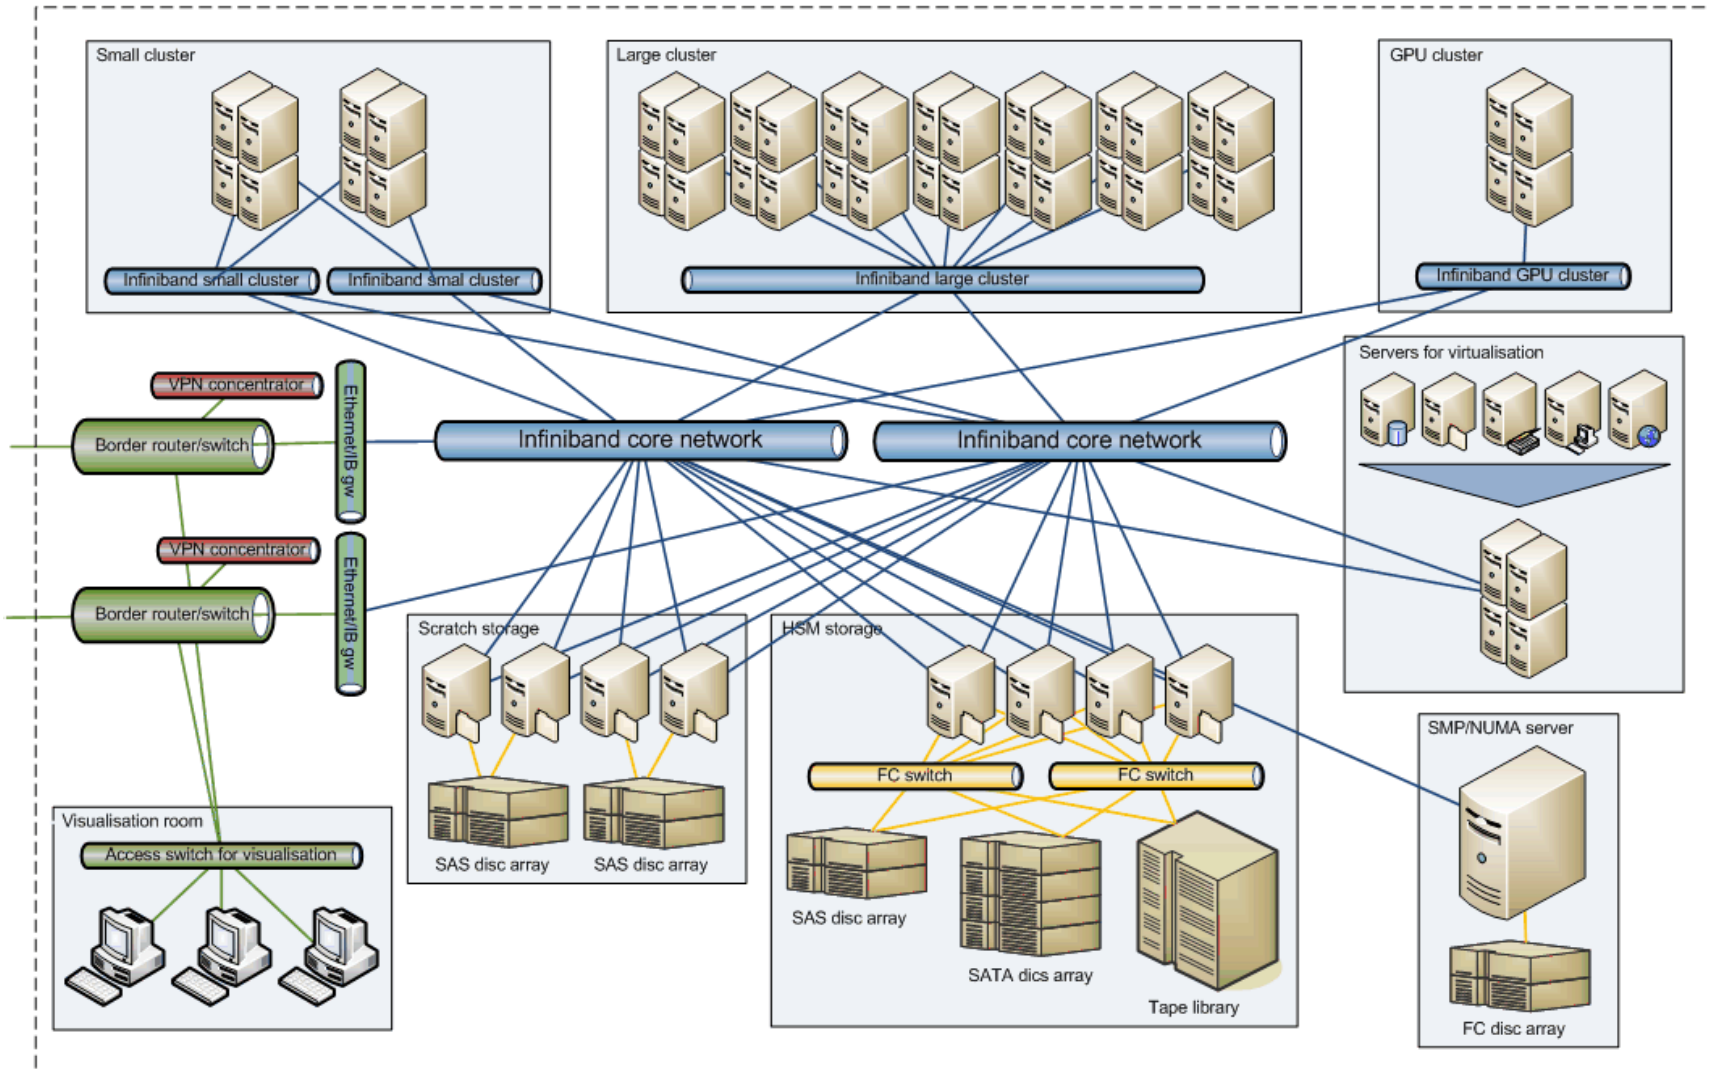


Data centre parameters

- 450 sq. m.
- Total power input 2,7MW
 - 4x DRUPS – 950kVA/760kW
 - N+1 redundancy
- PUE 1.34-1.5 (100-60%)
 - usage of “freecooling” (294 kW)



Final configuration of data centre

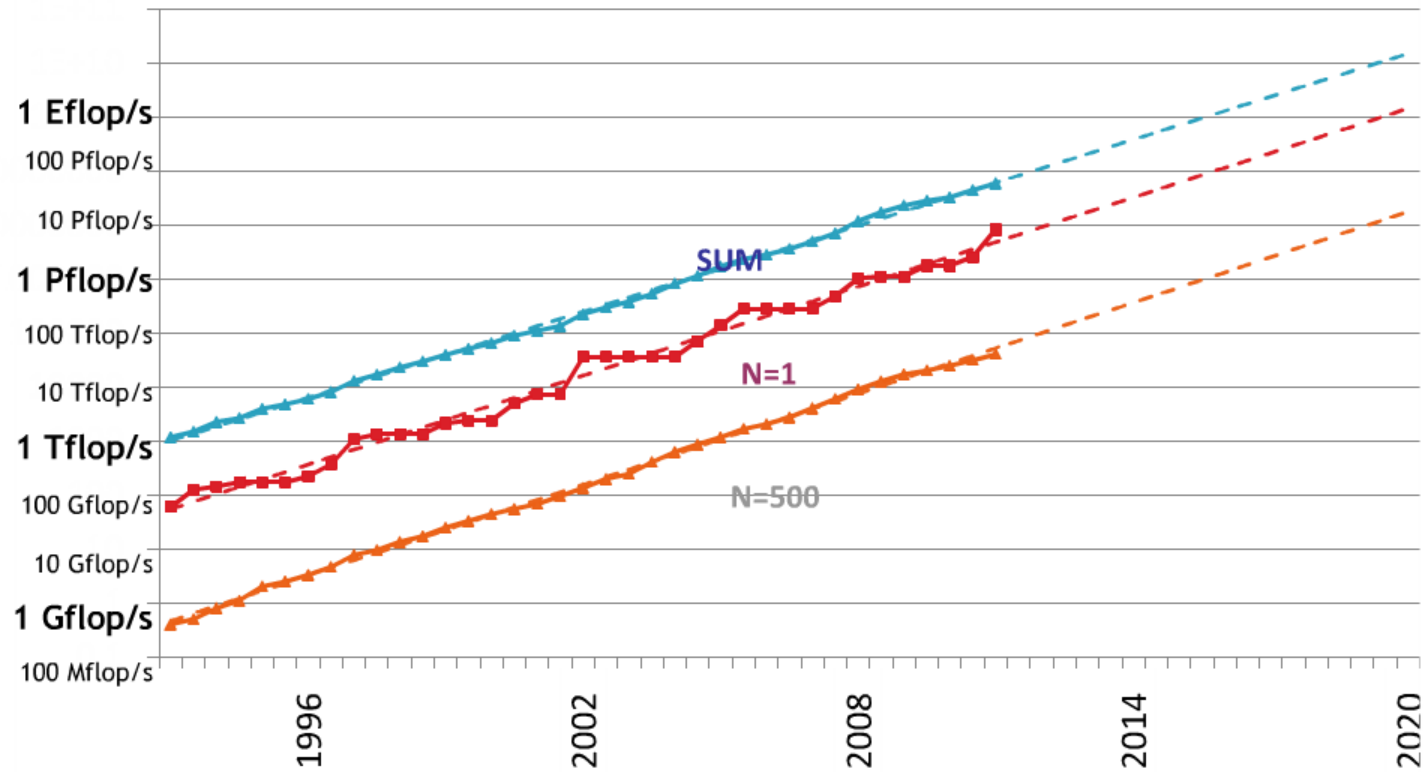




8	DOE/SC/LBNL/NERSC United States	Cray XE6, Opteron 6172 12C 2.10GHz, Custom / 2010 Cray Inc.	153408	1054.00	1288.63	2910.0
9	Commissariat a l'Energie Atomique (CEA) France	Bull bullx super-node S6010/S6030 / 2010 Bull	138368	1050.00	1254.55	4590.0
10	DOE/NNSA/LANL United States	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband / 2009 IBM	122400	1042.00	1375.78	2345.0
11	National Institute for Computational Sciences/University of Tennessee United States	Cray XT5-HE Opteron Six Core 2.6 GHz / 2011 Cray Inc.	112800	919.10	1173.00	3090.0
12	HWW/Universitaet Stuttgart Germany	Cray XE6, Opteron 6276 16C 2.30 GHz, Cray Gemini interconnect / 2011 Cray Inc.	113472	831.40	1043.94	
13	Forschungszentrum Juelich (FZJ) Germany	Blue Gene/P Solution / 2009 IBM	294912	825.50	1002.70	2268.0
14	National Supercomputing Center in Jinan China	Sunway BlueLight MPP, ShenWei processor SW1600 975.00 MHz, Infiniband QDR / 2011 NRCP CET	137200	795.90	1070.16	1074.0



Projected Performance Development



Software equipment

- **Development software**
 - Compilers C/C++, Fortran, Java (GNU, Intel, PGI)
 - Libraries LINPACK, LAPACK, MPI, PETSc
 - Totalview, Allinea DDT
 - Matlab & Simulink
 - Mathematica
- **CAD/CAM/CAE software**
 - ANSYS, Fluent, LS-DYNA
 - MSc. Software – Nastran, Patran, Adams, Marc
 - Comsol Multiphysics
- **Geology, hydrology, geoinformatics**
 - GRASS, ARC-GIS
 - DHI MIKE, HEC-HMS, HEC-RAS
 - MODFLOW, FEFLOW
- **Chemistry, physics**
 - Materials Studio
 - Gaussian, MolPro
- **Software for SCC**
 - Backup software
 - Monitoring software
 - Administration software
 - Schedulers

Resources management

= Business model

50% internal access for own research programmes

30% open access for external institutions

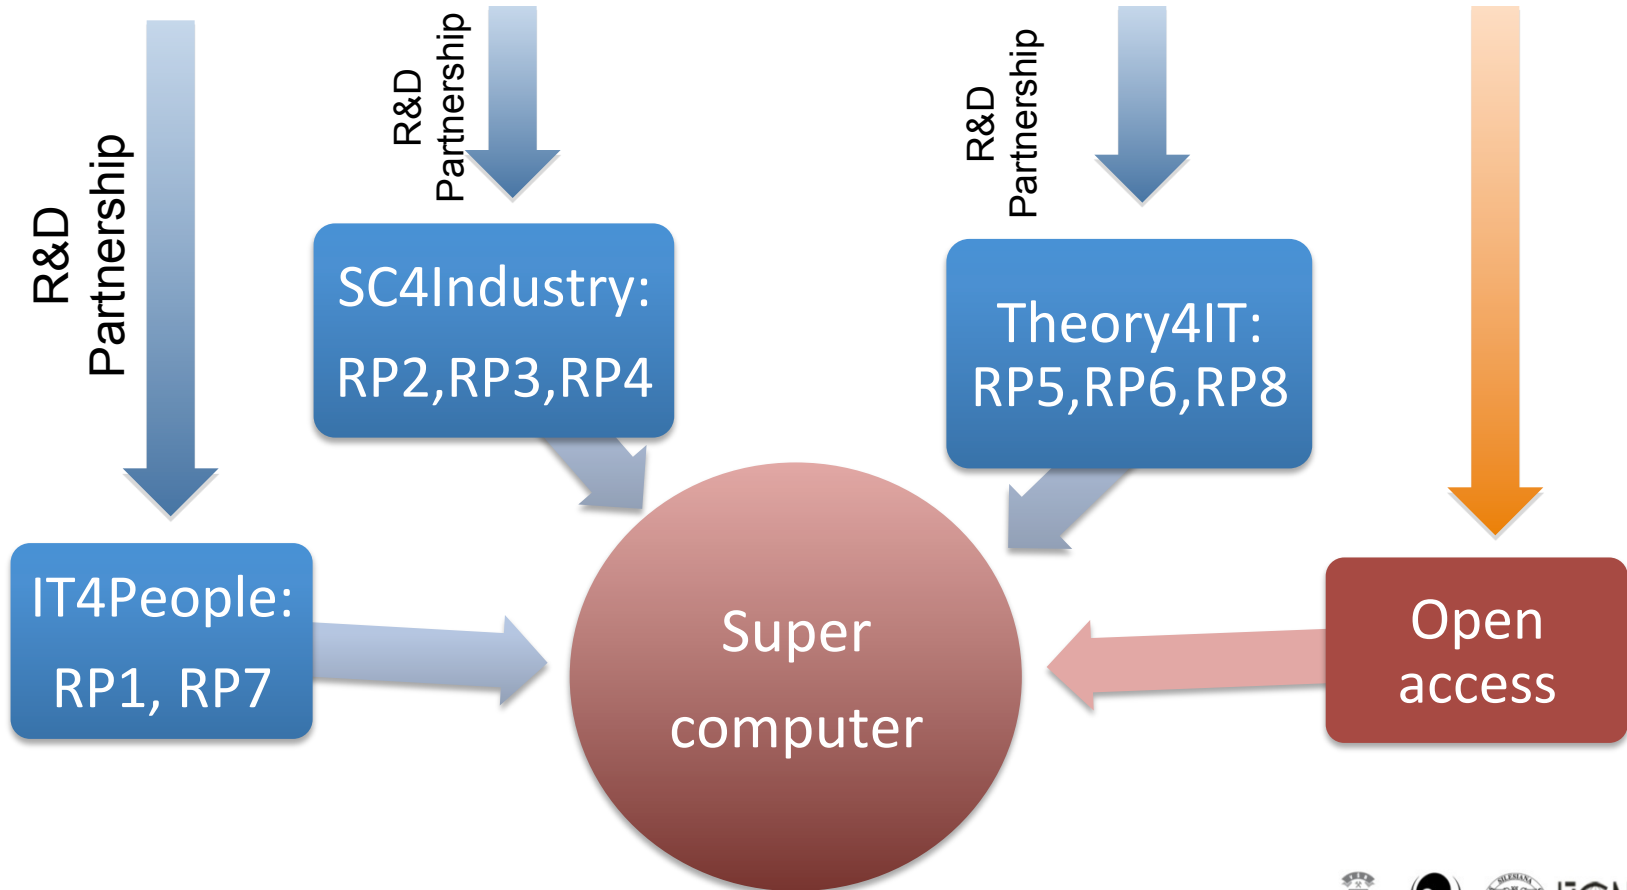
- access based on “grant” competition

20% dedicated access for projects of national importance

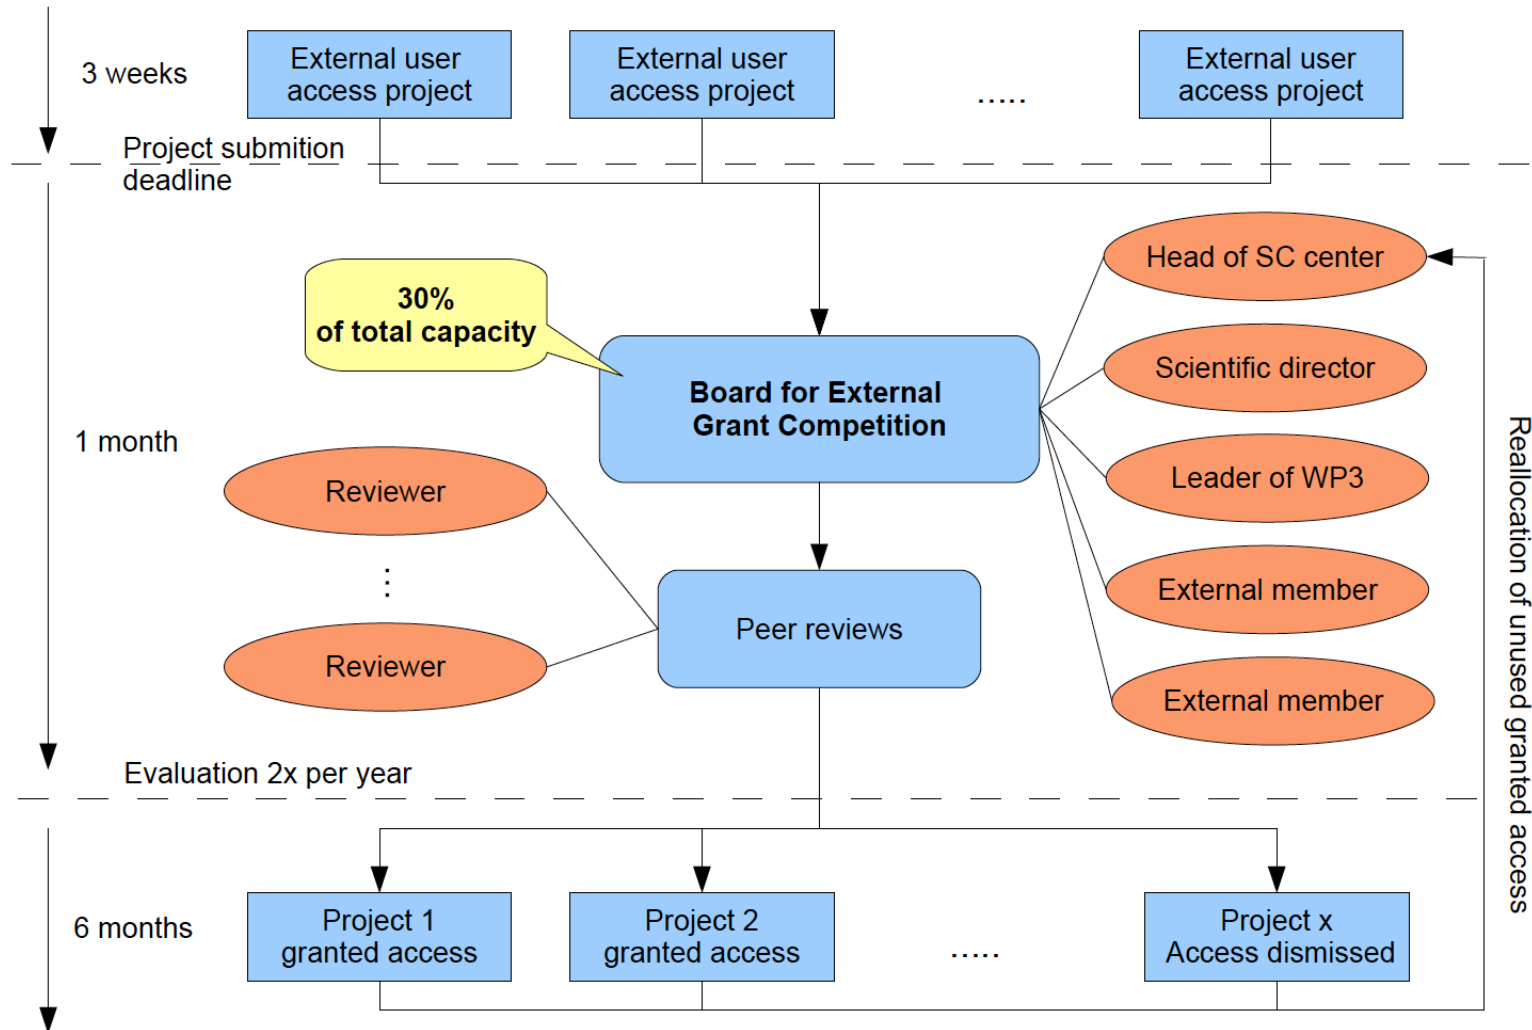




External organizations



External users access scheme



External open access

- Types of allocations
 - Small (up to 5.000 bu per month + 10GB disk storage)
 - Testing purposes, evaluation 4x per year, no peer review
 - Medium (up to 50.000 bu per month + 50GB disk storage)
 - Advanced projects for parallel computing, 4x per year, no peer review, possible sharing of operational cost
 - Large (up to 500.000 bu per month + individual disk storage requirements)
 - Large projects for parallel computing, 2x per year, peer review, possible sharing of operational cost
 - Grand Challenge (more than 500.000 bu per month)
 - Possible use of PRACE resources)
 - 1 bu = 1 Billing Unit = 1 core hour



**PARTNERSHIP
FOR ADVANCED COMPUTING
IN EUROPE**

PRACE – Building a new e-infrastructure in Europe






PRACE

The Partnership for Advance Computing in Europe is *the* European HPC Research Infrastructure

- PRACE enables world-class science through large scale simulations
- PRACE provides HPC services on leading edge capability systems on a diverse set of architectures
- PRACE operates up to six Tier-0 systems as a single entity including user and application support
- PRACE offers its resources through a single pan-European peer review process
- PRACE is providing services since August 2010
- The first Tier-0 system is the fastest supercomputer in Europe

The European HPC Ecosystem

- PRACE prepares the creation of a persistent pan-European HPC service, consisting of several tier-0 centres providing European researchers with access to capability computers and forming the top level of the European HPC ecosystem.
- Since 2010 – Czech Republic member of PRACE
- IT4Innovations – Tier-1 system in European HPC Ecosystem





Current status

- 21 European countries are currently part of PRACE
- 4 hosting partners
- The PRACE first implementation phase started on July 1. 2010
- The PRACE second implementation phase started on September 1. 2011





PRACE RI – Tier-0 Systems

- The first production system, a 1 Petaflop/s IBM BlueGene/P (**JUGENE**) at GCS (Gauss Centre for Supercomputing) partner FZJ (Forschungszentrum Jülich) is available for European scientists.
- The second production system called **CURIE** at CEA-GENCI by Bull has been available for European scientists since January 2011. It will reach its full capacity of 1.6 Petaflop/s by late 2011.
- The third production system, a 1 Petaflop/s Cray (**HERMIT**) at GCS partner HLRS (High Performance Computing Center Stuttgart) will be available for European scientists by the end of 2011 with an upgrade to 4-5 Petaflop/s in 2013.
- The fourth production system, a 3 Petaflop/s IBM (**SuperMUC**) at GCS partner LRZ (Leibniz-Rechenzentrum) will be available for European scientists starting in mid 2012.
- Italy has announced the deployment of its Tier-0 systems for 2012, and Spain will follow in 2013.

FZJ

2010 1st PRACE System - JUGENE

- BG/P by Gauss Center for Supercomputing at Juelich
 - 294,912 CPU cores, 144 TB memory
 - 1 PFlop/s peak performance
 - 825.5 TFlop/s Linpack
 - 600 I/O nodes (10GigE) > 60 GB/s I/O
 - 2.2 MW power consumption
 - 35% for PRACE



GENCI

2011 2nd PRACE system – CURIE

- Bull, 1.6PF, 92160 cores, 4GB/core
- Phase 1, December 2010, 105 TF
 - 360 four Intel Nehalem-EX 8-core nodes, 2.26 GHz CPUs (11,520 cores), QDR Infiniband fat-tree
 - 800 TB, >30GB/sec, local Lustre file system
- Phase 1.5 Q2 2011
 - Conversion to 90 16-socket, 128 core, 512 GB nodes
- Phase 2, Q4 2011, 1.5 TF
 - Intel Sandy-Bridge
 - 10PB, 230GB/sec file system



HLRS

2011 3rd PRACE System – HERMIT

- Cray XE6 (Multi-year contract for \$60+M)
 - Phase 0 – 2010
10TF, 84 dual socket 8-core
AMD Magny-Cours CPUs,
1344 cores in total, 2 GHz,
2GB/core,
Gemini interconnect
 - Phase 1 Step 1 – Q3 2011
AMD Interlagos, 16 cores, 1 PF
2 – 4 GB/core
2.7 PB file system, 150 GB/s I/O
 - Phase 2 – 2013
Cascade, first order for Cray, 4- 5 PF



LRZ

2011/12 4th PRACE system


- IBM iDataPlex (€83M including operational costs)
 - >14,000 Intel Sandy-Bridge CPUs, 3 PF (~110,000 cores), 384 TB of memory
 - 10PB GPFS file system with 200GB/sec I/O, 2PB 10GB/sec NAS
 - LRZ <13MW
 - Innovative hot water cooling (60C inlet, 65C outlet) leading to 40 percent less energy consumption compared to air-cooled machine.



BSC and CINECA

- 2012/2013 5th and 6th PRACE Systems

CINECA
Target ~2.5 PF



Computing Facility
10 MW 2013



Accessing the PRACE RI

Access Model

- Proposals are evaluated in a single European peer review process governed by the PRACE Scientific Steering Committee.
- Three types of resource allocations
 - Test / evaluation Preparatory access
 - Project access – for a specific project, grant period ~ 1 year
 - Programme access – resources managed by a community
- Free-of-charge

Funding

- Mainly national funding through partner countries
- European contribution
- Access model has to respect national interests (ROI)

Project access calls for proposals

- Access to PRACE resources is open to researchers from European academic institutions and industry
- The first combined call for Tier-0 and Tier-1 resources was open
 - The main target of the PRACE-2IP project is to merge Tier-1 services into the PRACE Research Infrastructure; these were previously provided by the DEISA projects (Distributed European Infrastructure for Supercomputing Applications).
- Find out more at

<http://www.prace-ri.eu/hpc-access>

Past project access calls for proposals

- **PRACE Early Access call**
May 10 – June 10, 2010
 - **10 proposals were granted a 332 million core hours**
- **PRACE Project Access - 1st call for proposals**
June 15 – August 15, 2010
 - **9 proposals requesting 362 million core hours were granted access**
- **PRACE Project Access – 2nd Call for Proposals**
November 1, 2010 – January 11, 2011
 - **17 proposals requesting 397.8 million core hours were granted access**
- **PRACE Project Access – 3rd Call for Proposals**
May 2, 2011 – November 1, 2011
- **PRACE Project Access – 4th Call for Proposals**
November 2, 2011 – May 1, 2012

Preparatory access calls for proposals

- Continuously open with cut-off every 3 months
- 3 types of preparatory access
 - Type A – **Code scalability testing** (max 100.000 core hours JUGENE, 50.000 core hours CURIE)
 - Type B - **Code development and optimization by the applicant** (max 250.000 core hours JUGENE, 200.000 core hours CURIE, max 2 months, max 6 months)
 - Type C - **Code development with support from experts** (max 250.000 core hours JUGENE, 200.000 core hours CURIE, max 6 months)

		Type A			Type B			Type C		
	Description	4 th cut-off	5 th cut-off	6 th cut-off	4 th cut-off	5 th cut-off	6 th cut-off	4 th cut-off	5 th cut-off	6 th cut-off
	Requests	14(3)	8	7	4	5	1	8	0	2
Awarded	JUGENE	5(2)	3(1)	1(1)	3(1)	0	1(1)	3(2)	0	0
	CURIE	7(2)	6(1)	6(1)	2(1)	4	1(1)	4(2)	0	2
	HERMIT	NA	NA		NA	NA	0	NA	NA	
KHours	JUGENE	500	300	100	600	0	250	600	0	0
	CURIE	350	300	300	500	800	200	1000	0	400
	HERMIT	NA	NA		NA	NA		NA	NA	

Thank you for your attention!

www.it4i.cz

www.it4i.eu

www.it4innovations.cz

www.it4innovations.eu

facebook.com/it4innovations

Kick-off meeting

November 29th, 2011

