# Integrated HPC Infrastructure for the Research Communities in South-East Europe

Assoc. Prof. Emanouil Atanassov
Head of Grid Technologies and Applications Department
Institute of Information and Communication Technologies – BAS
Email: emanouil@parallel.bas.bg

# Regional Cooperation - History

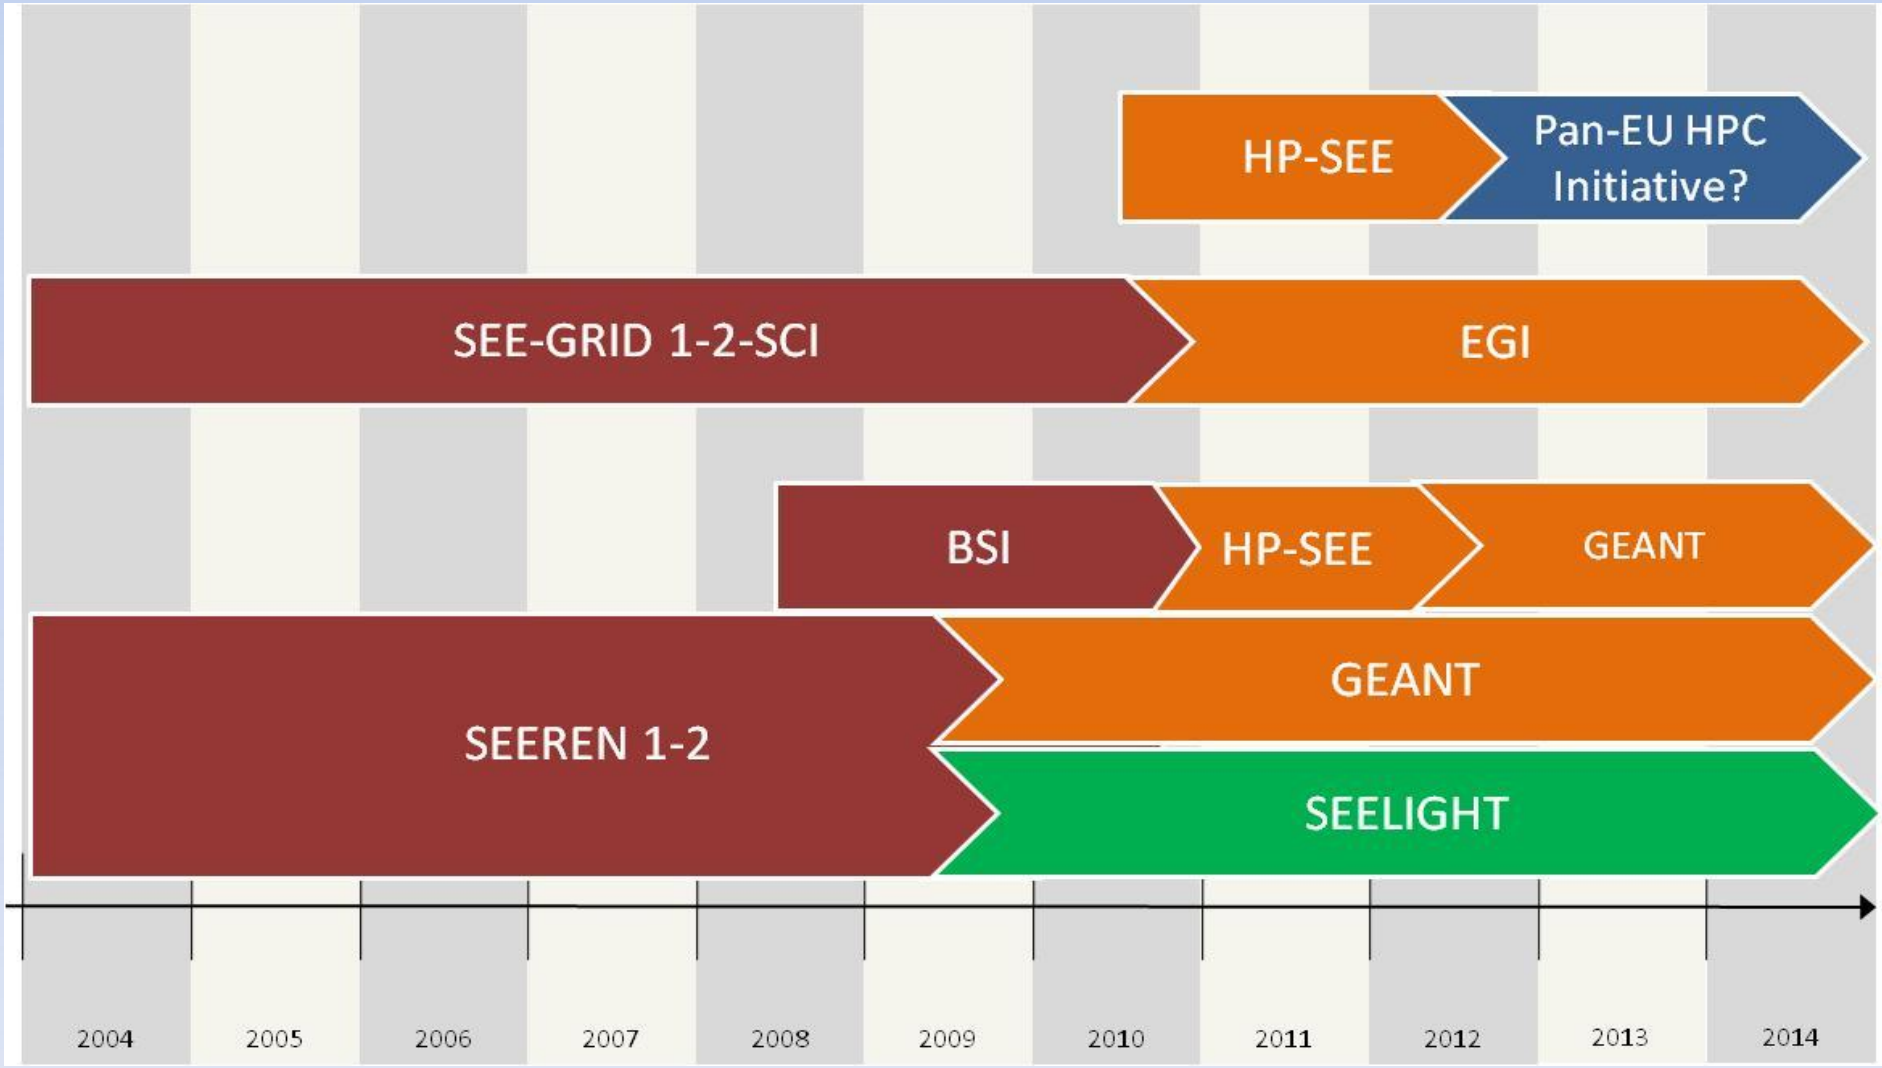

# Regional partnerships – from infrastructure to applications

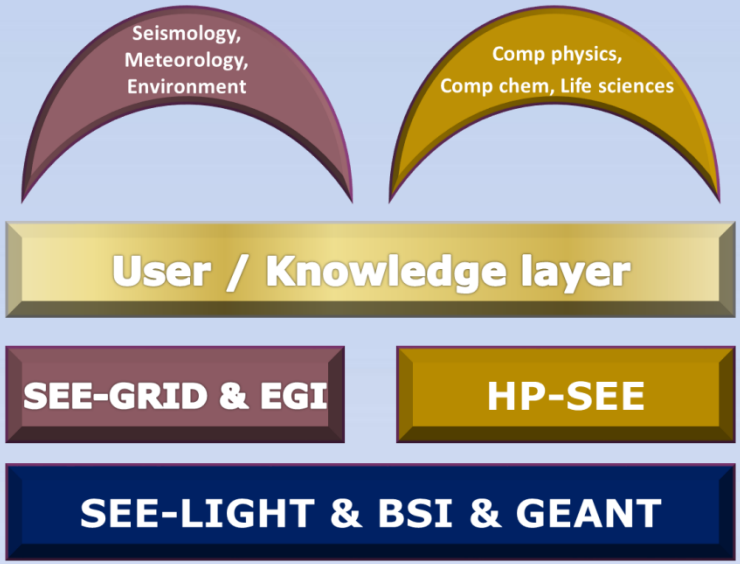


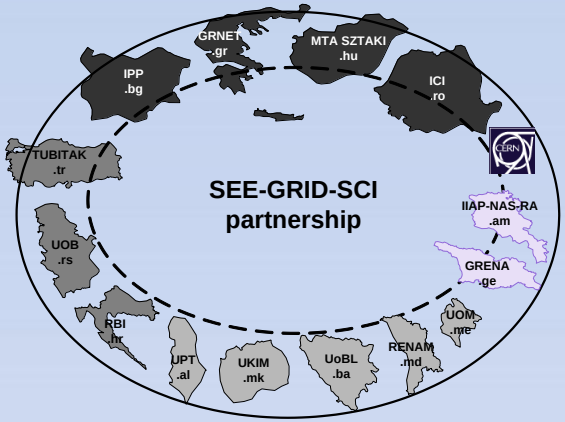


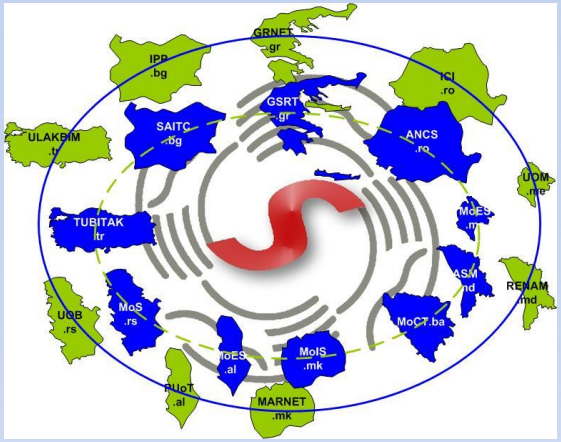


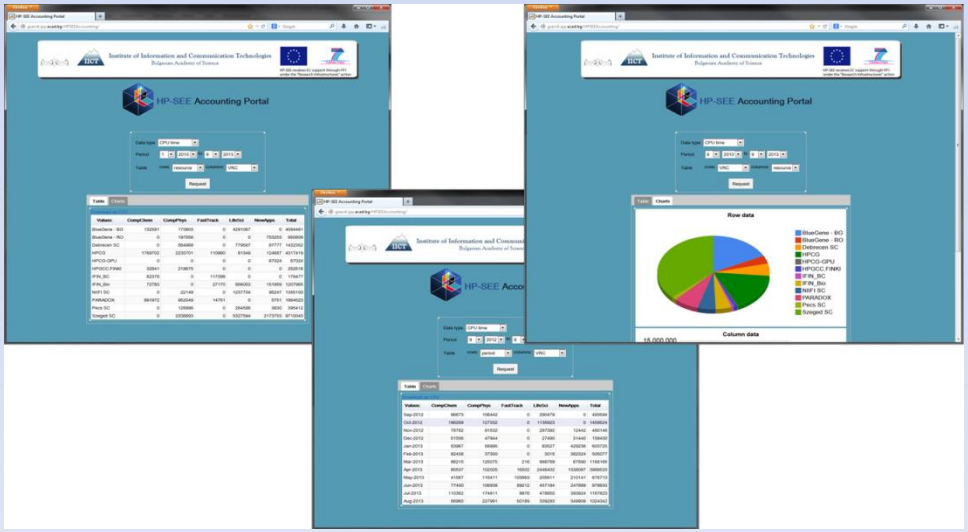


# National Roadmap for Research Infrastructures

## RI: National Center for HPC and Distributed Computing - NC4HPC&DC

Financial coordinator: Ministry of Education and Science; Ministry of Transport, Information Technology and Communications.

Scientific Coordinator: Institute of Information and Communication Technologies - Bulgarian Academy of Sciences (IICT-BAS)

| Consortium for supercomputers applications (since 2009 r.) | Consortium for distributed (Grid and Cloud) applications (set up - 2004, updated - 2012) |
|---|---|
| Sofia University "St Kliment Ohridski" | Institute for Nuclear Research and Nuclear Energy - Bulgarian Academy of Sciences (INRNE-BAS) |
| Technical University – Sofia (TU-Sofia | Sofia University "St Kliment Ohridski" |
| Medical University – Sofia (MU-Sofia) | Technical University - Gabrovo |
| National Institute of Geophysics, Geodesy and Geography -Bulgarian Academy of Sciences (NIGGG-BAS) | Institute of Molecular Biology "Roumen Tsanev" -Bulgarian Academy of Sciences (IMB-BAS) |
| | Institute of Organic Chemistry with Centre for Phytochemistry - Bulgaria Academy of Sciences ( IOCCP-BAS) |
| Institute of Mechanics – Bulgaria Academy of Sciences (IM-BAS) | National Institute of Geophysics, Geodesy and Geography -Bulgarian Academy of Sciences (NIGGG-BAS) |
| **Remark:** The consortium for supercomputer applications closely cooperates with Association "National Centre for Supercomputing Applications" (NCSA), which represents Bulgaria in PRACE | Institute of Mechanics – Bulgaria Academy of Sciences (IM-BAS) |
| | Institute of mathematics and Informatics – Bulgarian Academy of Sciences (IMI – BAS) |

# HPC resources for research in Bulgaria

## AVITOHOL at IICT-BAS

150x HP ProLiant SL250s Gen8 each with
- 2x Intel Xeon E5-2650 v2 (8C/16T),
- 64 GB DDR3-1866 RAM and
- 2x Intel Xeopn Phi 7120P

6x HP ProLiant DL380p Gen8 nodes with
- 2x Intel Xeon E5-2650v2 (8C/16T),
- 64 GB DDR3-1866 RAM

Infiniband 56 Gb/s FDR
Storage system with 96 TB

AVITOHOL

Total Performance:
RPeak:412.3 TFlop/s
RMax: 264.2 TFlop/s
Top 500 position: 389

## HPCG cluster at IICT-BAS

36 blades BL 280c(2x Intel X5560(4C/8T); 24GB DDR3);
8 management nodes HP DL 380 G6(2x Intel X5560(4C/8T); 32GB DDR3);
2 HP ProLiant SL390s G7(2x Intel E5649(6C/12T);96GB DDR3)
- 8x nVidia TESLA M2090 per server;

2 HP SL270s Gen8 (2x Intel Xeon E5-2650 v2(8C/16T); 128GB DDR3)
Total number of Xeon Phi 5110P coprocessors: 9
Total 132TBs of system storage

TOTAL PERFORMANCE:
RPEAK:22.94 TFlop/s

## NCSA IBM Blue Gene/P

8192 PowerPC 450 processors
4TBs of system memory
12TBs of system stotage
IBM proprietary interconnect with
2.5 μs latency and 10GBps bandwidth

TOTAL PERFORMANCE:
RPEAK:27.85 TFlop/s
RMAX:23.45 TFlop/s


## PHYSON at Sofia University

53 Intel Xeon x86_64 processors
524Gibs of system memory
6.5TBs of system storage
2x nVidia Tesla M2090 graphics processors

TOTAL PERFORMANCE:
RPEAK:3.57 TFlop/s
RMAX: 3.22 TFlop/s

## MADARA at IIOCCP-BAS

54 Primergy RX200 S5 servers with
2 Intel Xeon E5520(4C/8T) each
and a total of 800GB DDR3 1066MHz
20Gb/s DDR Infiniband
108TB System Storage by Fujitsu FibreCat SX100

# HPC Center “Avitohol” at IICT

| 150 HP Cluster Platform SL250S GEN8 servers with 2 Intel Xeon E 2650 v2 CPUs and 2 Intel Xeon Phi 7120P coprocessors | |
|---|---|
| Site | IICT-BAS/Avitohol |
| Manufacturer | Hewlett-Packard |
| Cores | 20700 |
| Interconnection | FDR InfiniBand |
| Theoretical Peak Performance | 412.3 Tflop/s |
| RMAX Performance | 264.0 Tflop/s |
| Memory | 9600 GB |
| Operation System | Red Hat Enterprise Linux for HPC |
| Storage capacity | 96 TB SAN |

Top500 List on 389 place (Nov 2015)
http://www.top500.org/system/178609

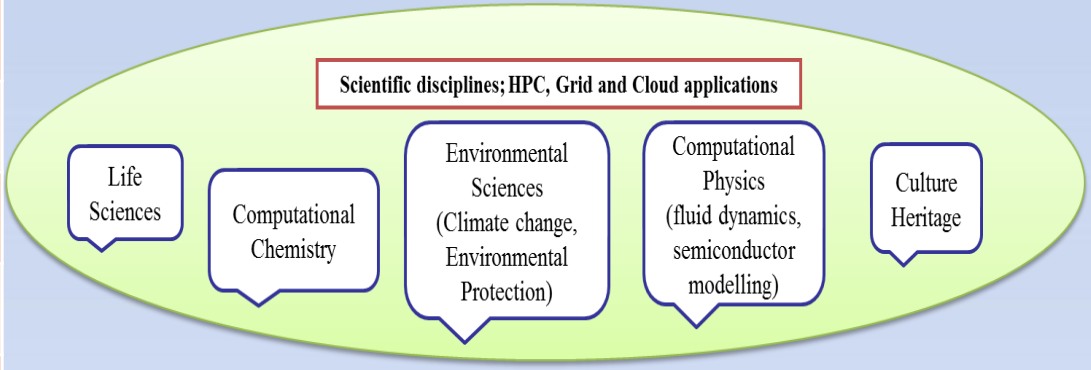


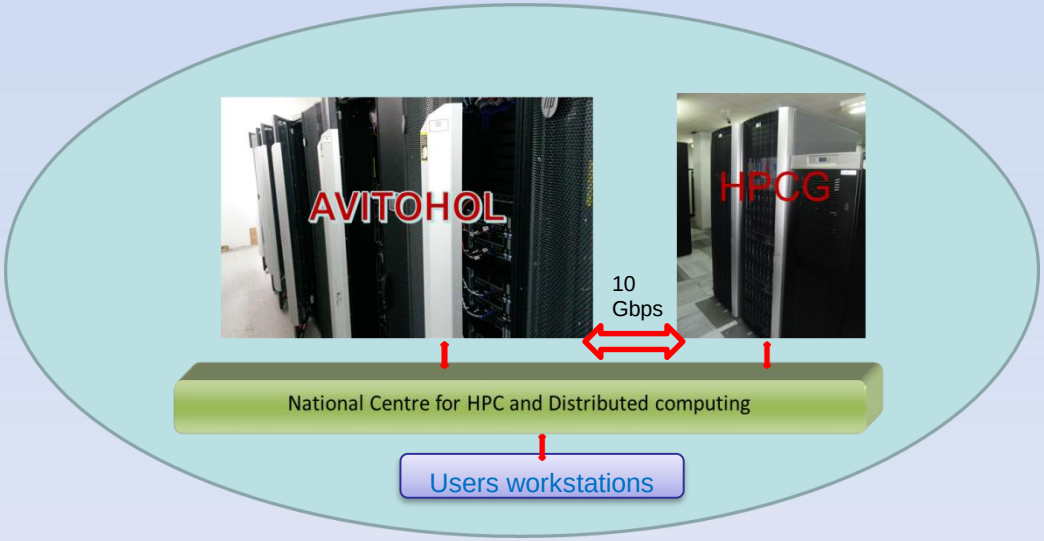

# The Supercomputer System Avitohol at IICT-BAS

- 8 water-cooled racks of type HP MCS 200, paired in couples.
- Each pair provides power and cooling for up to 50 kW of equipment
- About 90% of the computational power comes from the accelerators – one 7120P coprocessor achieves 1.25 TFlop/s in double precision.
- Total energy use at maximum load ~ 250kW.
- RHEL, Intel Cluster Studio, SAP HANA
- The most powerful system in the region

# Efficiency of sparse matrix – matrix multiplication on Intel Xeon Phi

**Computational time (s) of the algorithms on CPU by OpenMP. Mesh with N = 14259, M = 14259.**

| Number of processes | 8 | 16 | 32 |
|---|---|---|---|
| ALGORITHM 1 | 0.8494 | 0.7763 | 2.2047 |
| ALGORITHM 2 | 0.6047 | 0.7149 | 0.9330 |
| ALGORITHM 3 | 0.5619 | 0.3873 | 0.3994 |
| MKL | 0.7182 | 0.9029 | 0.9098 |

**Computational time (s) of the algorithms on Intel Xeon Phi by OpenMP. Mesh with N = 14259, M = 14259.**

| Number of processes | 30 | 61 | 122 | 183 | 244 |
|---|---|---|---|---|---|
| ALGORITHM 1 | 62.76 | 32.14 | 16.64 | 11.39 | 8.782 |
| ALGORITHM 2 | 0.4603 | 0.2982 | 0.2826 | 0.2853 | 0.3036 |
| ALGORITHM 3 | 0.3293 | 0.1664 | 0.1361 | 0.1822 | 0.4209 |
| MKL | 0.5469 | 0.3713 | 0.3174 | 0.3157 | 0.3285 |

# Regional HPC and distributed computing initiative – VI-SEEM

- VI-SEEM integrates all kinds of electronic infrastructure in the region – Grid, Cloud, HPC, BigData, in order to provide single point of access for scientists in the region of South Eastern Europe and the Mediterranean.
- Bulgaria is the leader in the HPC area.
  - Provides user-friendly integrated e-Infrastructure platform for Scientific Communities in:
    - Climatology,
    - Life Sciences,
    - Cultural Heritage

    for the SEEM region

    by linking compute, data, and visualization resources, as well as services, software and tools.

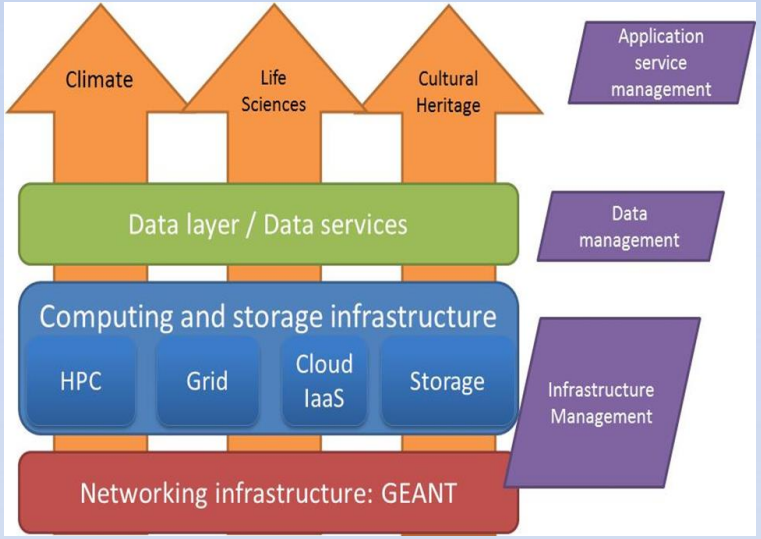


# Regional historical and geographical context

- VRE for regional Interdisciplinary communities in SEE and EM
- Merging of SEE and EM regions
- SEE: network SEEREN1-2, Grid SEE-GRID-1/2/SCI, HPC HP-SEE
- EM: HPC LinkSCEEM1-2

| Participant no. | Participant organisation name | Part. short name | Country |
|---|---|---|---|
| 1 (Coord) | GREEK RESEARCH AND TECHNOLOGY NETWORK S.A. | GRNET | Greece |
| 2 | THE CYPRUS INSTITUTE | CyI | Cyprus |
| 3 | INSTITUTE OF INFORMATION AND COMMUNICATION TECHNOLOGIES – BULGARIAN ACADEMY OF SCIENCES | IICT-BAS | Bulgaria |
| 4 | INSTITUTE OF PHYSICS BELGRADE | IPB | Serbia |
| 5 | NATIONAL INFORMATION INFRASTRUCTURE DEVELOPMENT INSTITUTE | NIIF | Hungary |
| 6 | WEST UNIVERSITY OF TIMISOARA | UVT | Romania |
| 7 | POLYTECHNIC UNIVERSITY OF TIRANA | UPT | Albania |
| 8 | UNIVERSITY OF BANJA LUKA | UNI BL | Bosnia and Herzegovina |
| 9 | SS CYRIL AND METHODIUS UNIVERSITY OF SKOPJE | UKIM | FYR of Macedonia |
| 10 | UNIVERSITY OF MONTENEGRO | UOM | Montenegro |
| 11 | RESEARCH AND EDUCATIONAL NETWORKING ASSOCIATION OF MOLDOVA | RENAM | Moldova |
| 12 | INSTITUTE FOR INFORMATICS AND AUTOMATION PROBLEMS OF THE NATIONAL ACADEMY OF SCIENCES OF THE REPUBLIC OF ARMENIA | IIAP-NAS-RA | Armenia |
| 13 | GEORGIAN RESEARCH AND EDUCATIONAL NETWORKING ASSOCIATION | GRENA | Georgia |
| 14 | BIBLIOTHECA ALEXANDRINA | BA | Egypt |
| 15 | INTER UNIVERSITY COMPUTATION CENTER | IUCC | Israel |
| 16 | SYNCHROTRON-LIGHT FOR EXPERIMENTAL SCIENCE AND APPLICATIONS IN THE MIDDLE EAST | SESAME | Jordan |

# VI-SEEM Regional Services

- Code Repository – GitLab
- GOCDB - topology and resources information.
- Monitoring – backend
- based on Nagios with
- custom frontend developed by GRNet
- Accounting portal – developed by IICT-BAS

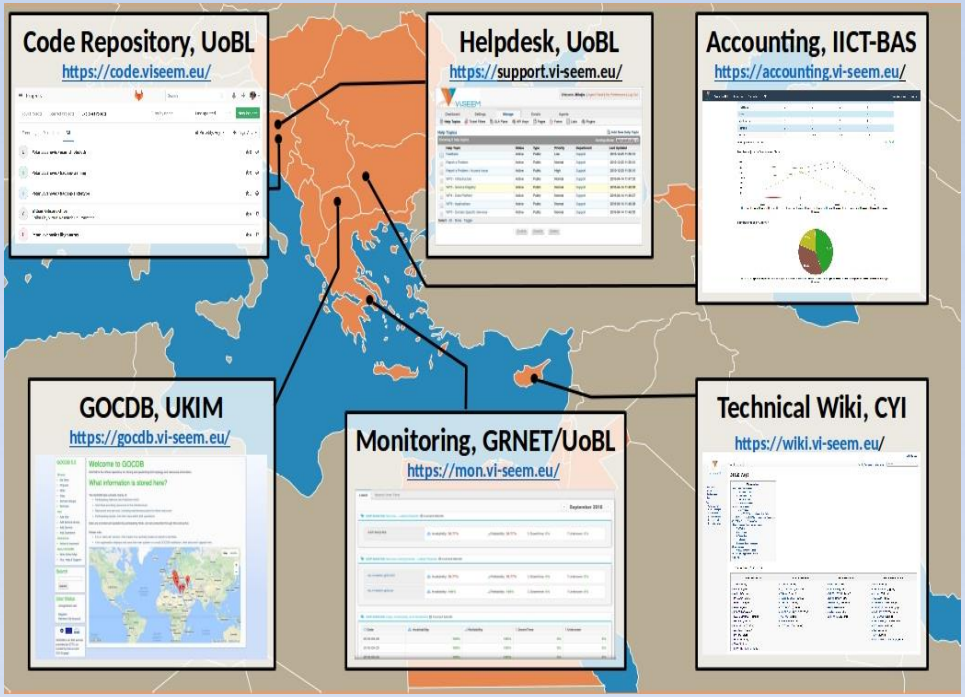

# GOCDB

- Overview of services deployed and registered in VI-SEEM GOCDB.


# VI-SEEM VRE Portal

- VRE Portal is available http://vre.vi-seem.eu/ Continuously populated with new content

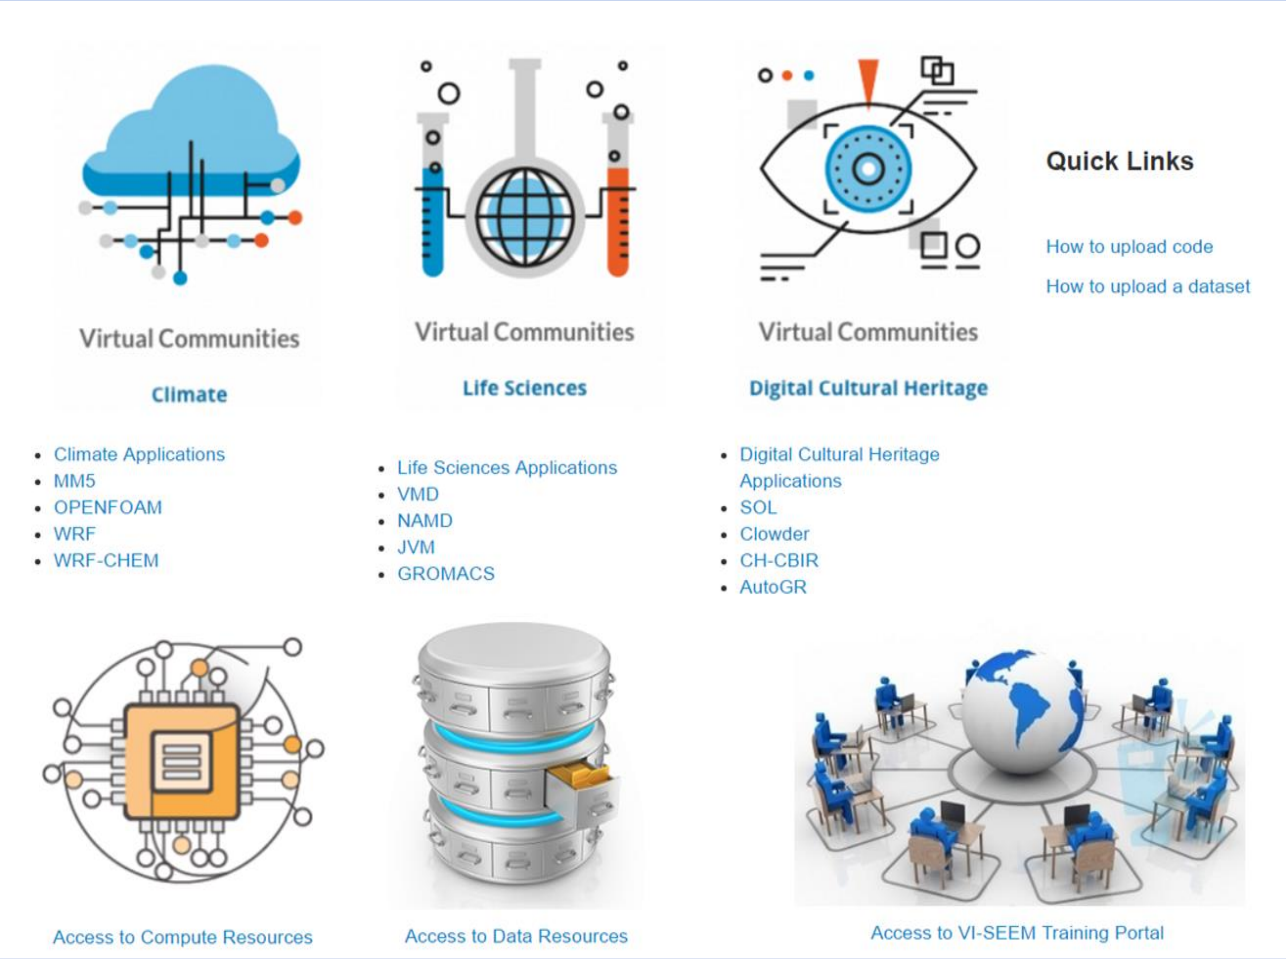

# Authentication and Authorization

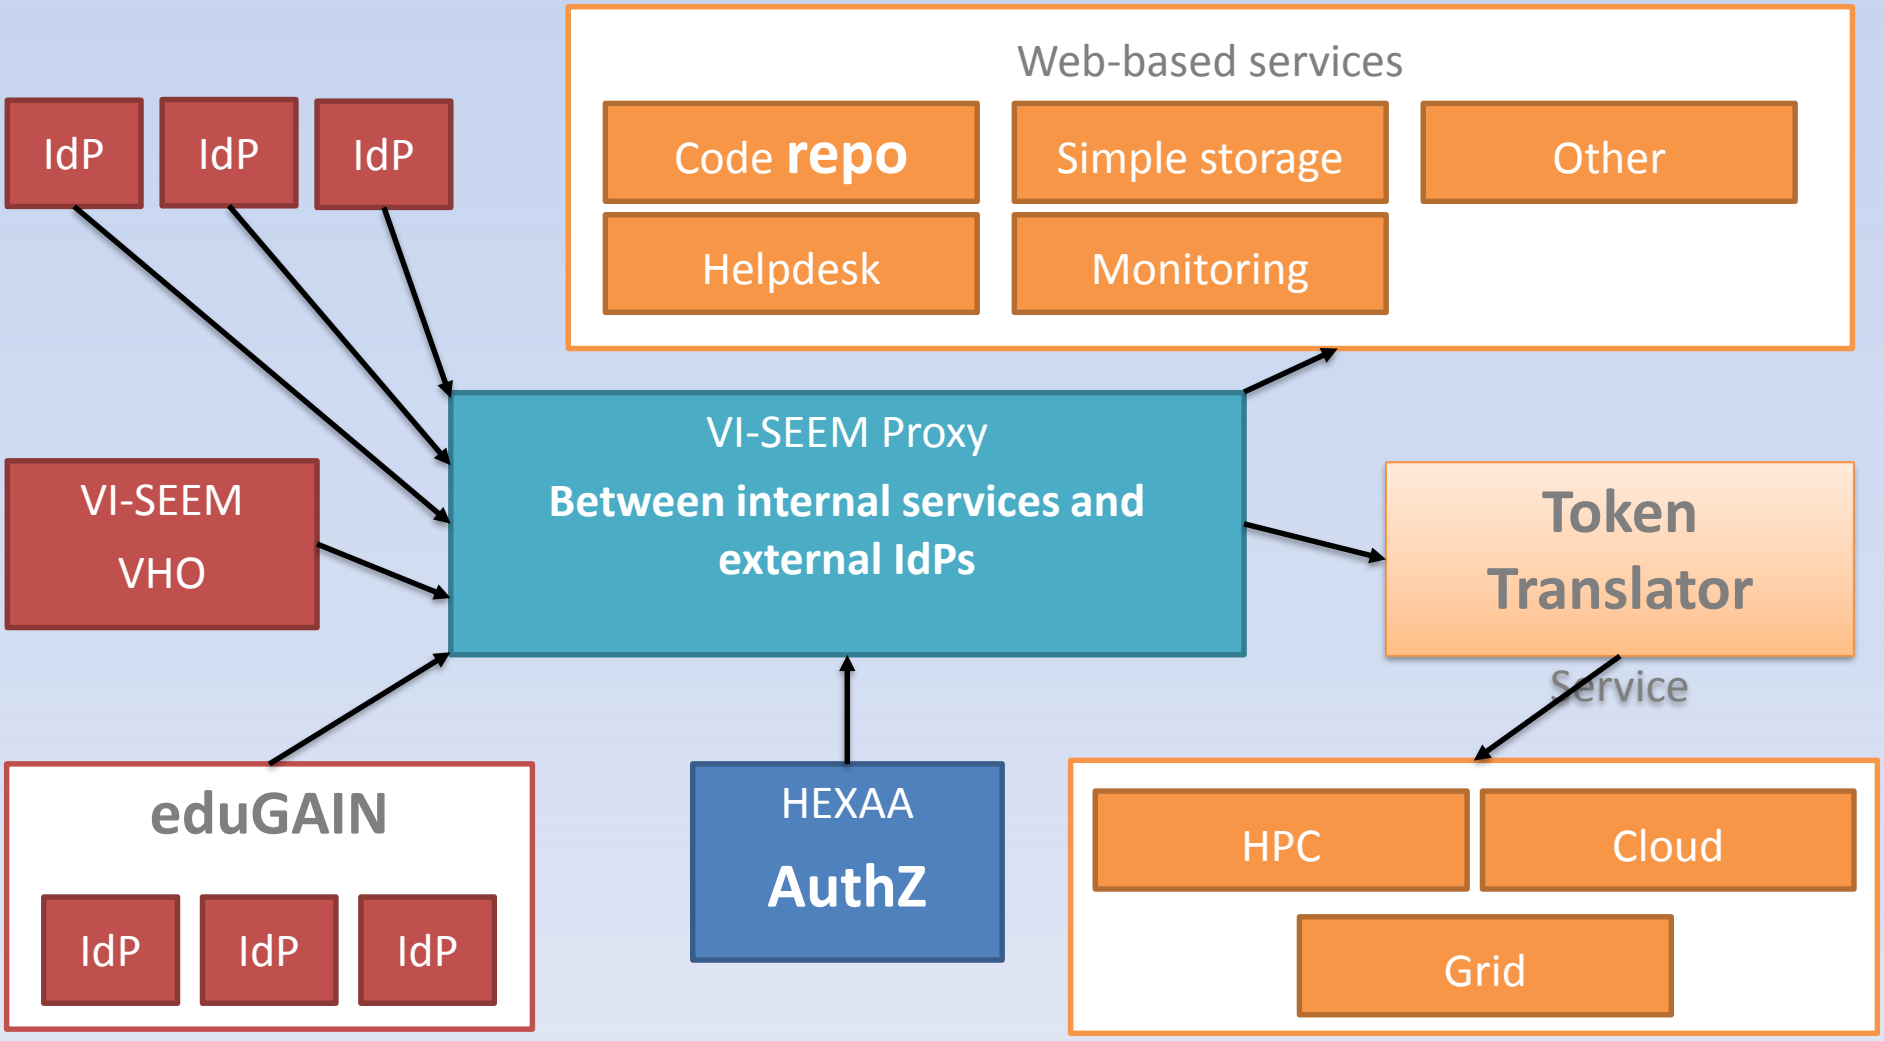

# HPC Resources in SEEM region

- Main HPC Resources

| Resource | Country | Total | | | Dedicated | | |
|---|---|---|---|---|---|---|---|
| | | CPU-cores | GPU-cores | Phi-cores | CPU-hours | GPU-hours | Phi-hours |
| **ARIS** | **Greece** | **8,520** | **-** | **-** | **3,000,000** | **-** | **-** |
| **Cy-Tera** | **Cyprus** | **1,392** | **16,128** | **-** | **1,829,088** | **21,192,192** | **-** |
| **Avitohol** | **Bulgaria** | **2,400** | **-** | **18,300** | **2,102,400** | **-** | **16,030,800** |
| **PARADOX** | **Serbia** | **1,696** | **108,544** | **-** | **742,848** | **47,542,272** | **-** |
| **NIIFI SC** | **Hungary** | **768** | **-** | **-** | **421,882** | **-** | **-** |
| **Leo** | **Hungary** | **1,344** | **628,992** | **-** | **588,672** | **275,498,496** | **-** |
| **InfraGRID** | **Romania** | **456** | **3,136** | | **350,400** | **5,494,272** | **-** |
| **ICAM** | **Romania** | **4,096** | **-** | **-** | **7,176,192** | **-** | **-** |
| **FINKI** | **FYROM** | **768** | **-** | **-** | **336,384** | **-** | **-** |
| BA-HPC | Egypt | 1,040 | - | - | 1,822,080 | - | - |
| **Gamma** | **Jordan** | **8** | **2,496** | **-** | **70,080** | **21,864,960** | **-** |
| | **Total** | **22,816** | **759,296** | **18,300** | **19,273,978** | **367,353,754** | **16,030,800** |

# Data management services

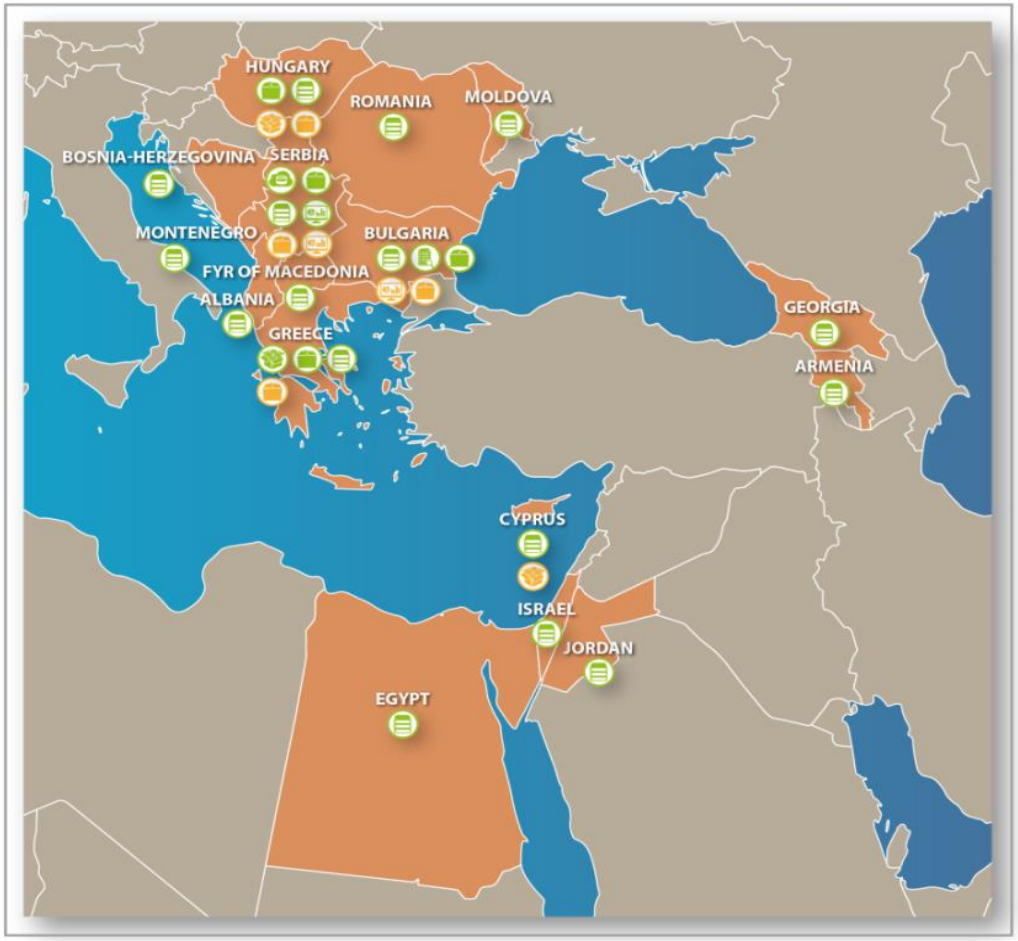


VI-SEEM Simple Storage Service (VSS)

VI-SEEM Repository Service (VRS)

VI-SEEM Work Storage Space / Local Storage And Data Staging (VLS)

VI-SEEM Archival Service (VAS)

VI-SEEM Data Discovery Service (VDDS)

VI-SEEM Data Analysis Service (VDAS)

SERVICES AVAILABLE FROM

INITIAL DEPLOYMENT PHASE

COMPLETE SETUP PHASE

# Accounting system

The accounting system covers HPC, Grid, Cloud, data resources and services

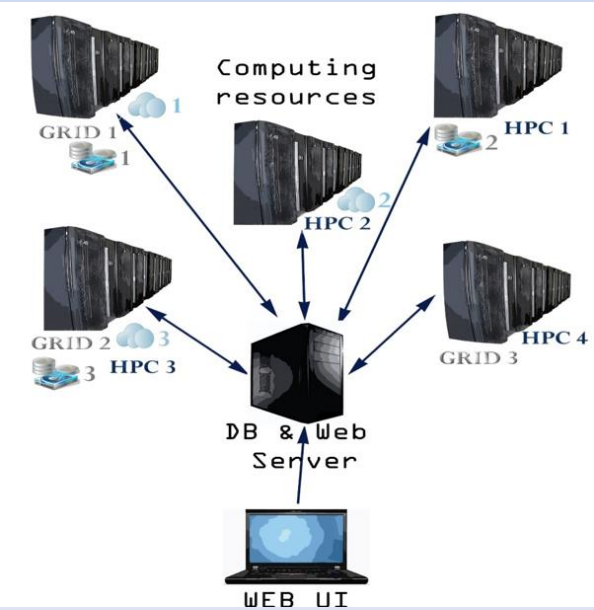


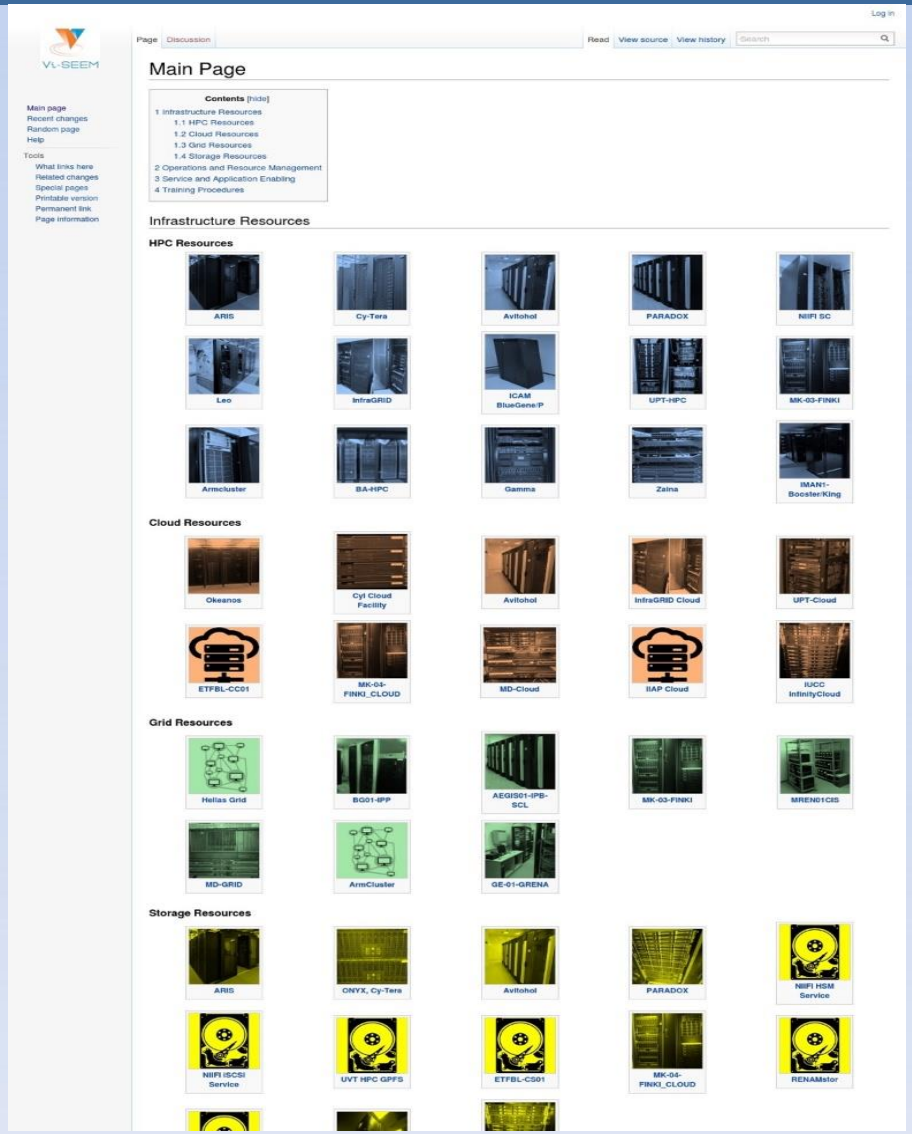

# Applications from Life Sciences

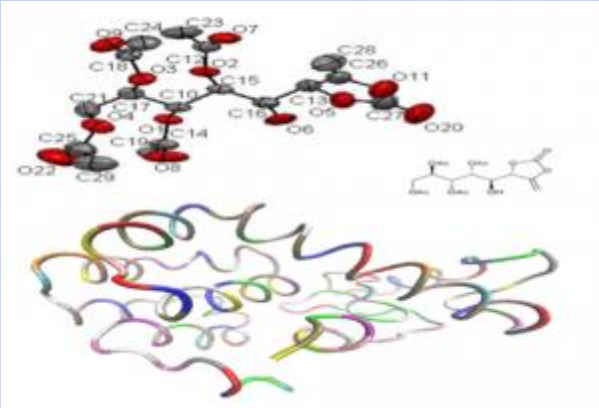

- Simulations can reduce the costs of developing new cures, which take 12-15 years and cost between 600 and 800 million US dollars, and accelerate process of compounds selection that are considered as part of new medicine.
- Molecular dynamics simulations using GROMACS and NAMD focused on the examination of interactions between synthetized molecule and proteins of the cell membrane.
- 17 new and previously unknown molecules were synthetized, which are based on the existence of enol carbonate structural units
- Cournia lab at the Biomedical Research Foundation, Academy of Athens participated in an international drug design competition. This competition was organized as a blind computational prediction of experimental results, and was coordinated by the [Drug Design Data Resource (D3R)], University of California. Cournia lab ranked 1st in the pose predictions out of 49 completed submissions with a median RMSD of 0.99 Å

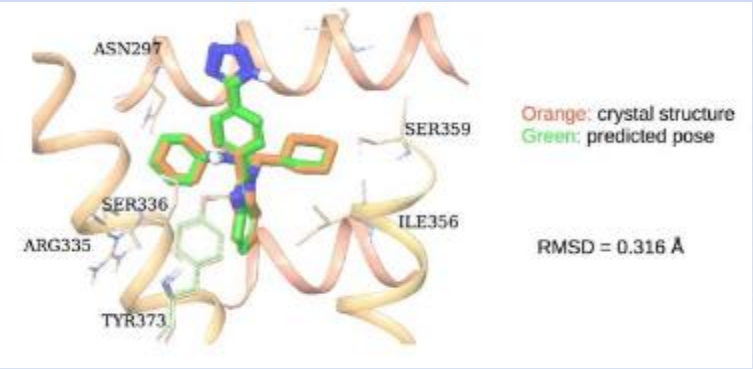

# Applications from climatology

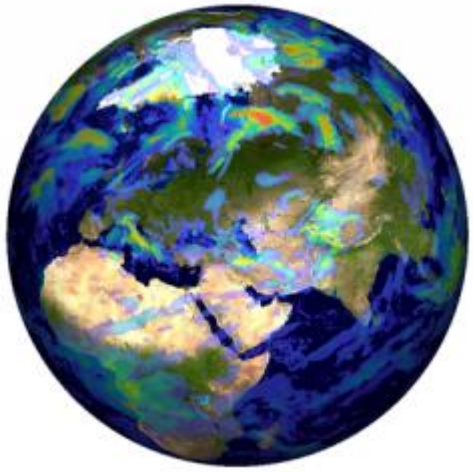

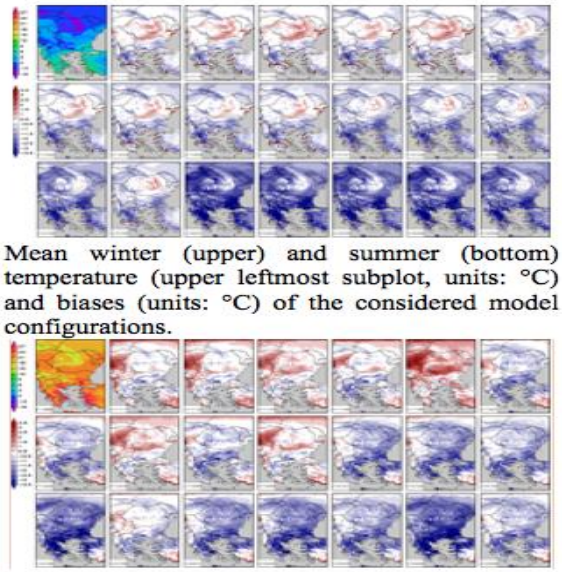

Mean winter (upper) and summer (bottom) temperature (upper leftmost subplot, units: °C) and biases (units: °C) of the considered model configurations.

Publicly released source code have been developed, using a hybrid GPGPU cluster, that automatically ports the kinetics calculations on GPU architectures.
Each GPU thread calculates the chemical concentrations of an individual atmospheric grid box.
The achieved performance showed up to 22.4× improvement of the kernel execution time

The latest version of the Regional Climate Model system RegCM4, is now fully supported by the ESP. It can be applied to any region of the World, with grid spacing of up to about 10 km and for a wide range of studies, from process studies to paleoclimate and future climate simulation. The study documents the performance of 20 different model configurations in representing the basic spatial and temporal patterns of the SE European climate for the period 1999-2009

# Applications from DCH

Fig. 1: Banatica collection in VI-SEEM DCH repository

BANATICA collection gathers together all the printed products considered monographs (brochures, books, yearbooks, calendars in volumes, prints with an individual cover, atlases, book-like printed scores etc.) which represent documentation sources for the culture and civilization on the Banat region.

anatica Virtual Library makes rare and old publications accessible again for a wider scientific audience in SEE region and enables further machine processing to be applied on the entire collection.

# Supercomputing in Science and Engineering

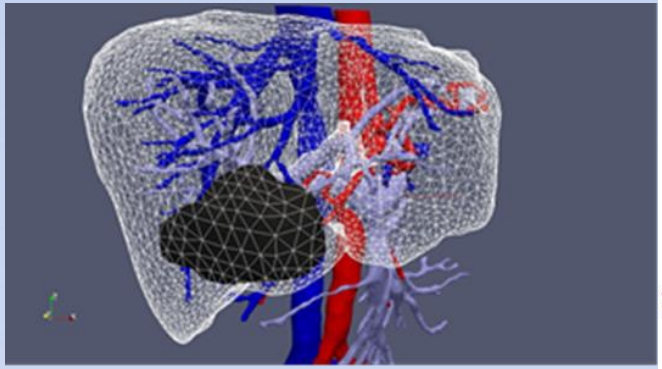

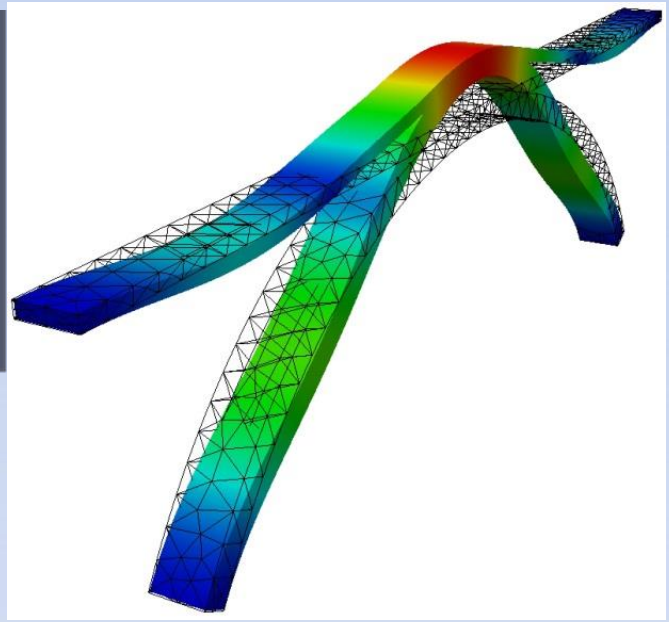

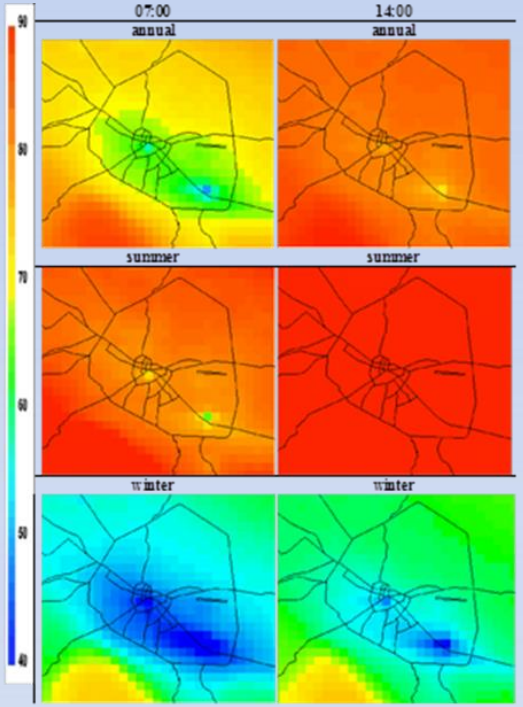


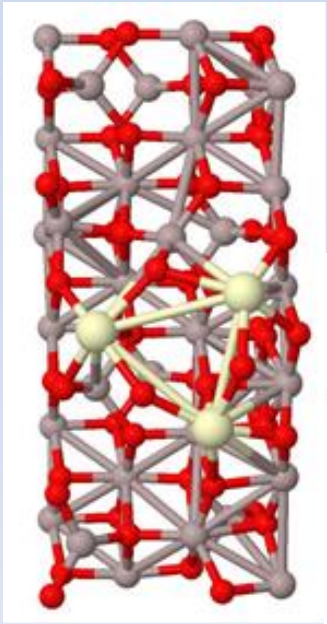

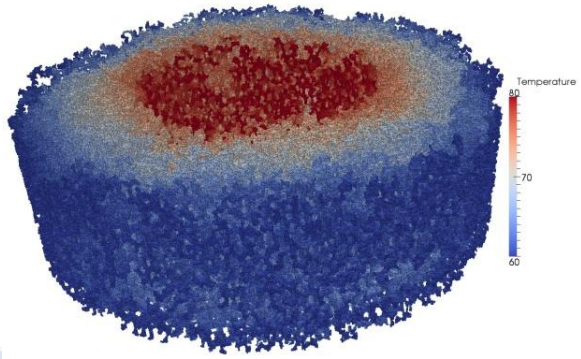


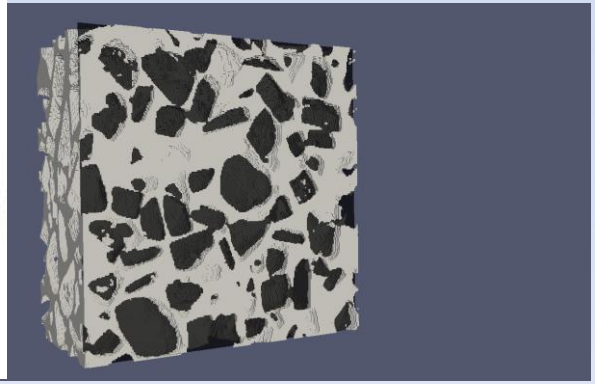

# Strong scalability of the shooting method

The parallel numerical tests are carried out on Avitohol supercomputer. The supercomputer has 300 CPUs (Intel Xeon E5-2650 v2 8C 2600 GHz) organized on 150 servers with non-blocking InfiniBand FDR. Each server has 64 Gb RAM.

A mesh of 5376 quadrilateral elements is used resulting into a system of 16 419 DOF.

| P | CPU [s] | Speed Up | Efficiency [%] |
|---|---|---|---|
| 16 | 6 169 | - | - |
| 32 | 3 152 | 1.96 | 97.85 |
| 64 | 1 582 | 3.90 | 97.45 |
| 128 | 838 | 7.36 | 91.99 |
| 256 | 446 | 13.82 | 86.37 |
| 512 | 262 | 23.51 | 73.45 |

# Conclusions and future work

- The establishment of an integrated HPC infrastructure in the region, which supports also Data and Cloud resources and has a comprehensive set of operational and user-facing services, enables high quality research.
- Apart from the initially designated user communities the resources were opened to other scientific areas and to SMEs in the region through open calls.
- The operational teams follow established models and distribute the work efficiently between partners, achieving economies of scale and providing an integrated portfolio of services, as envisaged for EOSC.
- The diversity of the underlying infrastructure allows for optimal distribution of the applications.