

Marian Bubak<sup>1,2</sup>, Piotr Nowakowski<sup>1,2</sup>, Jan Meizner<sup>1</sup>, Martin Bobák<sup>3</sup>, Ondrej Habala<sup>3</sup>, Ladislav Hluchý<sup>3</sup>, Viet Tran<sup>3</sup>, Adam S. Z. Belloum<sup>4</sup>,

Reginald Cushing<sup>4</sup>, Maximilian Höb<sup>5</sup>, Dieter Kranzlmüller<sup>5,6</sup>, and Jan Schmidt<sup>5</sup>

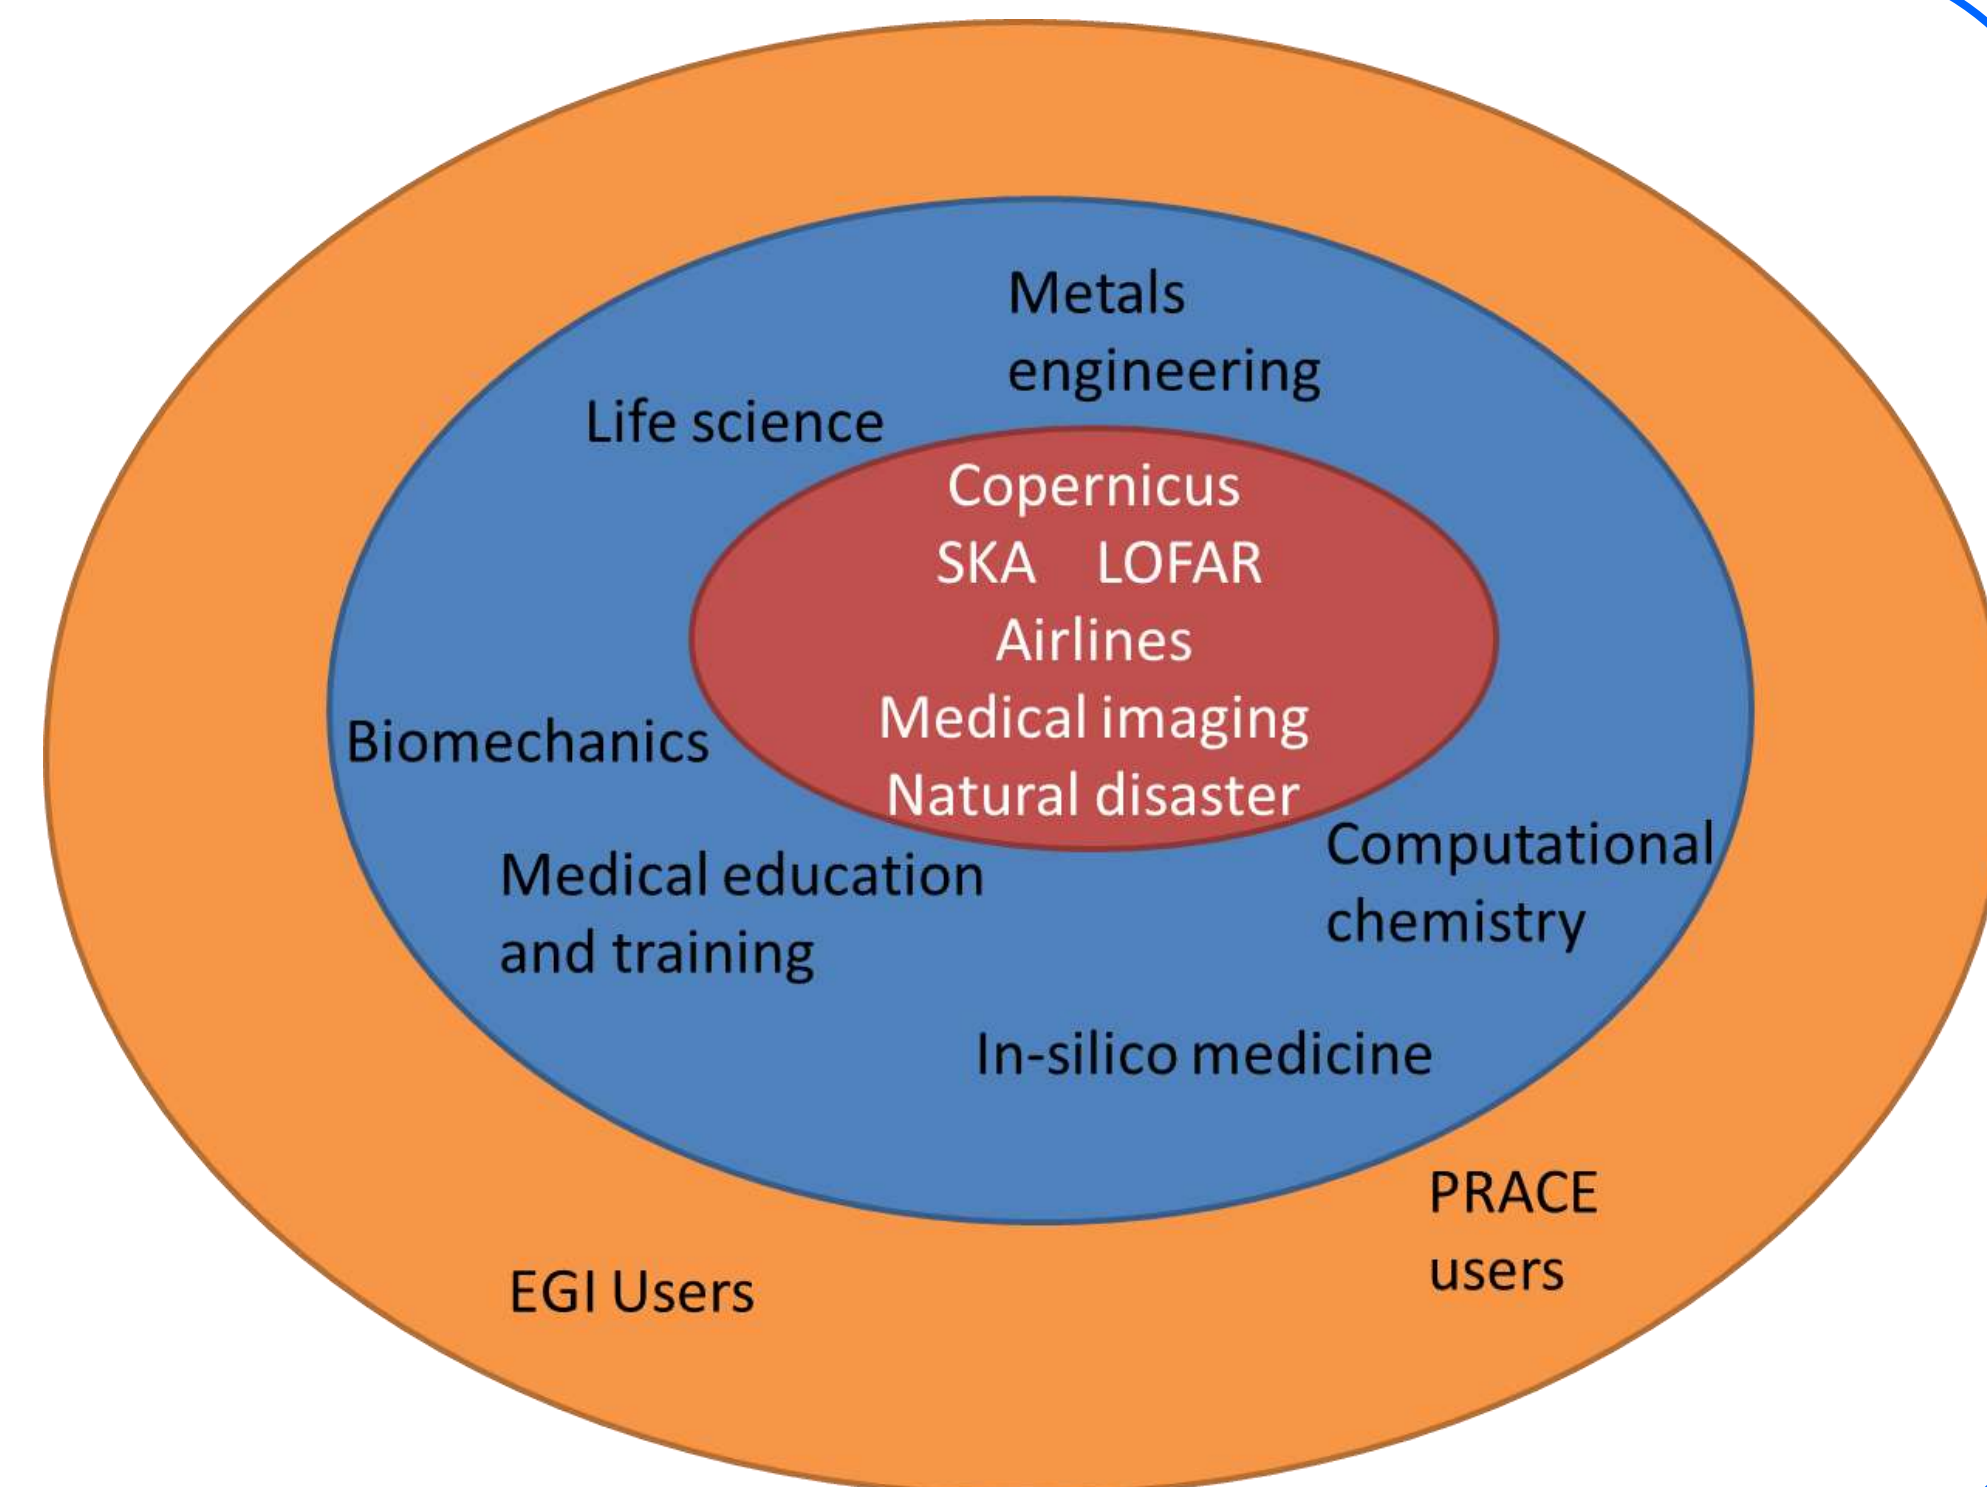
<sup>1</sup>ACC Cyfronet, AGH University of Science and Technology, Krakow, Poland; <sup>2</sup>Sano Centre for Computational Personalised Medicine, Krakow, Poland;

<sup>3</sup>Institute of Informatics, Slovak Academy of Sciences, Bratislava, Slovakia; <sup>4</sup>Informatics Institute, University of Amsterdam, Amsterdam, Netherlands;

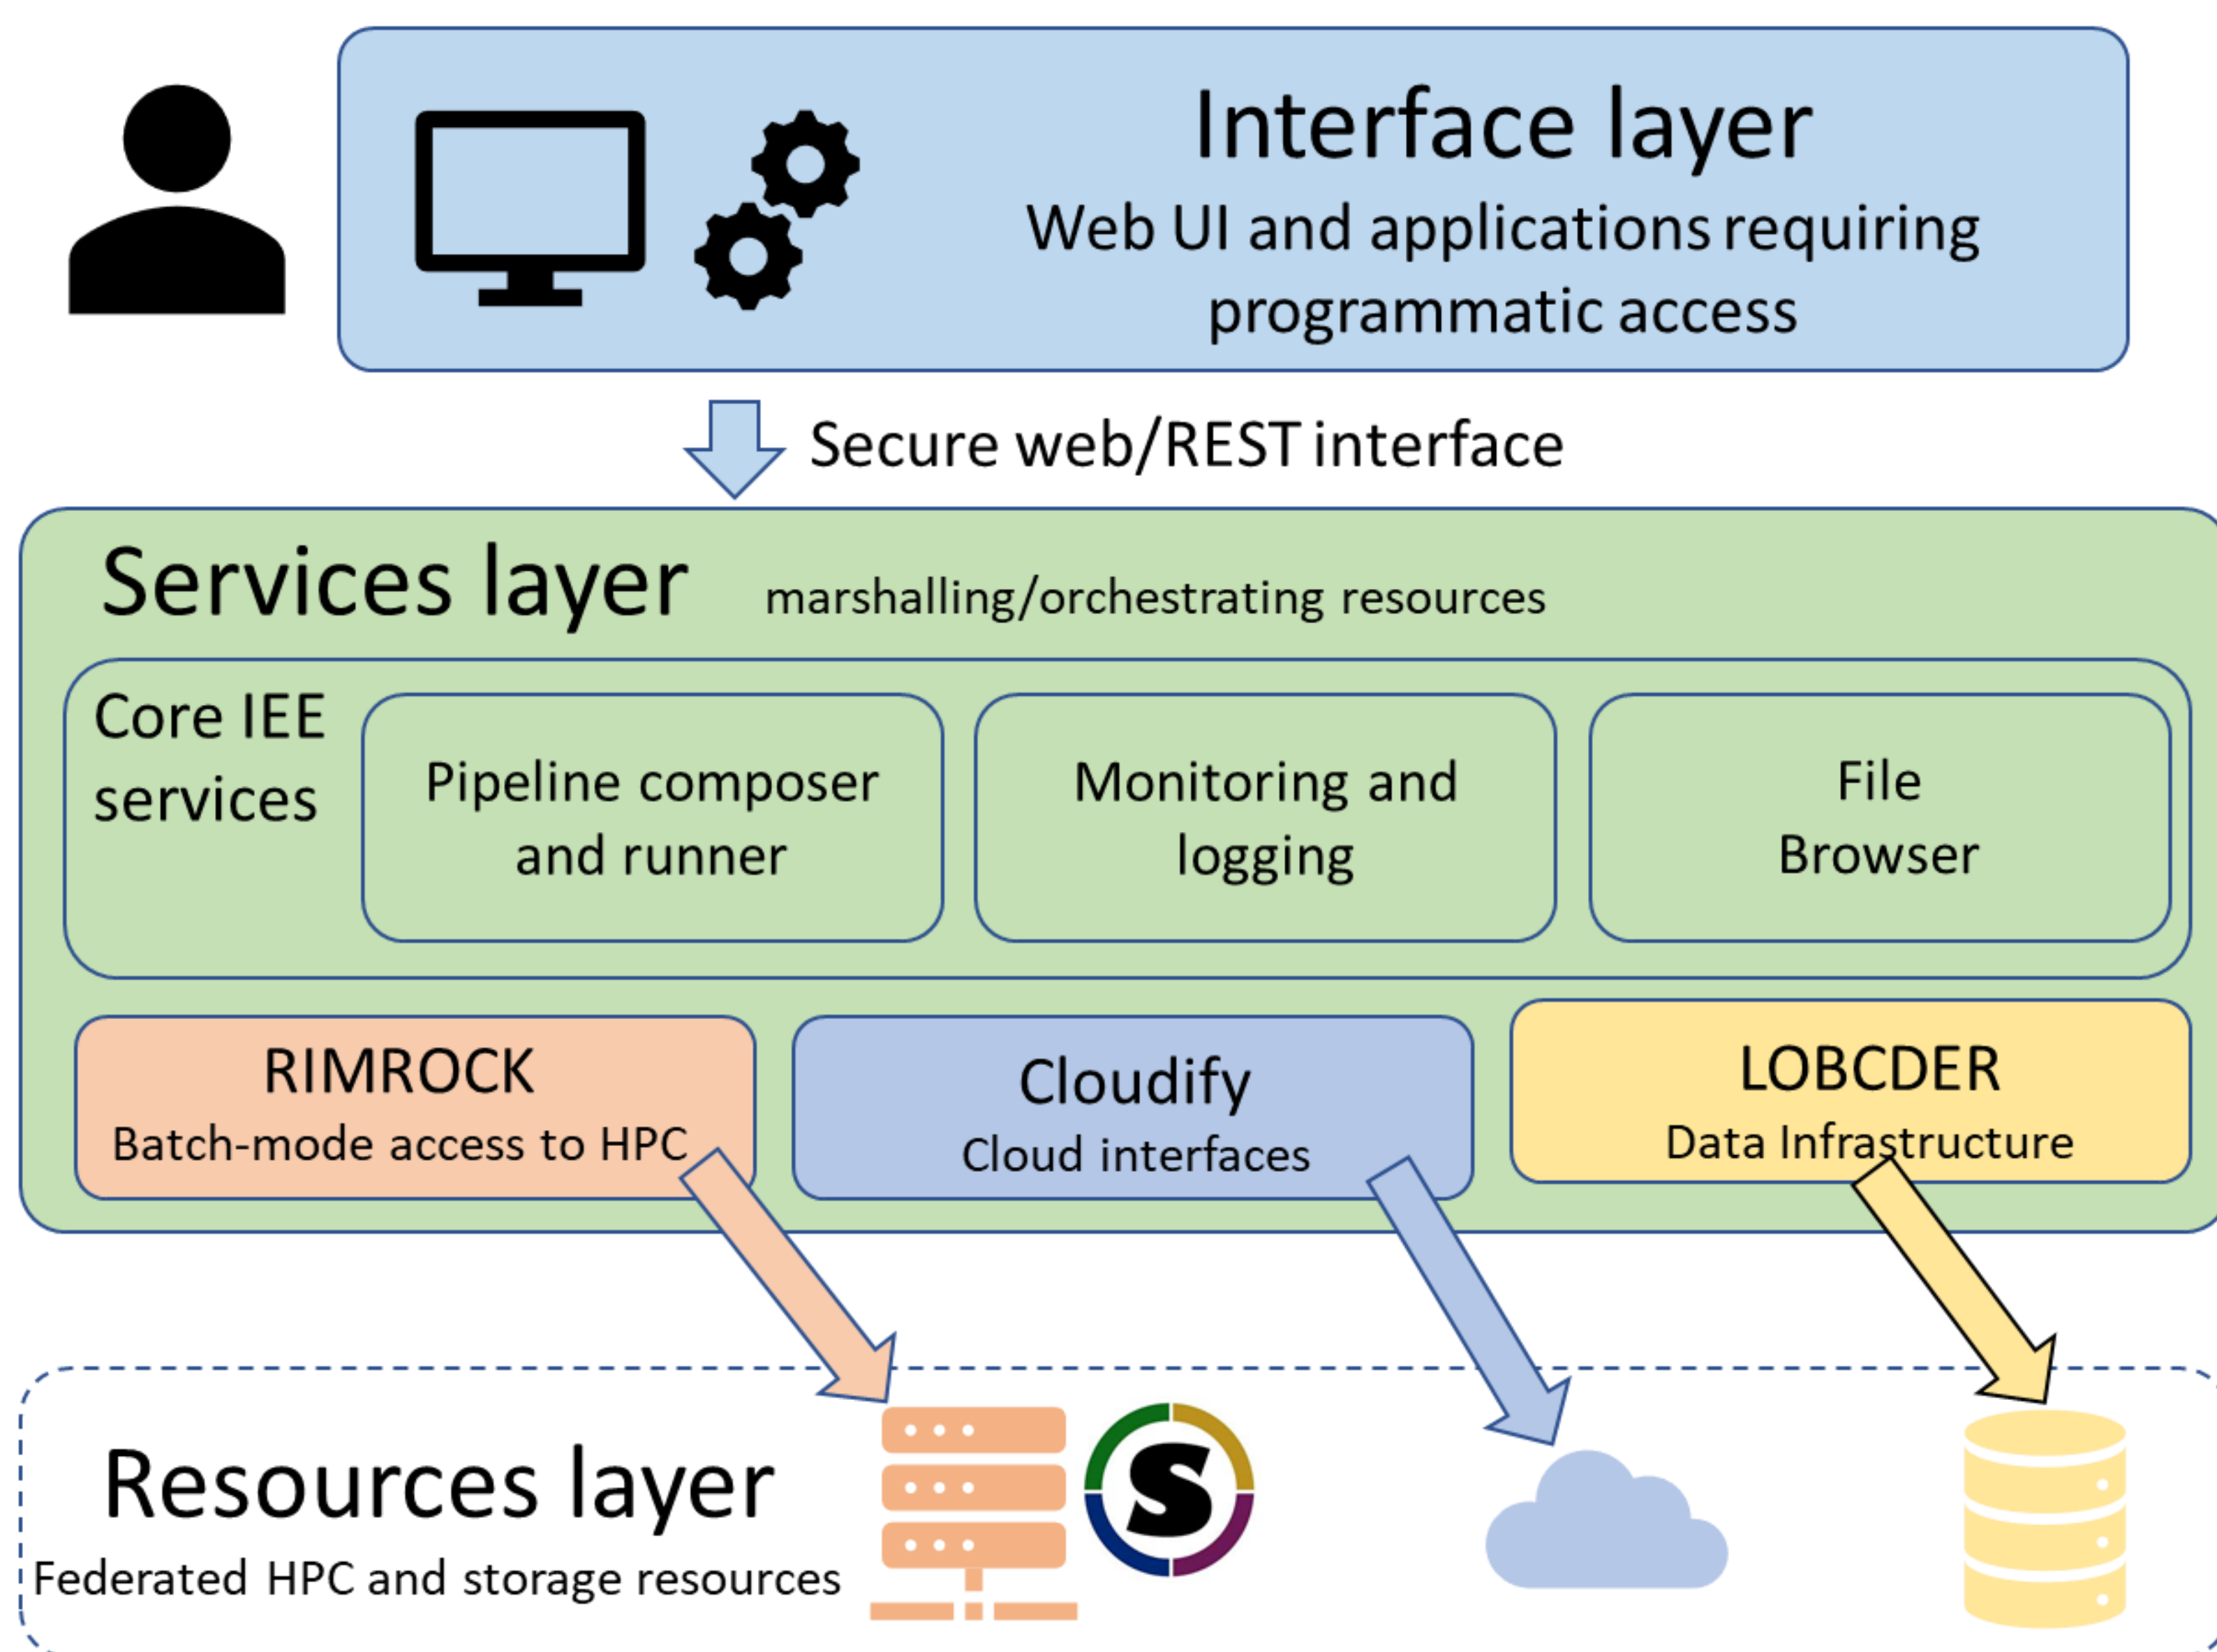
<sup>5</sup>MNM-Team, Ludwig-Maximilians-Universität, Munich, Germany; <sup>6</sup>Leibniz Supercomputing Centre, Garching, Germany

### Goals

- Provide exascale ready computational and data services that will accelerate innovation
- Validate the services in real-world settings of scientific research and in industry pilots:
  - Computation using data from the LOFAR telescope, technology preparation for the SKA challenges
  - Medical use-case utilizing AI on CPUs and GPUs for cancer diagnostic based on imaging data
  - Supporting innovation based on global disaster risk data
  - Business use-case for building models enabling revenue optimization and real-time calculations of prices
  - Agricultural simulations utilizing data from the part of Copernicus dataset collected from the Sentinel satellites



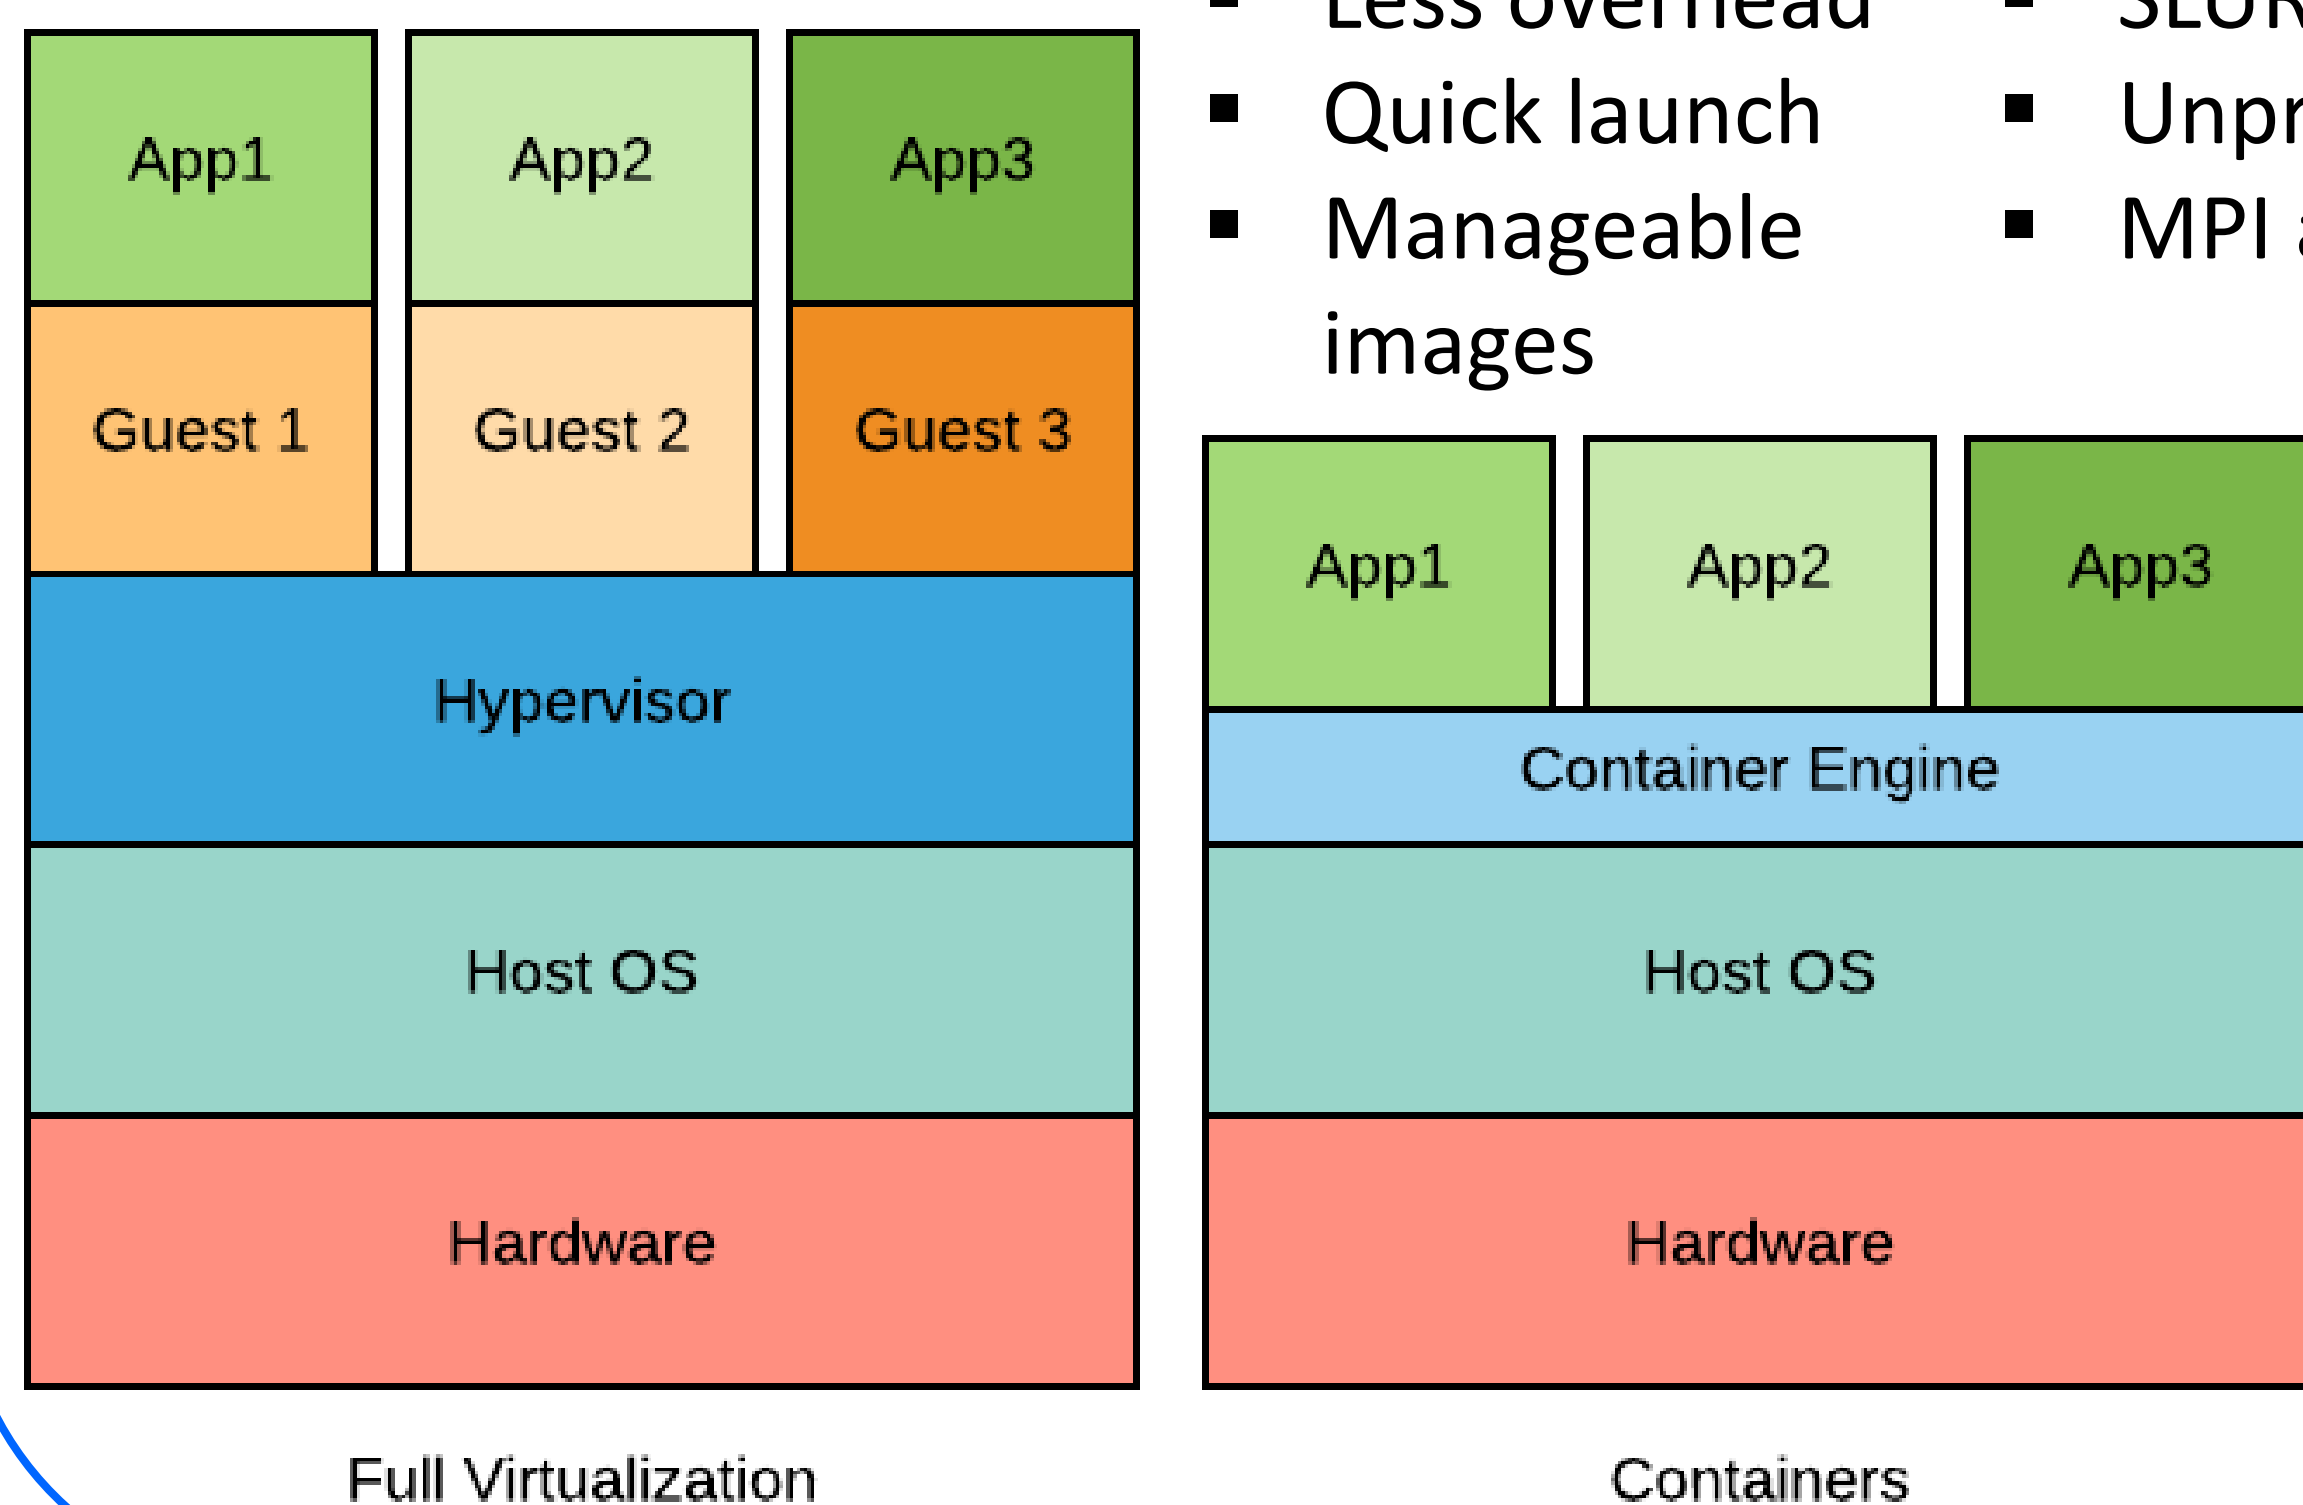
### Extreme Large Scale Computing Services



- Based on “focus on services and forget about infrastructures” idea
- Support computational activities: analysis, data mining, pattern recognition, etc.
- Use heterogeneous research datasets: input and output data from modelling, simulation, visualization and other scientific applications stored in data centers and storage systems available on European e-infrastructures
- Support HPC and cloud-based computations needed for various data analyses
- Utilize Docker, Singularity or Charliecloud containers as required by the given infrastructure type and location
- Follows simplified workflow:
  - User defines pipeline/steps in IEE via web UI
  - IEE reads a computation for execution or starts it automatically (as requested)
  - Pipeline steps are executed – container steps can be scheduled as HPC jobs via Rimrock
  - HPC queues jobs (SLURM) and executes them when resources become available
  - Inside each job, the appropriate code runs embedded containers with appropriate options
- Data transfers are handled automatically by the LOBCDER component and data infrastructure

### Why containers? Why Singularity?

- Small footprint
- Less overhead
- Quick launch
- Manageable images
- Built for HPC
- SLURM support
- Unprivileged
- MPI and GPUs



### Conclusions and future work

- The current status of the PROCESS platform:
  - integrated solution
  - running use cases on heterogeneous e-infrastructures
  - support for a range of core technologies (HPC, Cloud)
  - applications containerized for portability
  - multiple geographical locations
- In the future we will work on:
  - further stabilization of the solution
  - making it as user friendly as possible
  - streamlining data transfers across technologies and sites
  - concurrent execution of applications on multiple sites
  - supporting a multi-node MPI on Singularity (with Horovod)
  - scaling up pipelines to run on SuperMUC-NG (TOP500 #9)
  - improving interface between IEE and the LOFAR portal

### Acknowledgements



This work is supported by the “PROviding Computing solutions for ExaScale Challenges” (PROCESS) project that received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 777533.

This research was supported in part by PL-Grid Infrastructure. The authors also want to thank Mara Graziani, Balázs Somoskői, Jason Maassen and Henning Müller for their valuable insight on the presented subject.

### References

- M. Bobak, A. S. Z. Belloum, P. Nowakowski, J. Meizner, M. Bubak, M. Heikkinen, O. Habala, and L. Hluchý, “Exascale computing and data architectures for brownfield applications,” in Fuzzy Systems and Knowledge Discovery (FSKD), 2018 14th International Conference on. IEEE, 2018, pp. 461 – 468.
- M. Bobak, A. S. Z. Belloum, R. Cushing, J. Meizner, P. Nowakowski, V. Tran, O. Habala, J. Maassen, B. Somoskői, M. Graziani, M. Heikkinen, M. Höb, J. Schmidt, L. Hluchý, “Reference exascale architecture”, 15th International Conference on eScience, San Diego, USA
- Docker, <https://www.docker.com/>
- Singularity, <https://sylabs.io/singularity/>
- Charliecloud, <https://hpc.github.io/charliecloud/>
- The PL-Grid National Computing Infrastructure, [www.plgrid.pl](http://www.plgrid.pl).