

ERLANGEN REGIONAL COMPUTING CENTER



Prozessorientierte Dienststruktur
für Performance Engineering
von wissenschaftlicher Software
an deutschen HPC-Zentren

Jan Eitzinger
20.7.2017

Overview



Call:

Performance Engineering für wissenschaftliche Software

Partners:



Duration:

03/2017 – 02/2020

Coordination: Prof. G. Wellein (J. Eitzinger)

Associated Partners

- **KONWIHR**



Provide resources to supervise performance projects

- **TU Munich (Prof. Bungartz, Prof. Bader)**



Provide the algorithmic perspective on PE which is often far more important than an efficient implementation

- **Forschungszentrum Jülich**



Knowledge transfer within JARA-HPC SimLabs

Technical University Bergakademie Freiberg

Included in joint HPC support strategy



Current state

- HPC competence in German HPC centers distributed across country
- Gauss-Allianz is an initiative to integrate and organize TIER 2/3 HPC landscape in Germany
- Multiple local efforts and island projects:
bwHPC, KONWIHR, HKHLR, HLRN ...



Our contribution

- Similar targets as sketched in **GA Strategiepapier**, but focus on Performance-Engineering sub-topic

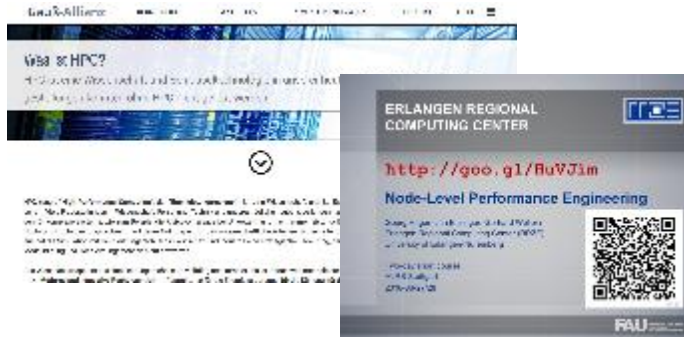
Integrate with and **built on** already existing efforts and further drive the final goal of an hierarchical and yet integrated German HPC infrastructure.

Major Building Blocks

- **Dissemination** – Increase publicity of project and raise general awareness for performance issues

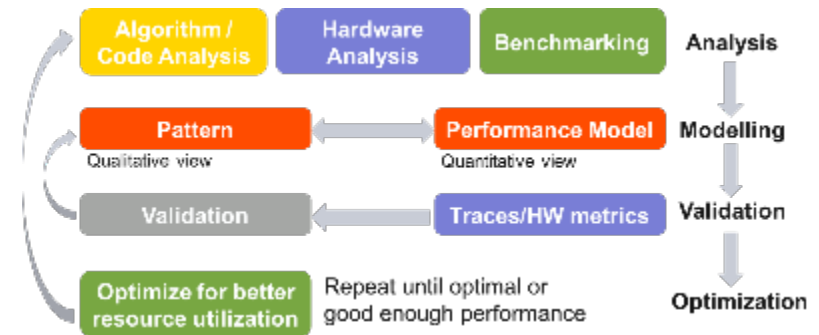


We want to talk with you about your PE problem!



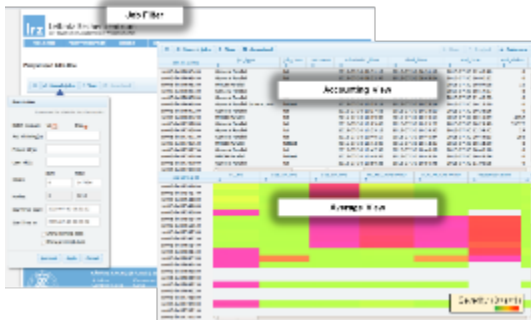
- **Documentation** – Build a central web offering, create content and provide resources to maintain it

- **Structured PE-Process** – Systematic bottleneck centric performance analysis and optimization process




Major Building Blocks cont.

- **PE Support Infrastructure** – Process blueprint for nation-wide aligned support effort



- **Application Monitoring and Analysis** – Automatic profiling and bottleneck analysis for all applications running on a HPC-System



- **HPC Curriculum** – Coordinated nation-wide Workshop and Tutorial program



Central Information Hub

Establish PE specific web-offering under GA umbrella:

- Central **curriculum** with local, regional and nation-wide events
- Downloadable **teaching material**: Videos, slides and exercise codes
- **FAQ** and **HPC-Handbook** section: Typical HPC tasks and activities organized in beginner/intermediate/advanced level
- **Case Studies** with example codes and whitepapers
- **Q&A** section: Ask questions and get expert answers ala StackExchange

Organize, create and review above content together with GA staff!



PE Support Infrastructure

- Multi-Tier distributed **support infrastructure** which allows to hand-over requests and allocate specialists from other centers
- Create a process for *Performance Projects* allowing to
 - **Keep track** of and **transfer** projects between sites and find the right expert for a specific problem
 - **Carry out** and **document** efforts and results in a **standardized coherent way**
 - **Pack** an already started **project** between sites so that experts can pick it up right away

Actual Doing

- Not enough **manpower** to fully unroll all of the points

We will create a **blueprint** and **develop** the necessary **tools** and **processes**.




Still it is required to **show-case** how this will look like by

- Carrying out real **Performance Projects**
- **Organize tutorials** between sites
- Organize **researcher exchanges** between sites

Application Performance Monitoring

Global automatic **application performance monitoring** is essential to improve **efficient** usage of HPC systems



Courtesy of LRZ

Targets:

- Give user immediate feedback on job runs
- Identify applications with high optimization potential or pathological performance behavior
- Create databases with **performance footprints** and **performance maps** to characterize applications and track HPC usage statistics

How will this be a success?

We need to create a **Win/Win** situation for a all participating HPC centers.


The centers must see enough **benefit** from the offerings to be happy to **contribute** and **use** it.

Reach critical mass!



Content and services must be **high quality, complete, structured** and **easy to find**, so that people will use the PE offerings without even thinking about it!

ERLANGEN REGIONAL COMPUTING CENTER



Thank You.