



DFG Deutsche Forschungsgemeinschaft

Grant Nr. 01IH13009 Duration: 03/2017 – 02/2020

ProPE – A joint effort to establish a unified service infrastructure for Performance Engineering in German HPC-Centers

G. Wellein, J. Eitzinger, T. Röhl
University of Erlangen-Nürnberg
Erlangen Regional Computing Center (RRZE)

M.S. Müller, A. Dammer, N. Filla, D. Schürhoff
RWTH Aachen University
IT Center

W. Nagel, R. Dietrich, F. Winkler
Technical University Dresden
Centre for Information Services and
High Performance Computing (ZIH)

Overview and Partners

HPC competence in German HPC centers is distributed across the country. The Gauss-Allianz is an initiative to integrate and organize TIER 2/3 HPC landscape in Germany. Furthermore there are multiple local efforts: bwHPC, KONWIHR, HKHLR, HLRN and JARA-HPC. Our contribution is to integrate with and built on already existing efforts and further drive the final goal of an hierarchical and yet integrated German HPC infrastructure with an emphasis on Performance-Engineering.

Partners

- RRZE (University Erlangen-Nuremberg)
- IT Center (RWTH Aachen University)
- ZIH (Technical University Dresden)

Associated Partners

- KONWIHR
- TU Munich (Prof. Bungartz)
- Forschungszentrum Jülich
- Technical University Bergakademie Freiberg

Documentation and Dissemination

Building a central hub, including a knowledgebase with HPC and Performance Engineering materials, sorted to addressing different target groups. Currently started by lowering the entry barrier into HPC providing a documentation aimed at beginners, comprehensive materials will follow, detailing the structured PE process complemented by related materials and case studies.

We want to talk with you about your HPC applications and PE Problem!

FAQ

Sh-file

General

Usage

`$ bash myscript.sh`

Tools

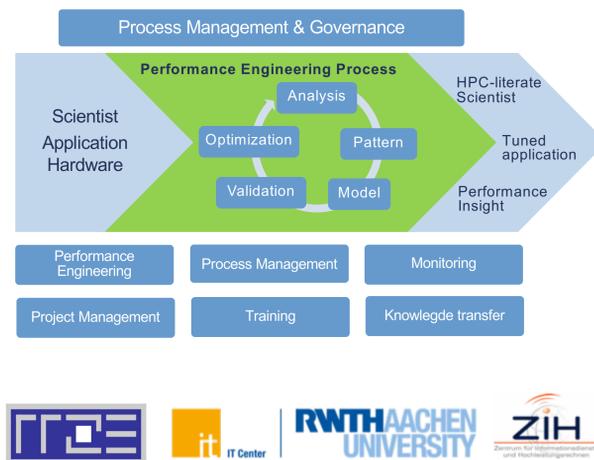
CASE STUDY

ProPE Project Structure

While the project does not have enough manpower to fully unroll all of the points it will create a **blueprint** and develop the necessary **tools** and **processes**.

Important ingredients will be **show-cases** by

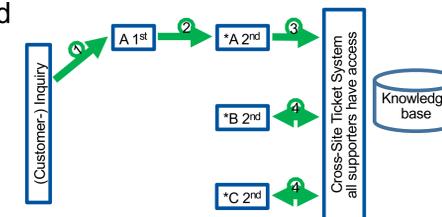
- Carrying out real **Performance Projects**
- **Organizing tutorials** and **researcher exchanges** between sites
- Establishing system-wide job specific **performance profiling** infrastructures



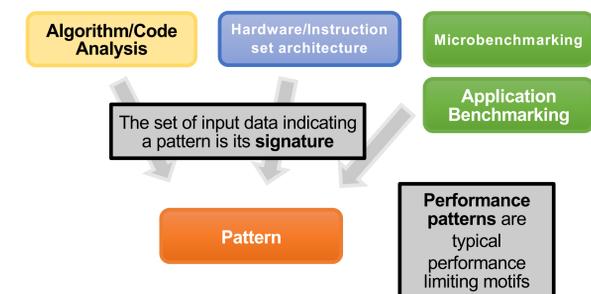
Distributed PE Support Infrastructure

Establish multi-tier distributed process management service and support structures that seamlessly integrate local and remote support levels to provide "in-depth" expertise distributed across sites.

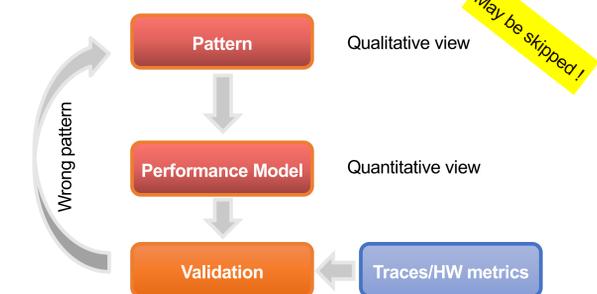
1. 1st level support of local site (A) receives customer request
2. Request is handed over to local 2nd level HPC support team
3. Request is transferred to cross-site ticket tool if additional competence is needed
4. Site (B and/or C) with expertise processes the ticket
5. 2nd level HPC support team processes ticket finalization



Structured PE-Process

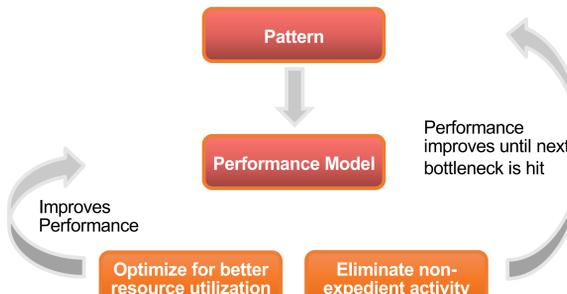


Step 1 **Analysis**: Understanding observed performance



Step 2 **Formulate Model**: Validate pattern and get quantitative insight

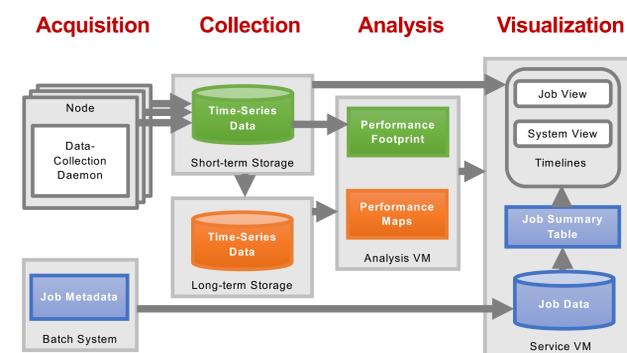
The core activity of analyzing and optimizing application performance is guided by a systematic PE-Process. At its core are typical performance limiting settings called performance patterns. To validate and to get a quantitative view of a pattern white box performance models are employed. Identifying a performance pattern is achieved by a set of hardware performance counter metrics but might also involve static code analysis and benchmarking results.



Step 3 **Optimization**: Improve utilization of available resources

Job Performance Monitoring

- Permanent light-weight monitoring
- Short-term and long-term data analysis
- Characterization of jobs and applications
- Detection of pathological performance behavior
- Live and post-mortem visualization



Integrated Job and System View

