Performance Optimization and Productivity

EU H2020 Center of Excellence (CoE)

1 October 2015 – 31 March 2018 (30 months)

POP CoE

- **A Center of Excellence**
  - **On Performance Optimization and Productivity**
  - **Promoting best practices in performance analysis and parallel programming**
- **Providing Services**
  - Precise understanding of application and system behavior
  - Suggestion/support on how to refactor code in the most productive way
- **Horizontal**
  - Transversal across application areas, platforms, scales
- **For academic AND industrial codes and users !**

Partners

- **Who?**
  - BSC (coordinator), ES
  - HLRS, DE
  - JSC, DE
  - NAG, UK
  - RWTH Aachen, IT Center, DE
  - TERATEC, FR

A team with

- Excellence in performance tools and tuning
- Excellence in programming models and practices
- Research and development background AND
  proven commitment in application to real academic and industrial use cases

Motivation

Why?

- Complexity of machines and codes
  → Frequent lack of quantified understanding of actual behavior
  → Not clear most productive direction of code refactoring
- Important to maximize efficiency (performance, power) of compute intensive applications and the productivity of the development efforts

Target

- Parallel programs , mainly MPI /OpenMP … although can also look at CUDA, OpenCL, Python, …

3 levels of services

? Application Performance Audit

- Primary service
- Identify performance issues of customer code (at customer site)
- Small Effort (< 1 month)

! Application Performance Plan

- Follow-up on the service
- Identifies the root causes of the issues found and qualifies and quantifies approaches to address the issues
- Longer effort (1-3 months)

✓ Proof-of-Concept

- Experiments and mock-up tests for customer codes
- Kernel extraction, parallelization, mini-apps experiments to show effect of proposed optimizations
- 6 months effort

Reports

Software demonstrator

Apply @
http://www.pop-coe.eu

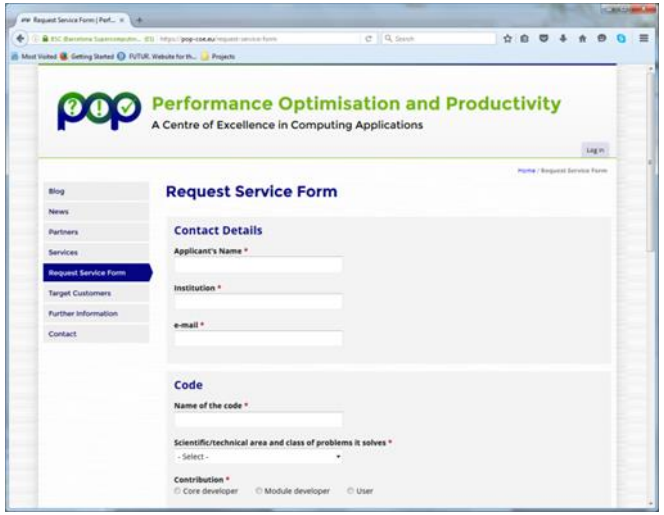

Target customers

- **Code developers**
  - Assessment of detailed actual behavior
  - Suggestion of more productive directions to refactor code
- **Users**
  - Assessment of achieved performance on specific production conditions
  - Possible improvements modifying environment setup
  - Evidences to interact with code provider
- **Infrastructure operators**
  - Assessment of achieved performance in production conditions
  - Possible improvements modifying environment setup
  - Information for allocation processes
  - Training of support staff
- **Vendors**
  - Benchmarking
  - Customer support
  - System dimensioning/design

Best practices in Performance analysis

- **Powerful tools …**
  - Extrae + Paraver
  - Score-P + Scalasca/TAU/Vampir + Cube
  - Dimemas, Extra-P
  - Other commercial tools
- **… and techniques**
  - Clustering, modeling, projection, extrapolation, memory access patterns, …
  - … with extreme detail …
  - … and up to extreme scale

- **Unify methodologies**
  - Structure
    - Spatio temporal / syntactic
  - Metrics
    - Parallel fundamental factors: Efficiency, Load balance, Serialization
    - Programming model related metrics
    - User level code sequential performance
  - Hierarchical search
    - From high level fundamental behavior to its causes
- **To deliver insight**
- **To estimate potentials**

Best practices in parallel programming

- **MPI and OpenMP**
  - Active members of OpenMP consortium (RWTH, BSC)
  - Active members of MPI Forum (JSC, RWTH)
- **Pushing application as early adopters and co-design drivers**
- **Promoting new features …**
- **… gathering feedback**

- **Promoting a throughput oriented methodology**
  - Task based programming
  - Asynchrony, overlap
  - Locality
  - Malleability, Dynamic Load Balancing
  - Nesting, recursion

Activities

- **External access**
  - WEB (www.pop-coe.eu)
    - Request form
    - Feedback questionnaires
    - News and blog
- **Internal organization**
  - CRM
  - TRAC ticketing system
  - Wiki

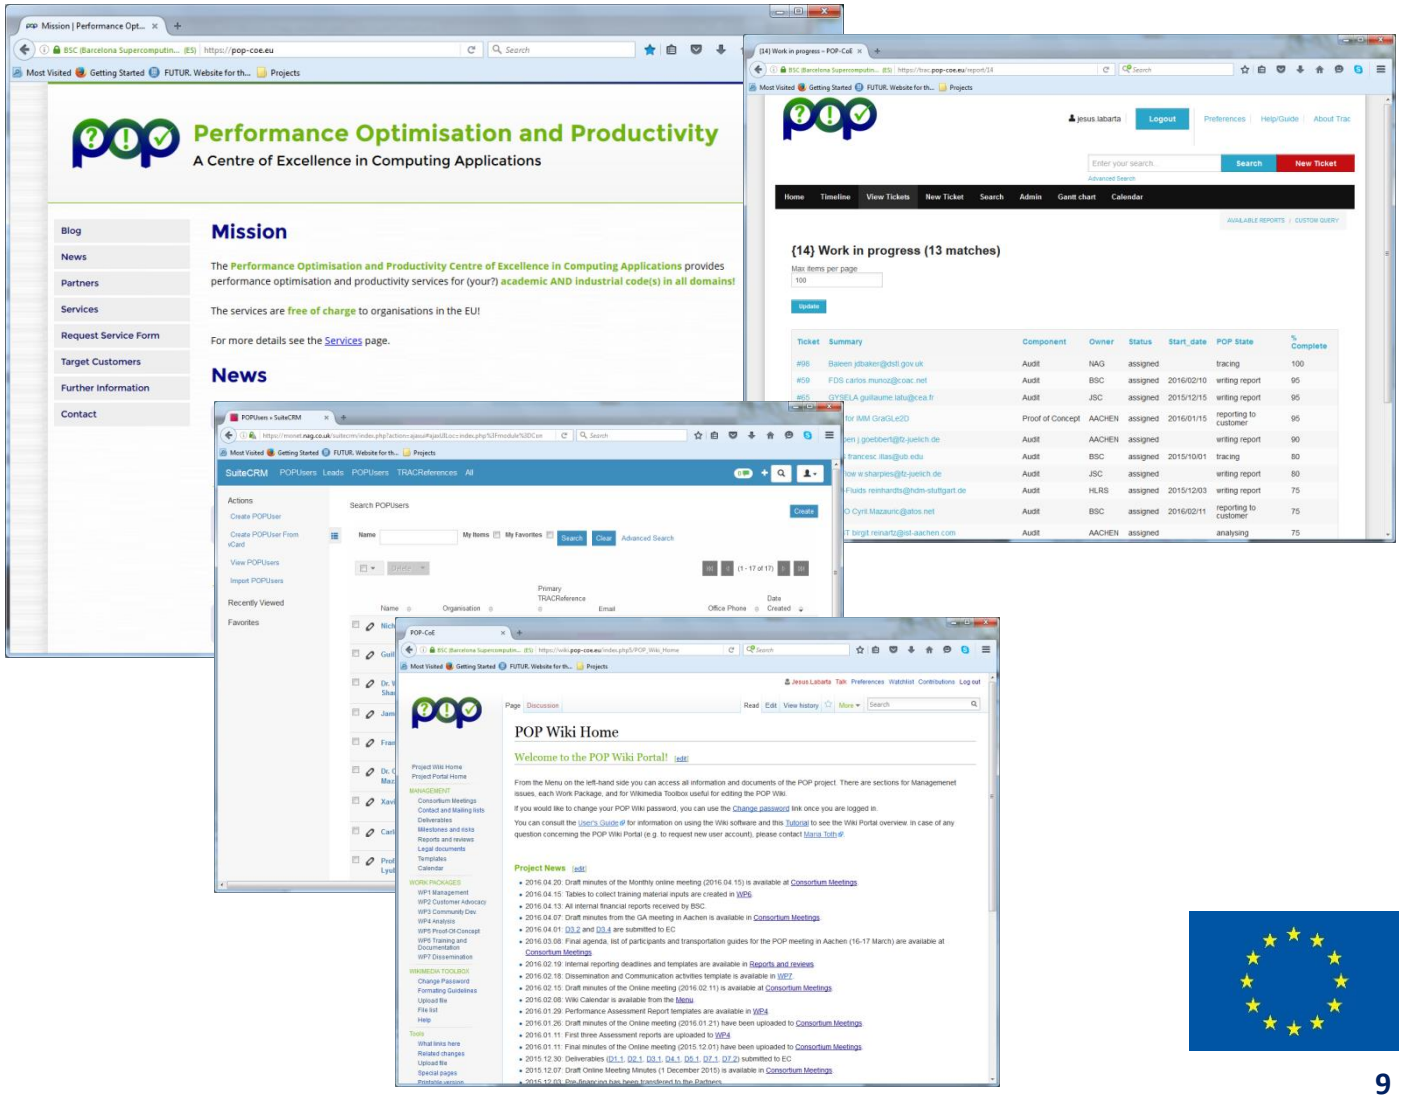

Activities

- **Services**
  - Completed/reporting: 15
  - Codes being analyzed: 6
  - Waiting user input. 13
  - Cancelled: 1

- **By type**
  - Audits: 31
  - Plan: 2
  - Proof of concept: 2

- **Reports**
  - 5 -15 pages

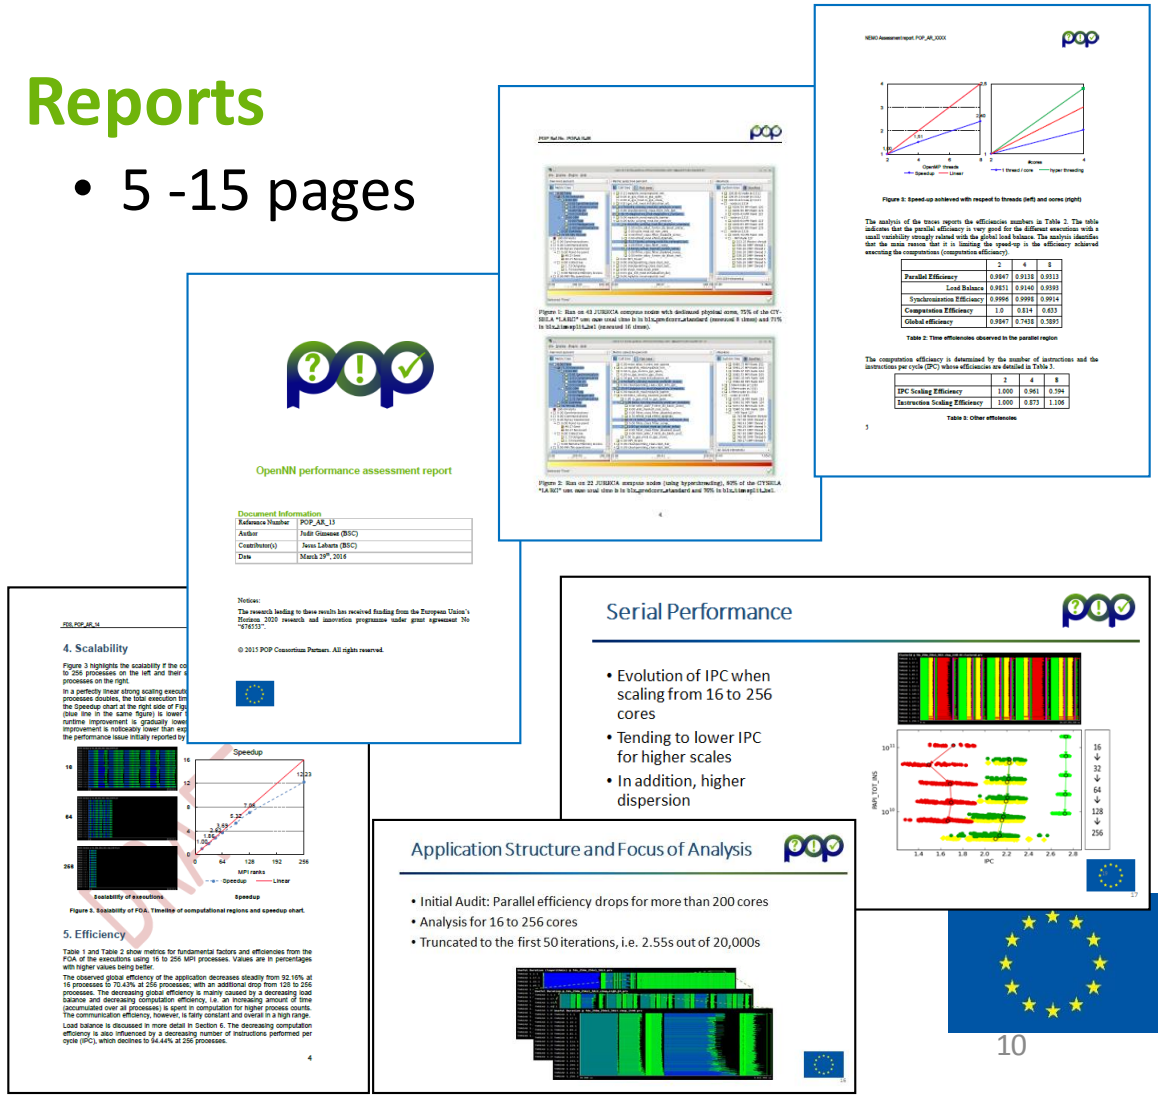

Other activities

- **Promotion and dissemination**
  - Market and community development
  - Dissemination material and events
- **Customer advocacy**
  - Gather customers feedback, ensure satisfaction, steer activities
- **Sustainability**
  - Explore business models
- **Training**
  - Best practices on the use of the tools and programming models (MPI + OpenMP)
    - Lot of interest … customers want to learn how to do it themselves

Answer to Questions

- **Presented what we offer, what we are doing**
- **Requests by EC**
  - Women participation: BSC: 4/6, HLRS: 1/3, RWTH: 1/3, NAG: 1/3, JSC: 0/3
  - Interaction between CoEs: Training EoCoE, Events (EsiWACE, …). Assessments to other CoEs
- **International cooperation**
  - we do have many activities (JLESC, VI-HPS,…) as individual partners.
  - Not at project level
- **PRACE scientific case, SRA, other FETHPC projects**
  - Involved on SRA at individual partners level
  - Have customers from other CoE and FETHPC projects. Our tools technologies are also used and partially developed in them

Conclusion

- We have established our internal operation infrastructure and procedures
- We have already performed 15 assessments and 19 are in the pipeline

- We believe the effort to unify our methodologies can become the core of best practices in performance analysis and programming practices that expand at international level

- We consider the POP CoE is progressing at a fairly good pace and results are already showing up.

www.pop-coe.eu

THANKS