



## D6.3 Tutorials Version 1.0

### Document Information

<b>Contract Number</b>	676553
<b>Project Website</b>	<a href="http://www.pop-coe.eu">www.pop-coe.eu</a>
<b>Contractual Deadline</b>	M30, March 2018
<b>Dissemination Level</b>	Public
<b>Nature</b>	DEC
<b>Author</b>	Christian Terboven (RWTH)
<b>Contributor(s)</b>	
<b>Reviewer</b>	Brian Wylie (JSC)
<b>Keywords</b>	Training, Tutorials



**Notices:**

*The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No "676553".*

© 2015 POP Consortium Partners. All rights reserved.



## Change Log

Version	Author	Description of Change
V0.1	Christian Terboven	First version for internal review
V0.2	Brian Wylie	Review revisions
V1.0	Christian Terboven	Further revisions
		<i>(Final Change Log entries reserved for releases to the EC)</i>



## Table of Contents

<b>Executive Summary</b> .....	<b>4</b>
<b>1. Introduction</b> .....	<b>5</b>
<b>2. Tutorials</b> .....	<b>5</b>
2.1 POP customer events.....	5
2.2 Tutorials at key HPC conferences .....	7
<b>3. Webinars</b> .....	<b>8</b>
<b>4. Results and impacts</b> .....	<b>10</b>
<b>Acronyms and Abbreviations</b> .....	<b>11</b>



## Executive Summary

This deliverable is part of work package 6 titled “Training and Documentation”. The POP work plan promised to organize at least five tutorials at customer sites or at key HPC conferences during the project’s lifetime. This report describes the structure of the events that were carried out as part of POP, summarizes the experiences and participant feedback from these events, and gives an overview about the additional training activities that went past the initial plan of WP6.

Six multi-day hands-on workshops were organized for application developers, serving a total of 86 participants and users of HPC facilities from several Centres of Excellence, as well as contributions to other training events and six tutorials at international conferences. These were complemented by a series of five webinars. Numerous POP service requests resulted from these events, as well as dissemination of use of the tools and performance analysis methodology of the POP Centre of Excellence.



# 1. Introduction

Work package 6 (WP6) with the title “Training and Documentation” is in charge of the general-purpose training activities supporting the outreach to POP customers and the applicability and sustainability of the results from performance audits, performance plans and proof of concepts. In summary, it provides customer-oriented events and documentation.

WP6 contained the selection of documentation material on the HPC performance tools and programming models employed in POP (see deliverable D6.1) and the creation of a description of the performance tuning workflow applied within POP (see deliverable D6.2). This final deliverable of WP6 reports on the training and tutorial events carried out by POP in the course of the 30-month project duration.

## 2. Tutorials

We differentiate between POP events targeted and customised for POP customers and tutorials at HPC conferences.

### 2.1 POP customer events

Task 6.3 from WP6 performed tutorials and workshop events for POP customers. Based on the customer’s needs and the availability of appropriate (local) HPC systems with necessary software, the form of the events varied.

#### POP-EoCoE Workshop: December 2015

The first EoCoE / POP workshop on benchmarking and performance analysis at Jülich Supercomputing Centre on December 8 to 11, 2015 brought together teams of code developers of the EoCoE Centre of Excellence application work packages with HPC experts associated with the EoCoE transversal basis work package and HPC experts from the POP partners BSC, JSC and RWTH. The goal was to familiarise the EoCoE developers with state-of-the-art HPC performance analysis tools, enabling the teams to make a preliminary identification of bottlenecks, and to initiate the standardisation of benchmark procedures for these codes within the EoCoE project. 13 EoCoE staff from 7 of their application code-teams participated, plus one from the MaX Centre of Excellence.

A very valuable outcome was the exchange of respective ideas and needs between code developers and HPC experts, as this helped clarifying the issues from either perspective and enabled both sides to interact more smoothly with a well-defined focus on the next actions to be taken. For example, the requirements for a full code 'audit' from the EoCoE and POP perspectives were clarified: here it was decided that the initial benchmarking would take place within and immediately after the workshop by EoCoE transversal basis members, whereas more in-depth follow-up analyses could



be passed on to POP at a later stage. POP Performance Audits of the Gysela, ParFlow and Shemat applications were initiated by this workshop.

#### POP-EoCoE Workshop: June 2016

The second joint EoCoE/POP workshop on benchmarking and performance analysis took place at Maison de la Simulation (Saclay - France) from 30th May to 2nd June 2016. 12 application developers were supported by POP experts from BSC and JSC with their performance analysis tasks.

The codes were integrated within the JUBE workflow environment and their execution performance analysed with tools from POP to establish a performance baseline and make a preliminary identification of bottlenecks.

#### POP Workshop: October 2016

POP partners BSC & JSC contributed two days of training as part of the Program Analysis and Tools Workshop from 25 to 27 October 2016 at the DKRZ (German Climate Computing Center) in Hamburg, Germany. During the workshop, performance analysis tools and POP's analysis methodology were presented, demonstrated using the ICON code from the ESiWACE Centre of Excellence. The main emphasis of this workshop was to present tools that help users of the new Mistral cluster to analyse and optimize their code for current and future CPU technology, for instance Intel's new KNL architecture. About 15 code developers participated in the workshop.

#### POP Tutorial: December 2016

On December 14 – 15, POP partners from BSC, JSC and RWTH gave a tutorial about performance analysis and optimization at IT4Innovations in Ostrava, Czech Republic. During the tutorial, performance analysis tools and POP's analysis methodology were presented, demonstrated with an audit of the local k-Wave code. Furthermore, the tutorial was rounded off by long hands-on sessions where code developers and POP experts worked on the analysis of several applications codes from workshop participants. About 15 code developers from the Czech Republic and Poland participated in the workshop.

#### POP-PATC-VIHPS Workshop: February 2017

From 8 to 10 February, POP provided instructors from BSC, JSC and NAG for the 24th VI-HPS Tuning Workshop which was held at the University of Southampton, England, co-organised by POP and the UK PRACE Advanced Training Centre at EPCC. 16 attendees participated, mostly from the hosting university and the rest of England, bringing their MPI and OpenMP parallel applications to analyse on the ARCHER Cray XC30 at EPCC with the performance tools from BSC, JSC and Alinea/ARM. Following initial demonstrations and hands-on exercises with prepared examples to familiarise participants with using the tools, they received assistance and advice for their own application codes, concluding with an overview of POP services. The event facilitates the review of existing audits and resulted in new audit requests.



## POP-EoCoE Workshop: April 2017

In a joint effort, the two centres of excellence EoCoE and POP have once again hold a hands-on workshop on HPC benchmarking and performance analysis on April 24 to 26, 2017. It was the third event of its kind and has been held at BSC in Barcelona supported by the French and the Spanish PATCs. Improving again an already proven concept, it has brought together 17 experts from topical fields in energy research and tools and 11 experts from HPC science including POP tool experts from BSC and JSC in order to tackle the transition of current R&D codes and applications towards Exascale. In particular, two application teams from the E-CAM Centre of Excellence joined the workshop, resulting in a POP Performance Audit of CP2K.

Experts from topical fields could really learn how to use advanced performance evaluation tools, get insight of the performance bottlenecks of their applications and take home the JUBE-based benchmarking tool configurations to repeat, in a reproducible manner, this analysis on future optimised versions of their code.

Without explicitly being listed here, POP staff supported the PATC/VI-HPS (Virtual Institute for High Productivity Supercomputing, [www.vi-hps.org](http://www.vi-hps.org)) events and bring-in-your-own-code activities at RWTH and TERATEC with staffing and material.

## 2.2 Tutorials at key HPC conferences

In order to gain international visibility for the POP training and performance analysis activities, we promised to contribute to selected tutorials at key HPC conferences with high visibility (i.e. ACM/IEEE SC, ISC). POP staff and other staff from POP member organizations contributed to the following tutorials.

No.	Year	Conference	Tutorials with POP contributions
1	2015	ACM/IEEE SC	<ul style="list-style-type: none"> <li>Advanced OpenMP: Performance and 4.1 Features</li> <li>Practical Hybrid Parallel Application Performance Engineering</li> <li>Debugging and Performance Tools for MPI and OpenMP 4.0 Applications for CPU and Accelerators/Coprocessors</li> <li>Efficient Parallel Debugging for MPI, Threads, and Beyond</li> </ul>
2	2016	ISC	<ul style="list-style-type: none"> <li>Practical Hybrid Parallel Application Performance Engineering</li> <li>Advanced OpenMP: Performance and 4.5 Features</li> </ul>
3	2016	ACM/IEEE SC	<ul style="list-style-type: none"> <li>Advanced OpenMP: Performance and 4.5 Features</li> <li>Practical Hybrid Parallel Application Performance Engineering</li> <li>Efficient Parallel Debugging for MPI, Threads, and Beyond</li> </ul>
4	2017	ISC	<ul style="list-style-type: none"> <li>Practical Hybrid Parallel Application Performance Engineering</li> <li>Advanced OpenMP: Performance and 4.5 Features</li> </ul>
5	2017	ACM/IEEE SC	<ul style="list-style-type: none"> <li>Advanced OpenMP: Performance and 4.5 Features</li> <li>Practical Hybrid Parallel Application Performance Engineering</li> </ul>
6	2018	Supercomputing Frontiers Europe	<ul style="list-style-type: none"> <li>Introduction to the Tools and Performance Analysis Methodology used in POP CoE</li> </ul>

**Table 1: Tutorials with POP contributions**



### 3. Webinars

A webinar is a live web-based video conference that uses the internet to connect the individual hosting the webinar to an audience – the viewers and listeners of the webinar from all over the world. Hosts can show themselves speaking, switch to their computer screens for slideshows or demonstrations, and even invite guests from other locations to co-host the webinar with them.

By the time of this writing, five POP webinars of 30 minutes plus questions have been carried out. The webinars were well-attended and have been recorded. The recordings have been made available on the POP project's homepage under <https://pop-coe.eu/blog/tags/webinar>. Furthermore, the recordings are available in the POP youtube channel under <https://www.youtube.com/channel/UCUdzljg5dfPd3X5cO7482ow>. The technical infrastructure to carry out and record the webinars was provided by NAG.

#### Webinar 1: How to improve the Performance of Parallel Codes (June 2017)

POP's first webinar tackled the subject of how to improve the performance of parallel codes. It presented a systematic approach to optimising these codes, while pointing out various factors that should be considered. As well as a live demo of performance tools from Barcelona Supercomputing Center, the talk was illustrated with practical examples from various POP performance assessments. Over 80 people attended the webinar, with questions being answered both during the session and afterwards by e-mail, on topics such as the profiling of Python codes and the accuracy of profiling data. The webinar was presented by Jon Gibson, a Technical Consultant at the Numerical Algorithms Group (NAG). He has 25 years' experience in scientific programming and has worked for two national HPC services in the UK.

#### Webinar 2: Getting Performance from OpenMP on NUMA (August 2017)

Most contemporary shared memory systems expose a non-uniform memory architecture (NUMA) and this has implications for application performance. However, the OpenMP programming model does not provide explicit support for that. This webinar discussed the approaches to getting the best performance from OpenMP applications on such machines. The webinar covered the characteristics of cc-NUMA architectures, the OpenMP Thread Affinity model and the operating system mechanisms of memory placement. It then explained how to use this understanding to achieve performance optimization. The talk included practical examples showing best practices. The webinar was presented by Dr. Christian Terboven, a Senior Scientist and the HPC Group Manager at RWTH Aachen University. Dr. Terboven has been involved in the analysis, tuning and parallelization of several large-scale simulation codes for various architectures. He is a member of the OpenMP Language Committee and leads the Affinity subcommittee.





### Webinar 3: Performance in 3 Numbers (October 2017)

Measuring application performance often results in a large amount of profile data or traces that are difficult to handle or interpret beyond some trivial first observations. These analyses often do not provide the kind of insight that would really help a code developer determine the most appropriate direction to follow in order to improve the code. This webinar introduced and explained the model being used in the POP Project that achieves high semantic description of an application efficiency with just three metrics at the first level. It also demonstrated how these metrics can be easily computed with the BSC tools and Scalasca toolset as well as how the top three metrics can be derived from a standard MPI profiler output. It concluded by showing metrics for some example applications and how different behavioural effects of a system or application are reflected in the metrics. This webinar was presented by Jesus Labarta, a full professor of Computer Architecture at the Technical University of Catalonia (UPC) since 1990. Since 2005, he has been responsible for the Computer Science Research Department within the Barcelona Supercomputing Center (BSC). His major directions of current work relate to performance analysis tools, programming models and resource management. His team distributes the Open Source BSC tools (Paraver and Dimemas) and performs research on increasing the intelligence embedded in the performance analysis tools.

### Webinar 4: OpenMP Tasking (December 2017)

With the increasing prevalence of multi-core processors, shared-memory programming models are essential. OpenMP is a popular, portable, widely supported and easy-to-use shared-memory model. Since version 3.0, released in 2008, OpenMP offers tasking to support the creation of composable parallel software blocks and the parallelization of irregular algorithms. However, the tasking concept requires a change in the way developers reason about the structure of their code and hence expose the parallelism of it. This webinar gave an overview of the OpenMP tasking language features and performance aspects, such as introducing cut-off mechanisms and exploiting task dependencies. A number of questions from attendees were addressed by the speakers at the end of the webinar to conclude what was an informative and interesting talk. The webinar was presented by Dr. Christian Terboven, a Senior Scientist and the HPC Group Manager at RWTH Aachen University, and Dr. Dirk Schmidl, a member of the HPC Group at RWTH Aachen University. Dr. Terboven has been involved in the analysis, tuning and parallelization of several large-scale simulation codes for various architectures. He is a member of the OpenMP Language Committee and leads the Affinity subcommittee. Dr. Schmidl has participated for several years on the OpenMP Language Committee to help standardize the OpenMP programming model.

### Webinar 5: Parallel I/O Performance with Darshan (February 2018)

The 5th POP webinar on I/O profiling using Darshan was presented which had over 80 attendees joining. The webinar started with an introduction to parallel file and storage systems, followed by discussion of five different I/O models available for MPI codes. A live demo of the Darshan tool was presented which



included a discussion on favourable and unfavourable performance characteristics to observe in the Darshan report. The webinar concluded with programming hints and tips on writing efficient file I/O. A lively series of questions from the attendees included the performance of I/O libraries, file formats, POSIX I/O and how the recently reported CPU bugs will affect I/O workloads. The webinar also listed some references which include more in-depth information on parallel file I/O. The webinar was given by Wadud Miah, employed at NAG and working on the Performance Optimisation and Productivity (POP) project. He has worked in multiple roles in HPC, helping researchers improve their productivity.

During the registration process of the 5<sup>th</sup> webinar, we recorded gender ratios. That one was attended by 12 women, 64 men and 47 who prefer to not say (total of 123 registrations).

## 4. Results and impacts

Summarizing the explicit POP customer events, we performed 6 events with a total of 86 participants (10 female). About 50% of those can be considered as young researchers or students.

Along with the partner events with the EoCoE CoE, we served participants from the ESiWACE, MaX and E-CAM CoEs.

The training events supported the acquisition of POP customers, resulting in POP service requests.



## Acronyms and Abbreviations

- BSC – Barcelona Supercomputing Center
- D – deliverable
- EC – European Commission
- HLRS – High Performance Computing Centre (University of Stuttgart)
- HPC – High Performance Computing
- Juelich – Forschungszentrum Juelich GmbH
- M – Month
- MS – Milestones
- PM – Person month / Project manager
- POP – Performance Optimization and Productivity
- RWTH Aachen – Rheinisch-Westfaelische Technische Hochschule Aachen
- USTUTT (HLRS) – University of Stuttgart
- WP – Work Package