D2.4 Final Customer Feedback Report
Version 1.1

Document Information

Contract Number	676553
Project Website	www.pop-coe.eu
Contractual Deadline	M30, March 2018
Dissemination Level	PU
Nature	R
Author	JM Morel (Teratec)
Contributor(s)	D. Loureiro (INRIA)
Reviewer	Cristina Calatayud (BSC)
Keywords	Customer Advocacy, Customer Feedback, User Forum

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".

Change Log

Version	Author	Description of Change
V1.0	JM Morel	Initial Version
V1.1	JM Morel	Revised version taking into account Cristina's remarks

Table of Contents

Executive Summary

This deliverable summarizes the findings of the customer advocate collected by both the questionnaires filled in by the customers and the interviews we conducted with several customers. It summarizes also the corresponding recommendations which have been presented to the project management all along the project. Most of them have already been implemented but should anyway be carefully considered in the operation of the Center of Excellence after the end of the project in order to maximize customer satisfaction.

1. Introduction

As stated in the proposal *"The Customer Advocacy is a way of ensuring that the activities of the project are really performed to the full satisfaction of the customers. POP will deliver value to application developers, infrastructure operators and scientific users of parallel applications. […] The Customer Advocacy will gather their feedback and make sure they play an important role in driving the operation of the Centre of Excellence. To this end, The Customer Advocacy will carry out a process for measuring Customer Satisfaction and will organise events where customers can freely express their feedback and suggestions."*

As mentioned in D2.2 and D2.3, at the very beginning of the Project, we:

- designed various questionnaires in order to get feedback from the end users for each specific action carried out by POP experts (Audit of code, Performance Plan, Proof-of-Concept), about the Performance Analysis Tools and about the global ROI (Return on Investment measured as the ratio of the benefits resulting from the performance improvement of the application to the cost of the code optimization),
- designed various scripts of interview (according to the type of service provided to the customer) and organized to interview most end users by phone for getting more details about their experience of POP services, get their agreement on the interview report and identify customer quotes,
- and created a set of facilities to streamline the interaction with the users as well as a wiki table to follow-up the progress of these surveys, and share the results with other members of the Project (the POP experts who deliver the services, the people in charge of Dissemination, the Management, …).

Then, using this material we kept:

- following up the status of the POP services, soliciting end users to fill in appropriate questionnaires as soon as the service they have asked for

was finished (i.e. once they have received from the POP expert the corresponding Audit Report, Performance Plan, or PoC Findings),

- interviewing customers by phone to better understand how much they appreciated the POP services, to know about their intention to apply for complementary services, and to ask them if they would accept to pay for such services in the future,
- compiling and synthesizing the feedback from the questionnaires and from the interviews, uploading this information on the POP Wiki, and reporting regularly to the Operational Management meetings (monthly audio conferences, General Assemblies, …).

2. Customer feedback

2.1 Customer feedback through surveys

Each time an Audit, a Performance Plan or a Proof-of-Concept is finished, the corresponding survey is sent to the end user who benefited from this POP service, and once the survey is filled in, the compilation of the answers is sent back to the end user, and in most cases an interview by phone is organized to get more details.

The answers to the survey and the interview minutes are also sent to the corresponding POP experts in order to give them some insights on how the customer perceived the service that he or she has been provided.

Overall, we got a pretty good return rate of the completed surveys (~80%) regarding POP services (81/101 Audits; 12/15 PP; 10/16 PoC). This high response rate have occurred partly because these questionnaires are quite straightforward to complete and because nonresponding end users are systematically and rapidly called back.

As of today (March 26, 2018), in total, 124 answers have been gathered: 81 regarding audit reports, 12 for Performance Plan, 10 for Proof-of-Concept, 13 for the Performance Tools Evaluation, and 8 for the assessment of Performance improvement.

The questionnaire regarding the Performance Tools has only been sent to customers who reported they had used the Performance Tools by themselves (14%) or in cooperation with the POP expert (25%) to produce the traces.

Since the questionnaire regarding the assessment of performance increase and the actual gains resulting from code improvement versus the cost of effort spent (i.e. the Return on Investment) can only be answered by the customer once the code has been refactored and the various benefits assessed and measured. Therefore, it is difficult to know when the customer is able to answer this questionnaire: the time to modify the code may vary from a few days to several months according to the complexity of the code modification to

be carried out, and the developer often has higher priorities, while some of them report that the POP expert's recommendations will be taken into account only when developing a new version of their application). Moreover, the figures required (e.g. the actual value – in euros - of the gain obtained by being able to handle larger simulations is not easy to measure). However, we could get interesting data and several customers who did not answer the questionnaire were proud to give us interesting figures when they were interviewed.

The respective data summaries for the Analysis Report, the Performance Plan and the Proof-of-Concept can be found in the annex 1, 2 and 3 respectively.

If we take a look at the answers related to the general perception of POP experts and their reports for the Audit service, the feedback is very good:

- A vast majority (90%) of the end users found the Audit Report Clear and easy to understand (see Figure 1):

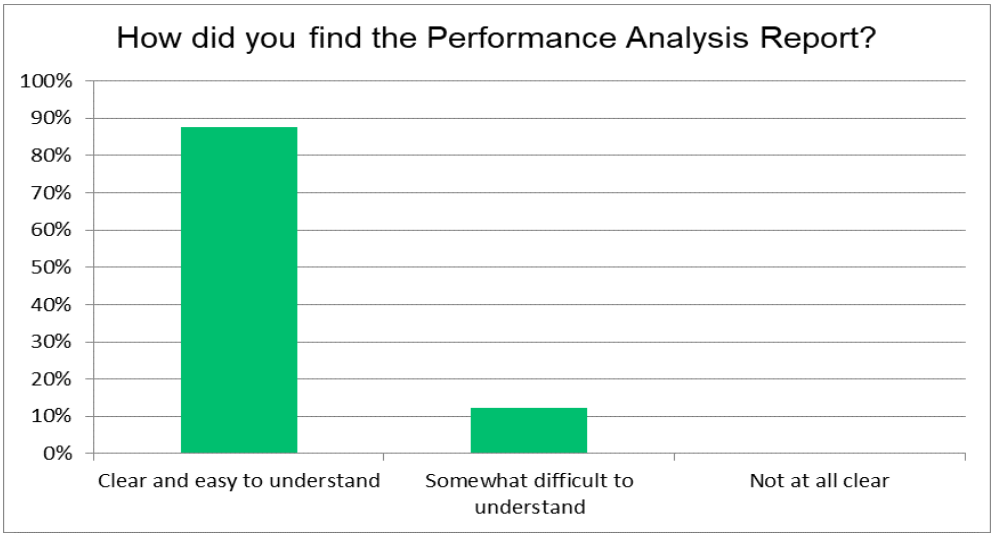

Figure 1 Readability and intelligibility of the report

- The end users found that the POP experts were found very or extremely responsive when questioned about the Audit (see **Erreur ! Source du renvoi introuvable.**)

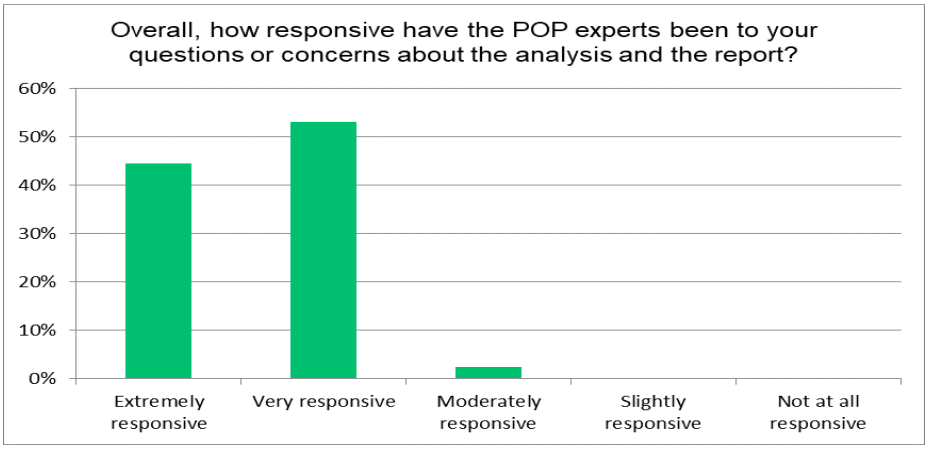

Figure 2 POP experts' responsiveness during Performance Audit

- The customers found that the answers of the POP experts to their questions were from good to excellent (see **Erreur ! Source du renvoi introuvable.**)

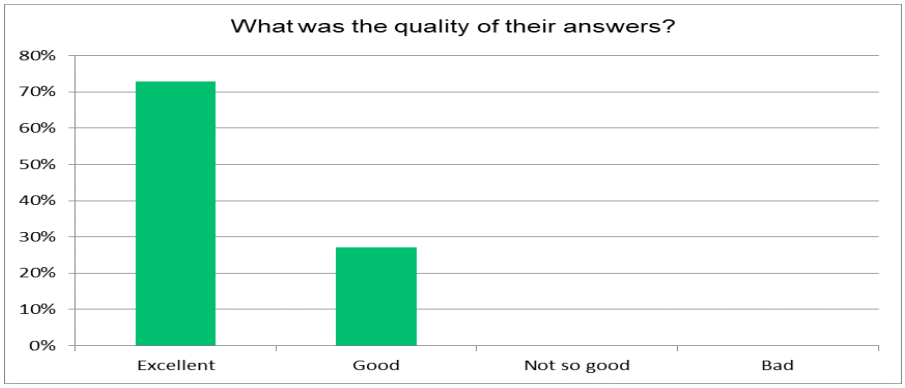

Figure 3 Quality of POP experts' answers

- Among the last 50 customers who filled in the questionnaire regarding the Performance Audit, 21 declared having enough information, while 20 declared they would like to get a Performance Plan, and 7 a PoC to get more precise guidance for enhancing their code.
- Also, one out of four customers declares to have other applications to audit. This means other potential requests, thereby a continuation of the POP Center of Excellence
- Moreover, about one third of them declare to be interested by training on performance analysis methods and tools, and on parallel code optimization.
- However, most customers declare that they are not sure that their organisation is ready to pay (see figure 4): some of them because they don't want to commit on behalf of their organisation, while others consider that the Performance Analysis report, which so far only describes the performance problems, should also give precise recommendations for code refactoring (something that requires much more efforts from the POP expert)

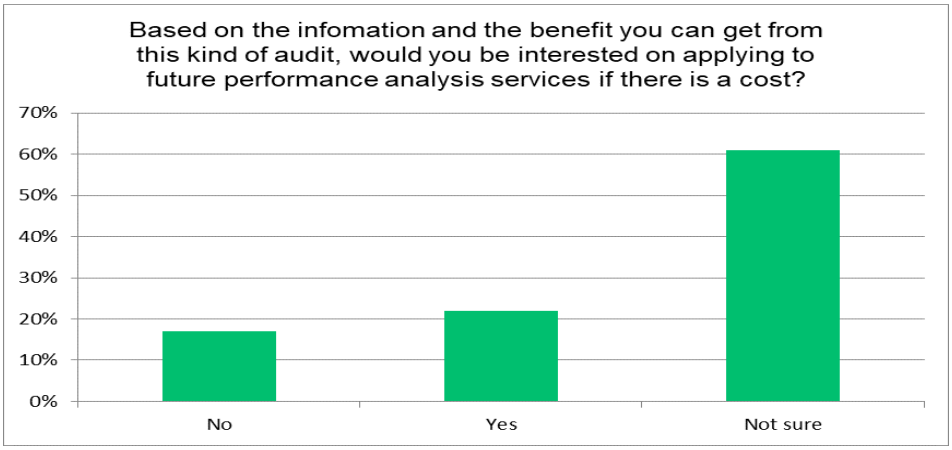

Figure 4 Readiness to pay for Performance Analysis services

The feedback regarding the **Performance Plan** service is also very encouraging:

- Most end users who benefited from a **Performance Plan** found it 'very useful' (17%) or 'useful' (67%) and are 'Very satisfied' (42%) or 'Satisfied' (42%).
- Most of them (75%) plan to implement the recommended modifications and all of them are ready to apply for similar service on other applications (see **Erreur ! Source du renvoi introuvable.**).

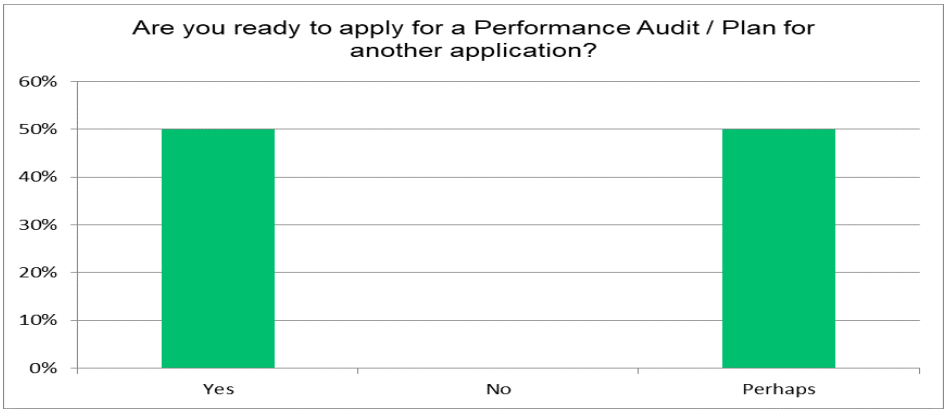

Figure 5 Readiness to apply for other Performance Analysis services

- However, while recognizing the value of such services, few of them declare to be disposed to pay. When they are asked about that during the interview, most of them declare they don't want to commit to pay on behalf of their organization.

The results are quite similar regarding the **Proof-Of-Concept services (PoC):**

- 90% of the customers who benefited from a PoC found it 'Very useful' (60%) or 'Useful' (30%).
- 100% consider that the POP expert was quite responsive (see Figure 6) and found their answers 'Very useful' or 'Useful'.

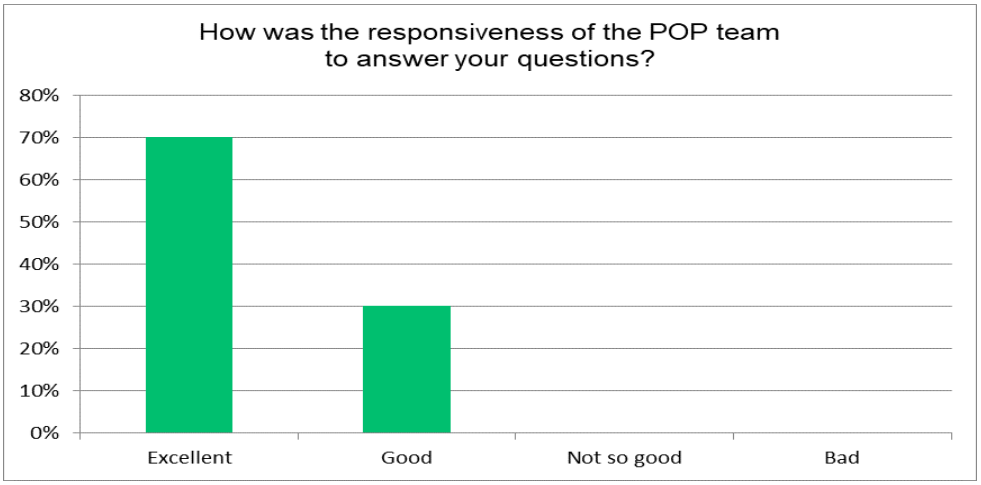

Figure 6 POP expert's responsiveness during PoC

- 90% did plan to implement the recommended code modifications, while all of them are 'Very satisfied' or 'Satisfied' with the set of services POP offered them:

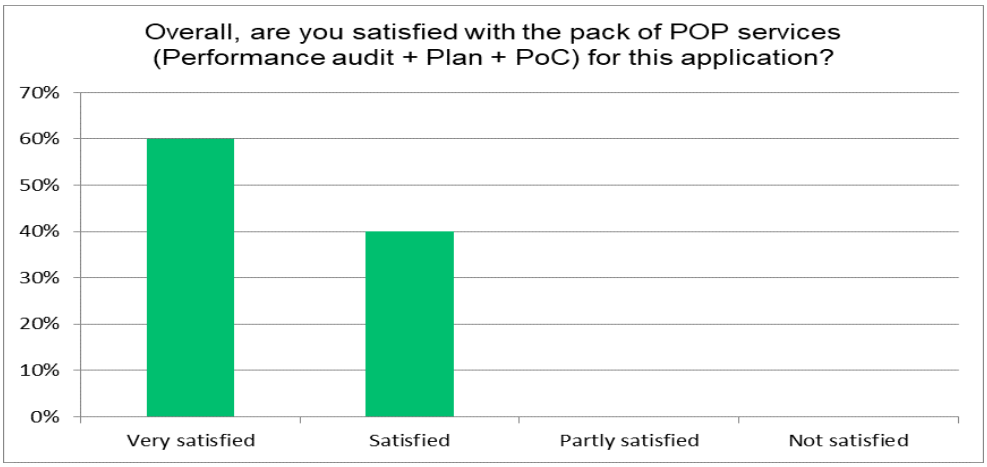

Figure 7 Overall POP customer satisfaction

- Finally, a large majority (70%) would be ready or possibly ready to pay for such services, but the price should be very low (see Figure 8).

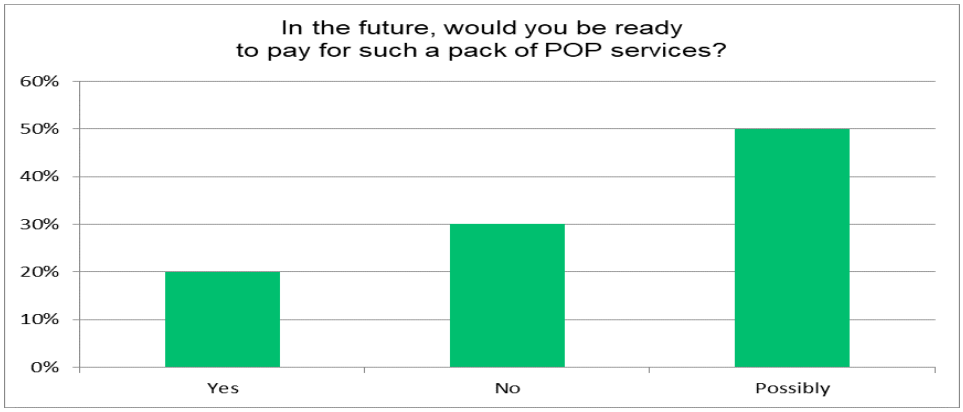

Figure 7 Readiness to pay of customers who benefited from a PoC

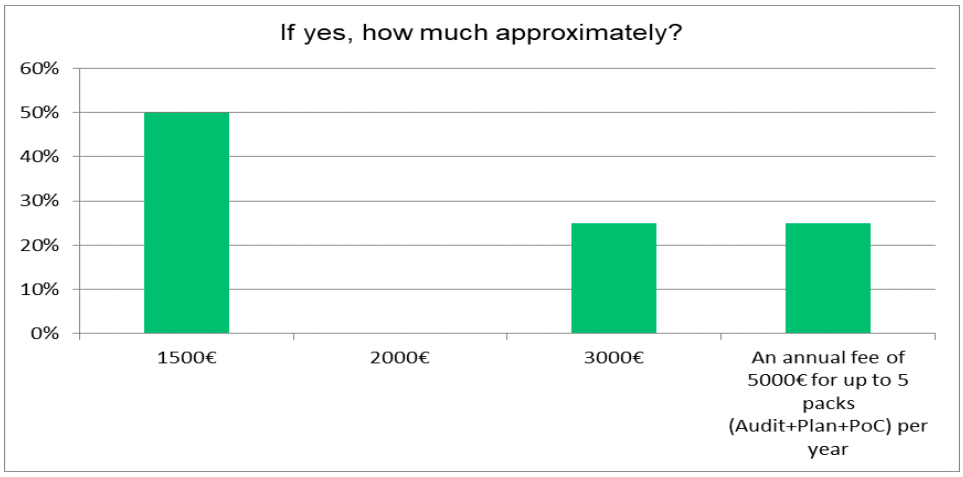

Figure 8 The price the POP customers would be ready to pay

The feedback collected via the questionnaires also includes interesting remarks like the following ones:

> "*The Report is well structured and showed me different points where to optimize my code. Some of them were clear to me but others were quite surprising, but helped me a lot to further improve the performance of our application".*
>
> *"The major impact of the audit was that it allowed us to validate the chosen implementation by giving us confidence on the choices we made in designing the code."*
>
> *"It is a good opportunity to get an impartial and objective point of view on the code when arriving."*
>
> *"Having the occasion to take some time with the expert to walk through the content of the report was great, and the team did learn a lot from the POP expert during the audit."*

Interestingly, one out of three customers did collect the traces by themselves or with the POP expert, and two users out of three ask for training in performance analysis. This clearly shows their will to better understand how to monitor the performance of their application as they evolve their code or have to run it on new architectures.

Feedback received from customers who used the **Performance Analysis Tools** shows that:

- The tools used were the following:

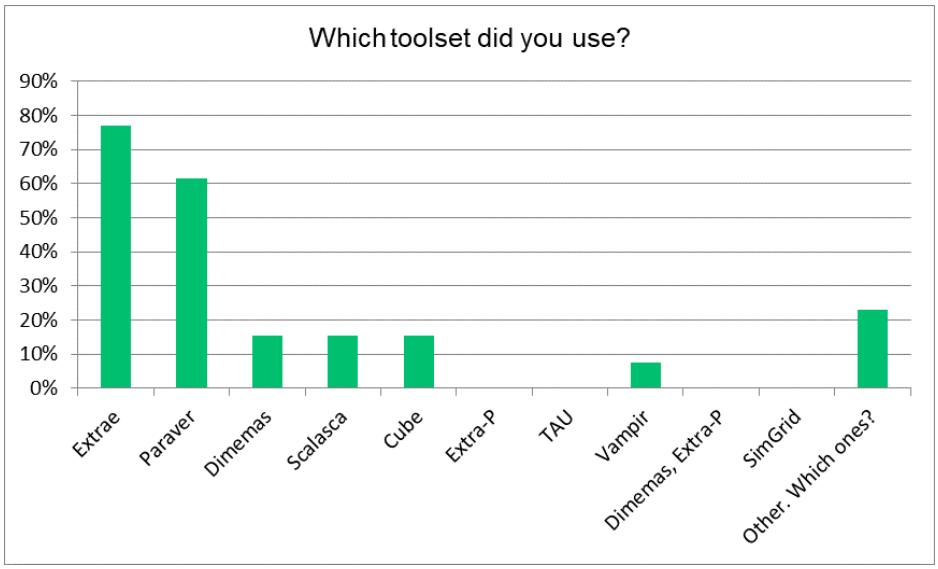

- These tools are not so easy to use:

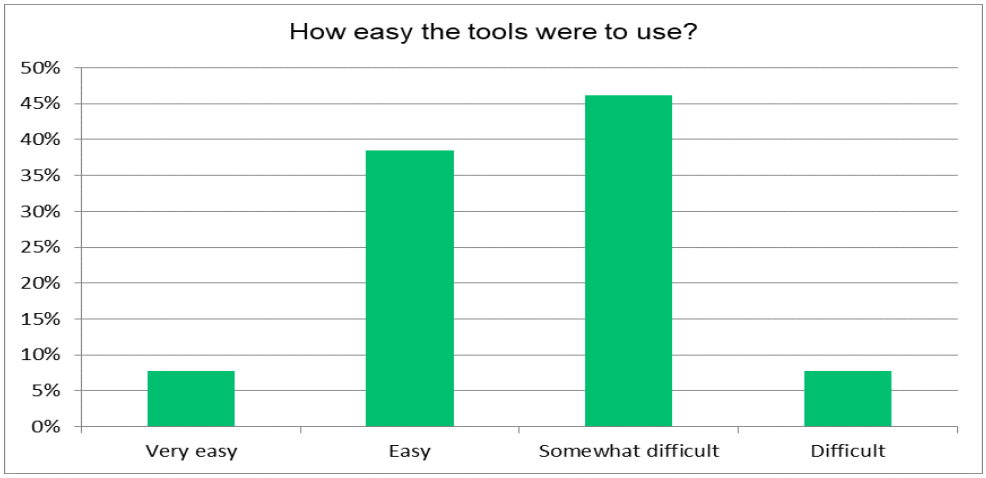

- But interpreting the output of these tools is quite difficult:

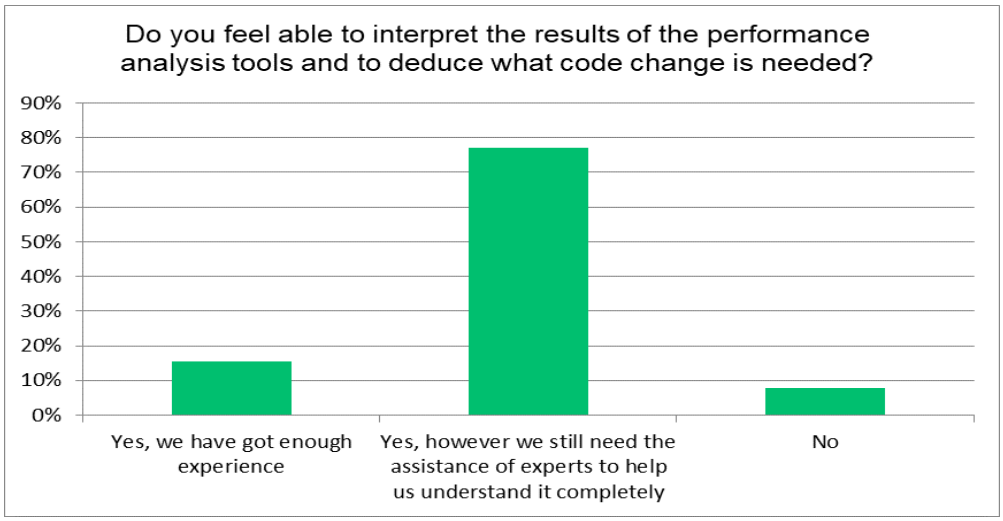

- However, users appear quite satisfied:

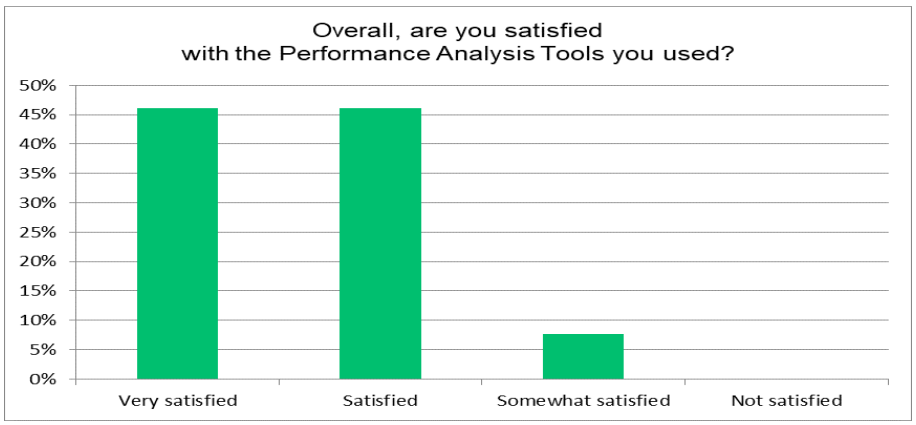

Last but not least, we got some interesting **feedback about the effort spent to modify the code and the resulting gains**:

- First it appears that the effort consumed to modify the code varies from a few days for to a few months:

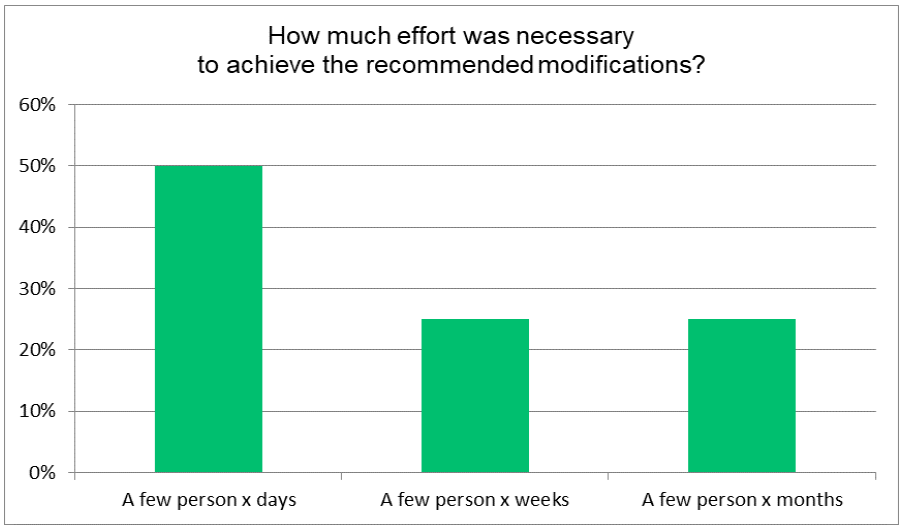

- But, there is no nasty surprises:

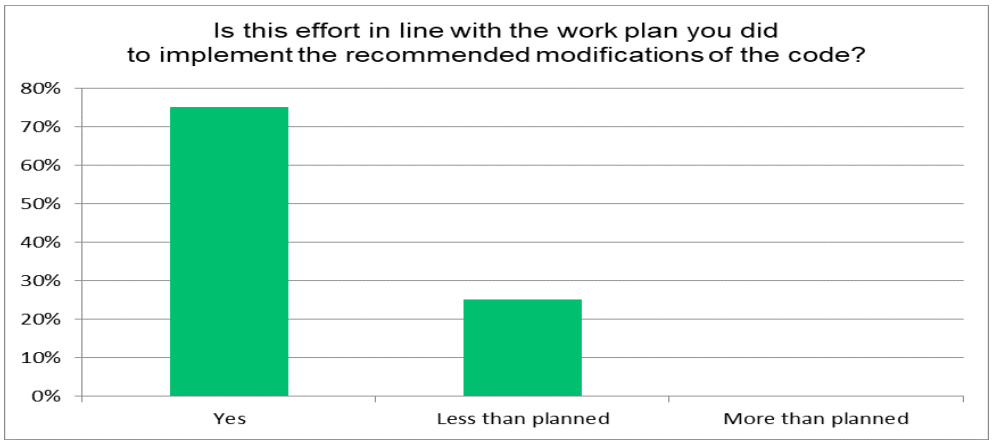

- One out of three customers adds functionality while refactoring its code:

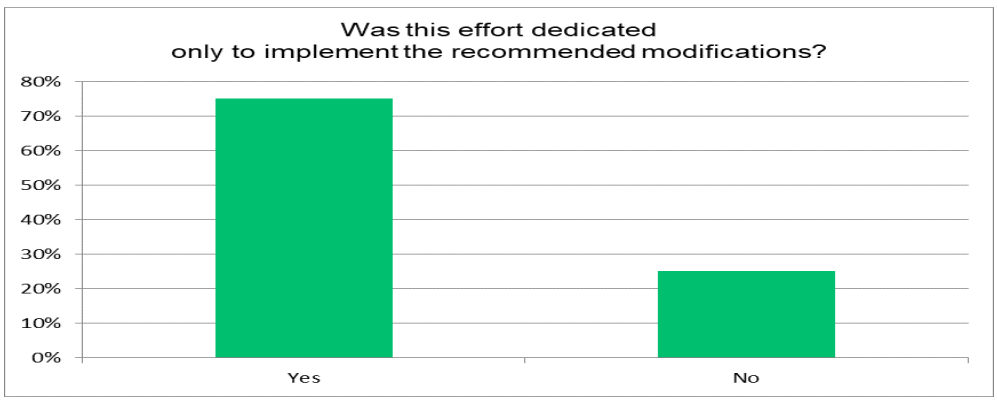

- Two out of three do use the Performance Tools to monitor the evolution of their code:

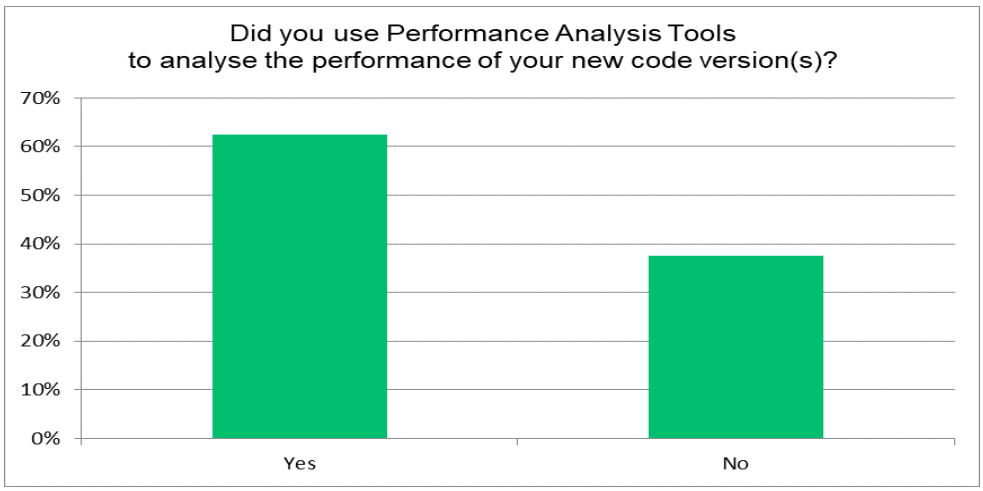

- The observed speed-up is interesting: the table below shows the speed-up reported by some of the customers who measured it.

• 25%
• 25%
• 20% overall, 50% for the given module
• 50-75% (case dependent)
• 12%
• Up to 62 %, depending on the use case.
• 6 - 47 % depending on the test case.
• 15%

- But there are other benefits: most customers realize that the code modifications not only improve the performance and the scalability of their application but also enable to treat larger problems and/or better exploit new architectures (mixing multi- and many-core servers):

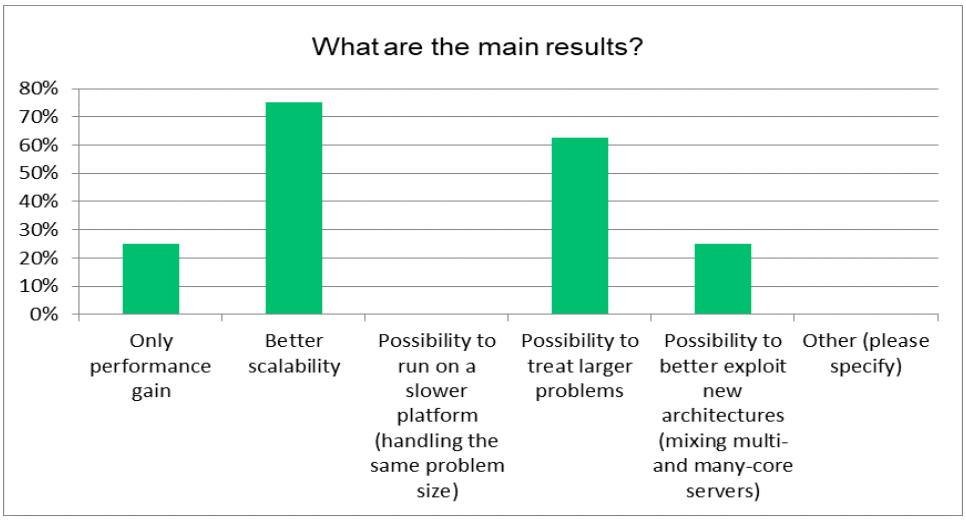

- And there are several sources of earned value:

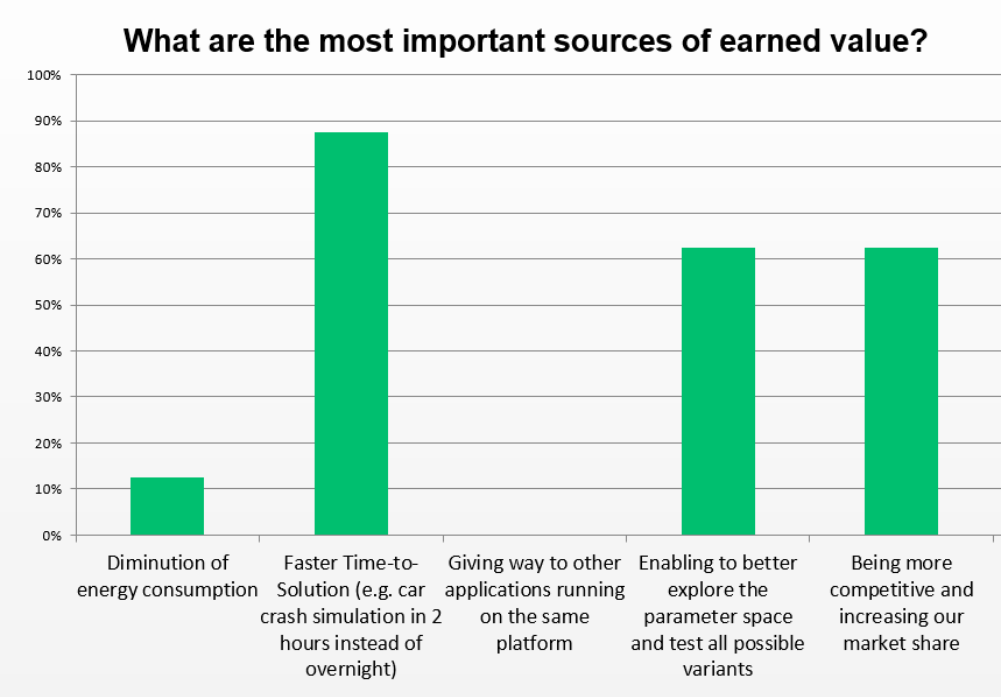

- Some customers did measure the Return on Investment by evaluating the gains they expect over one year (based on the gains they could measure over a few weeks) divided by the costs they spent to refactor their code:

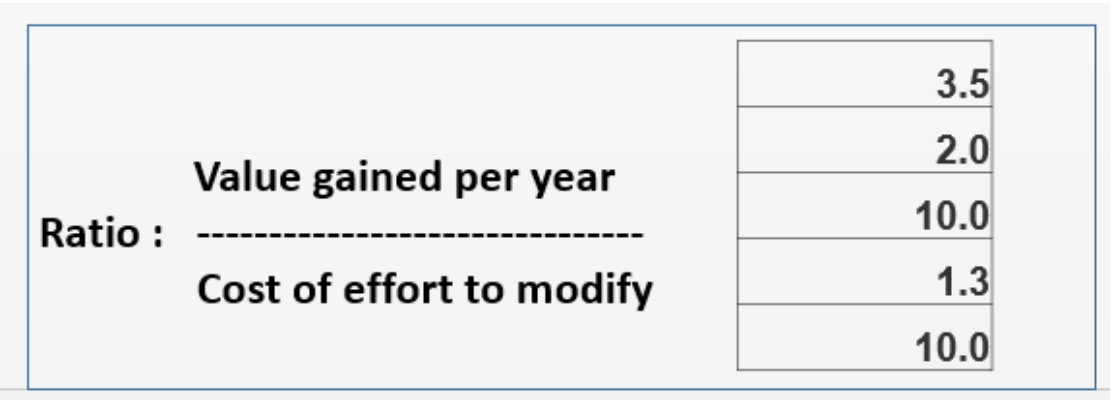

- In addition, they often mention that the application they have optimized is used by other people, e.g.:
  - *This application is run on hundreds or even thousands computers worldwide*
  - *We just have a few external users, only a part of them will benefit from those optimizations.*
  - *This open-source code is possibly used by other organizations, at least 3 of them.*
  - *This code will be the main code in several future projects.*

Complementary feedback on the gains obtained versus the effort spent has been collected through interviews and is described in the following section.

2.2 Customer feedback through Interviews

In order to get a deeper understanding of the customer experience, several scripts for interviews have been written, each one dedicated to a type of service (Audit, Performance Plan and Proof of Concept)

52 services have been the subject of interviews. All the customers have been very receptive to the action (only three of them refused to get interviewed).

Those interviews have been designed to get more feedback from the customers regarding the context of the POP services, how the end users knew about POP or any advice they could provide regarding the improvement of the POP services.

After each interview, the minutes of the interview are sent to the customer for validation and the final version is uploaded on the POP wiki, so that experts and partners in charge of communication (WP7) can build on this material.

Part of the information collected via surveys and interviews is registered in the WP2 page of the wiki to help follow-up the customers, and possibly propose them new services. In particular:

i) do they ask for some follow up (Performance Plan or Proof of Concept),
ii) do they ask for training (and the corresponding subject)
iii) do they have other codes to audit.

Through these interviews, it appears that:

- While first customers were often people directly or indirectly known to one POP member, we progressively got customers who discovered POP thanks to the efficient communication (mailings, newsletters, contacts in conference, POP web site,…) performed by the POP community development team (WP3).
- Almost all analyses were performed in the same context as the production one or at least in a similar context, while the input data were real production data for a vast majority or representative for others.
- All customers found that the POP services were performed in a reasonable time.
- All customers were pretty happy with the report they got at the end of the performance analysis even if some of them had to ask to additional information in order to get a complete understanding of all the figures and graphs.
- The major advice given by some customers to POP was to improve the communication with several points of synchronization during the performance analysis (at the beginning, at least one during the analysis in order to get some feedback and one at the end) with Face-to-Face meeting, if possible. Some customers were also quite happy with their report because they had the possibility to review it during its writing, allowing them to get a crystal clear comprehension of its content but

also to get all the information they need. This advice has been forwarded to all POP experts and corrective actions have been done to improve this communication.

- Most customers allowed POP to publish the reports. Very few refused.
- All customers stated, with no exception, that they would recommend POP to other people.

Remarks and suggestions for improvement collected during recent interviews include:

- The report should include an introductory chapter with the test case description and some background information about previous activities related to code performance and optimization.
- The report should also include a chapter describing the input and outputs evaluated to perform the analysis in order to have a proper traceability of the analyses.
- The steady interaction between the customer and the POP experts during all the duration of the audits is a key for success.
- Also, because some sections of the report may not be fully clear, a follow-up conference call should be systematically organized to help the customer understand all measures.
- Many customers say that Audits are a good opportunity to learn a lot on performance analysis methods and tools. Therefore, they would be interested in training on the analysis and the interpretation of the traces produced. They would like to learn how to produce such a report.
- It would be useful to make the difference between an audit for production codes and one for codes under heavy development which may not have a concrete example to use for the performance analysis.
- Some customers say that the major impact of the audit was that it allowed them to validate the chosen implementation by giving them confidence on the choices they made in designing the code.
- Having the occasion to take some time with the expert to walk through the content of the report was great, and the team did learn a lot from the POP expert during the audit.
- POP experts should also audit the IO part of the code.
- POP experts should propose alternative implementations, different libraries for example.
- POP should propose various levels of detail for the reports depending on the code and the person who submitted it.

We also got some interesting feedback about the ratio gains/costs of effort to optimize the code via the interviews:

- Modifications took 2-3 days and led to decrease memory usage so that the code is no longer memory bound and scales much better, thereby enables to simulate much larger problems.
- Modifications took 2 weeks and led to a speed-up improvement of 7x!
- Just modifying the workflow (renumbering the nodes) took 2 days and yielded a 25% improvement. (But before they had a hard time compiling Extrae on their cluster.
- We thought the code was quite optimized but in fact there was some room for improvement. Modifications took 6 months and led to a doubling of the speed up!
- The improvement was clearly worth the time spent on the audit.

2.3 Customers' remarks and quotes

We have got a pretty good set of remarks and quotes from customers. The table below contains the most interesting ones.

POP_AR_3	The POP experts did a great job during the performance audit. They identified the main problems thereby allowing the development team to understand the reasons behind some of these bottlenecks
POP_AR_5 POP_PP_01 POP_PoCR_1	The POP services helped us improve the efficiency and the scalability of our codes and more specifically for the non-trivial cases (when a lot of options are used). The team behind the codes never had time to take a closer look at complex test cases. The fact that the POP team was available to help them on the subject was a very good opportunity.
POP_AR_6	Thanks to the audit, we are now able to do much larger simulation (before we were able to simulate 1 million particles and now more than 60 millions).
POP_AR_8	This audit allowed us to identify the problems of the code, where to optimize it and thus put forward the performances of the machines that will run this code.
POP_AR_9	The major impact could be an improvement of the performance by a factor of 2 of the computation time. If the performance analysis can reduce the computation time and the energy consumption, this would be interesting for the end users.
POP_AR_10	If the customer is not an expert, such an audit can be of great value.
POP_AR_13	The major impact of POP was that with more performance, we can be better than the competition when run on customer desktop computer. And with better performance we hope to sell more licenses.
POP_AR_18	This audit helped us to realize that our code was close to the optimality and it also helped prepare input models for a large simulation we have to perform in a project starting in the near future.
POP_AR_19	The audit helped us to reach more rapidly the point where we wanted to go.
POP_AR_20	Thanks to POP experts' knowhow, which is missing in our

POP_PoCR_3	organization, we could carry on the parallelization of our code.
POP_AR_21	The major impact *(of the POP Audit)*, was that we could get a rather detailed look at the code, something we don't often do.
POP_AR_23	The recommendations most confirmed our suspicions, but it was good to quantify performance inefficiencies.
POP_AR_26	Those tools can help us to see how the performance can be deteriorated and if it is a major problem to address. This type of report is thus a good way to highlight the problems to solve.
POP_AR_30	The balance of the time spent in each routine was not clear before the audit. And with the audit this has been clearly identified. The routines were most of the time is spent will probably be rewritten. They now know where the slowdown was coming from.
POP_AR_32	The major impact of the audit was a clear insight into the bottlenecks of the code, mainly memory bounds problems. This audit allowed our team to know where to focus the work. Without the audit the team may have guessed memory bounds problems but not where to make improvements.
POP_AR_33	The major impact of the audit lies in the fact that it gave us a lot of information we can directly use to improve the performance of the code. From an organization point of view, any improvement which is directly linked to cost and computing time reduction can lead to direct cost savings.
POP_AR_34	HPC (POP) experts became interesting as far as they could combine various HPC metrics, and not just one or two, in order to analyse the performance of a code
POP_AR_35	The audit gave (the customer's) team good insights into the weak spots of the code. The code developers already started integrating the recommended improvements.
POP_AR_37	Thanks to the audit, we were able to know the parts of the code that needed to be improved which allowed us to list where we have to concentrate our work.
POP_AR_39	The major impact of the audit was that we got the certitude that we were executing the code in an efficient way and the scalability threshold of the code. We now know how it will behave if we try to augment the size of the grid to simulate. It's an important point to know the size of the problem the code can handle.
POP_AR_48 POP_PoC_6	POP expert's answers were meticulous and helped me getting more insight into both OpenMP and CUDA. Our application speed up considerably, and we have shown that OpenMP is a valuable choice for our case, while CUDA is not suitable, despite having achieved some speed up even with the GPU-enabled implementation.
POP_AR_60	From previous tests we had carried out ourselves we knew about the relatively poor scaling with increasing numbers of MPI processes, and that a lot of time could be spent by some threads in MPI_Wait while other threads were still running. This study told us these same things, but then also gave us more information on the range of different

	compute times for each MPI thread, and highlighted that it was just a couple of MPI threads which held up the program (presenting the possibility that we could garner large reductions in compute time by applying fixes to just a few patches in the model domain).
POP_AR_77	We are really grateful with the POP team due to their help in both the installation of their tools and the good explanation about their analyses.
POP_AR_87	Both the report and our communications with the POP expert were smooth and clear. The report focused on the main parts we were interested in and gave a clear understanding of the status of the code. Overall, I and other developers of this package are very satisfied with the report. It confirmed our experience with the code, and added quantitative understanding that we hope to use as a benchmark while continuing development.
POP_AR_93 POP_PoC_11	This PoC was useful because of easily understandable summary of an in-depth analysis of application/communication timing.
POP_AR_72 POP_PoC_16	The PoC improved the load balancing which was a major issue for the application. This will enable to run the application on larger datasets with a larger number of cores thanks to better scalability with the aim to analyze thoroughly real data from our customers.

3. Suggestions to POP management

If we take a deeper look at the answers provided by the end users, some points could be improved:

- In the previous deliverable, we stated that more explanations in the reports were necessary (mainly regarding figures and graphs).

  ⇨ **Significant progress has been achieved.**

- A customer reported in a comment that “Analysis should be made for a larger number of threads / MPI cores. It is not clear whether conclusions drawn from a very small number of cores can be extended to real-world simulations with thousands of cores.”

  ⇨ **This has been solved by updating the request service questionnaire which now asks for a detailed description of both production runs and development test case configurations**

- In a broader way, users may be asking for various levels of analysis. It will depend on:
  - Whether the user is a developer or a user of the code
  - If various datasets must be used to check all the parts of the code that need to be covered by the analysis
  - If the codes are currently under heavy development or if it’s a production code

  - If scalability must be investigated beyond the number of cores currently available on the customer's platform.
  - If the customer is an HPC expert of not

  ⇨ **More precise requirements must now be specified at the beginning so that the POP expert can take into accounts specific needs.**

- Another customer comment showed that "More frequent and early exchange on arising questions regarding obstacles found in the code. Some issues could have been resolved easily by an earlier communication."

  ⇨ **POP experts have been asked to interact regularly with the end user, so that the customer keeps involved all along the performance analysis.**

- Some end users found that the performance analysis were a good opportunity to learn how to use the performance analysis tools. However, a fraction of the customers found them hard to install or to use. Those two feedbacks could be solved with training that could be provided to customers with audited code (some of them explicitly asked for it) in order allow them to replicate the analyses but also the reports.

  ⇨ **In POP2, some effort will be dedicated to enhance the installation and the usability of the performance Tools.**

- So far, POP experts only audited the computing part, and not the IO part of the code. There could also be some room for improvement there. Depending on how the IO works, the computing performance could drop drastically.

  ⇨ **POP experts recently started using Darshan, a scalable HPC I/O characterization tool.**

4. Conclusions

During the period of the project, the number of services performed has greatly increased, and the Customer Advocate managed to solicit the feedback from all end users by sending systematically the appropriate survey and by carrying out many interviews.

This feedback was very good and showed that the users were satisfied with the services provided by POP. More and more often the end users who benefited from an Audit service choose to go forward with a Performance Plan and/or a Proof-of-Concept Service.

Of course, the collected feedback shows that there is still some room for improvement. We analysed the remarks and suggestions for improvement we got and passed them to the POP Management and to POP experts to help

them to be even more efficient by providing services that fully answer end users' needs in a cost effective way.

Moreover, we also incorporated all this feedback into the POP2 proposal to align the services provided with the end users' needs and requirements and guarantee that POP is useful for the software community and achieves a major impact.

Acronyms and Abbreviations

- AR – Audit Report
- BSC – Barcelona Supercomputing Center
- CA – Consortium Agreement
- CAdv – Customer Advocate
- DoA – Description of Action (Annex 1 of the Grant Agreement)
- EC – European Commission
- GA – General Assembly / Grant Agreement
- HLRS – High Performance Computing Centre (University of Stuttgart)
- HPC – High Performance Computing
- IPR – Intellectual Property Right
- Juelich – Forschungszentrum Juelich GmbH
- KPI – Key Performance Indicator
- MS – Milestones
- PEB – Project Executive Board
- PM – Person month / Project manager
- PoC – Proof-of-Concept
- POP – Performance Optimization and Productivity
- PP – Performance Plan
- RV – Review
- RWTH Aachen – Rheinisch-Westfaelische Technische Hochschule Aachen
- USTUTT (HLRS) – University of Stuttgart
- WP – Work Package
- WPL – Work Package Leader

List of Figures