D3.6 Market Review and POP 2 Business Plan Version 2.1

Document Information

Contract Number	676553
Project Website	www.pop-coe.eu
Contractual Deadline	M24, September 2017
Dissemination Level	PU
Nature	R
Author	Sally Bridgwater (NAG)
Contributor(s)	Mike Dewar, Nick Dingle, Jonathan Boyle, Jon Gibson, Wadud Miah (NAG)
Reviewer	Jean-Marc Morel
Keywords	Sustainability, commercial services, funding model

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
V0.1	Sally Bridgwater	Initial Draft
V1.0	Sally Bridgwater	Draft after consortium discussion
V1.1	Sally Bridgwater	Draft after Telco discussion
V2.0	JM Morel	Add review remarks
V2.1	Sally Bridgwater	Accepted review changes

Table of Contents

1 Executive summary 5

2 Updated Market Analysis 5

2.1 POP Customer Base 7
2.2 Current Marketing Approach 12
2.2.1 Events 12
2.2.2 Direct Marketing 12
2.2.3 Webinars 12

3 Lessons Learnt 12

3.1 The Big Picture 13
3.2 Tools 15
3.3 Industrial Customers 16
3.4 Services 17
3.4.1 Cost of Services 17
3.5 Action Plan 18

4 D3.6 POP 2 Business Plan 19

4.1 Mission 19
4.2 Key Expertise 19
4.3 Customers 19
4.4 Market 20
4.5 SWOT Analysis 20
4.6 Value Added and Competitive Advantage 20
4.7 Marketing Mix 23
4.7.1 Product 23
4.7.2 Price 24
4.7.3 Promotion 26
4.8 Work Volume 27
4.9 Revenues 27
4.10 Operational Strategy 27
4.11 Operational Goals 28

5 Sector categorisation 29

6 Acronyms and Abbreviations 30

Figure 1 POP assessments by Country 7
Figure 2 Share of POP assessments by country compared with HPC server revenue from 2016 7
Figure 3 Comparison of the leads entered into the CRM on 1/10/2016 and 4/8/2017 by country. 8
Figure 4 Number of assessments by sector 9
Figure 5 European server vertical markets in 2016, excluding the government and university markets. 9
Figure 6 3D plot showing the proportion of assessments that are classified as commercial, open source or inhouse by sector. 10
Figure 7 Organisation classification for POP studies 11
Figure 8 Categorisation of the usability of the performance analysis tools from data gathered by the POP partners NAG, HLRS and JSC. 15

Figure 9 Data from WP2 user questionnaire asking, "How easy were the tools to use?", based on 7 responses. 16
Figure 10 Data gathered from WP2 questionnaire asking, "What are the most important sources of earned value?", based on 6 responses. 22
Figure 11 WP2 data gathered by questionnaire asking, "What are the main results?", based on 6 responses. 22
Figure 12 WP2 data gathered by questionnaire asking, "Do you feel able to interpret the results of the performance analysis tools and to deduce what code change is needed?", based on 7 responses. 23
Figure 13 Cumulative data from WP2 questionnaires asking a variation of, "Would you pay for the POP service you received?", based on 19 responses. 26

1 Executive summary

This document presents an updated market analysis and compares POP's market with recent data. It shows that there is a wide range of vertical HPC markets and that software has a large opportunity for investment and improvement that POP is well placed to take advantage of. The current customers served by POP cover the breadth of European HPC and we are continually working to improve the representation and coverage of the service.

Existing discussions regarding the sustainability of the project are presented along with lessons learnt during the execution of the project and their relationship with regards to sustainability. The project has made great progress in its first stage, we have successfully developed a service that provides clear value to organisations while improving the competitiveness of European science. However, in its current state the POP services are expensive, so customers may be unwilling to take the risk of paying for them at the current cost although these services do provide them with clear Return on Investment (ROI). From the work done to date, it is clear how we should work towards sustainability and decrease our costs. Nevertheless, some questions still emerge regarding the viability of a fully commercial service to add value to the full European HPC ecosystem as certain groups will have less ability to pay than others.
Following this discussion, a business plan is presented for the next phase of the project along with a long-term view of the sustainability of the project and how we aim to work towards this.

2 Updated Market Analysis

The HPC market is large and continues to grow. The European HPC ecosystem was worth €6.8 billion in 2016, 20% of the worldwide HPC Market[1]. Spending on HPC servers is by far the largest proportion at 51.7% of the total in 2016, with software being a relatively small part of the HPC Ecosystem with 15.9% of the total. However, IDC 2015 report[2] "considers software the single most important category for future HPC leadership". Due to the focussed investment in hardware for HPC, the software is struggling to exploit it efficiently. Similarly, the large number of new technologies (ARM, Xeon Phi, Power, many-core, multi-core, and vectorization) increase the burden on the software developers to take advantage of the many levels of parallelism to achieve efficient performance.
The IDC 2015 report states that the world-wide market for HPC software was €4.4 billion in 2013, and roughly €1.2 billion for Europe. It is estimated to expand to €6.4 billion world-wide in 2018 and €1.7 billion in Europe in 2018. IDC estimates that European independent software vendors (ISVs) today

[1] Hyperion Research 2017 "*Trends in the Worldwide HPC Market*" ISC17 slides
[2] IDC 2015 report on "*High Performance Computing in the EU: Progress on the Implementation of the European HPC Strategy*" DOI 10.2759/034719

represent 15-20% of the global HPC market for ISV software and 25-30% of the European market.

IDC's 2011 study on highly parallel software in Europe[3] found that the majority (83%) of the most important parallel software applications in use at the surveyed European HPC sites were created in Europe. Intellectual property rights for a substantial majority of the sites' most important application codes (66%) were exclusively owned by European organizations. But many of these important codes are used only by one or a handful of HPC sites. This means that there is a large but widely dispersed software market in Europe.
Some of the main markets for spending on HPC are:

- Manufacturing, which accounts for 16% of GDP and spent €450 million on HPC in 2013
- Oil and gas is responsible for €440 billion of EU GDP
- Health which includes pharma, bio, and bioinformatics accounts for 10% of EU GDP, spent about €416 million on HPC in 2013.
- Weather forecasting spent €173 million on the HPC ecosystem for weather and climate in 2013 and it is estimated they will spend around €230 million in 2018.

Since POP is transversal the whole breadth of the software market for HPC in Europe is available to us.

The Solve 2014[4] report looked at the challenges for US companies with respect to performance improvements and found that in general the scalability and cost of the software were seen as the significant limiting factors, as well as limited expertise. These are the areas where POP can add value to companies. Most US companies looked to their technology vendors for help in significantly improving the scaling of their applications. For SMEs, technology vendors are unlikely to be their first port of call since they are unlikely to have such a good relationship and may not even own their own cluster; this is where POP can position itself. Overall the Solve 2014 report found that "Software is the greatest gap in efficiency for industrial HPC users today" and there is an "acute need for software to leverage the improved hardware capabilities". Since the bulk of HPC spending is on hardware, software is lagging and this is why initiatives like POP and the other CoEs are so important and have such room for impact.
Due to the transversal nature of HPC there are many niche markets which support ISVs, which are generally very specialised and fragmented. This makes them harder to reach as they often will not self-report as being part of the HPC community but it also means there is room and need for a community to help them navigate the challenges particular to HPC, irrespective of their vertical markets. Goldbeck consulting[5] studied the materials modelling market and found a flat market of ISVs with new players

[3] IDC 2011" *Financing a Software Infrastructure for Highly Parallelised Codes"*
[4] Solve 2014, "*The Exascale Effect: the Benefits of Supercomputing Investment for U.S. Industry*"
[5] Goldbeck Consulting Ltd 2012, "*The economic impact of molecular modelling: Impact of the field on research, industry and economic development*."

entering the market being balanced out by the consolidation of others. The number of end users was orders of magnitudes larger than the number of ISVs, and so by POP improving the efficiency of a small number of codes we can have a wide impact.

2.1 POP Customer Base

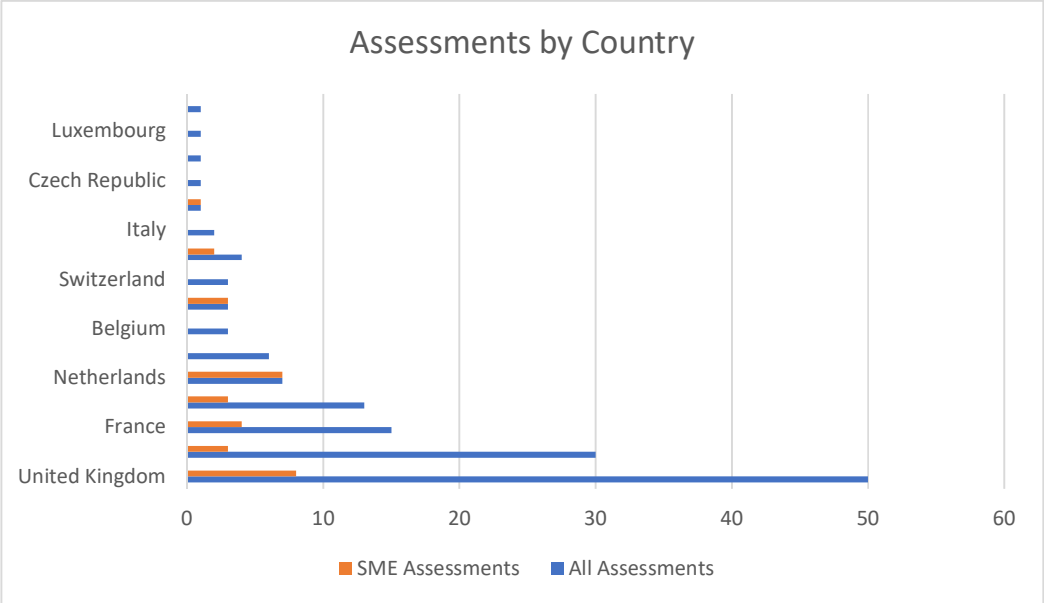

Figure 1 POP assessments by Country

As of 04/08/2017 POP has worked with customers from 16 different EU+ countries, as shown in Figure 1. The UK, Germany and France are where most of our customers come from, given the initial advantage that these countries are where the existing customer bases of the partners tended to be. However, we have clearly branched out from just consortium countries and are servicing a wide range of countries with varying degrees of HPC usage.

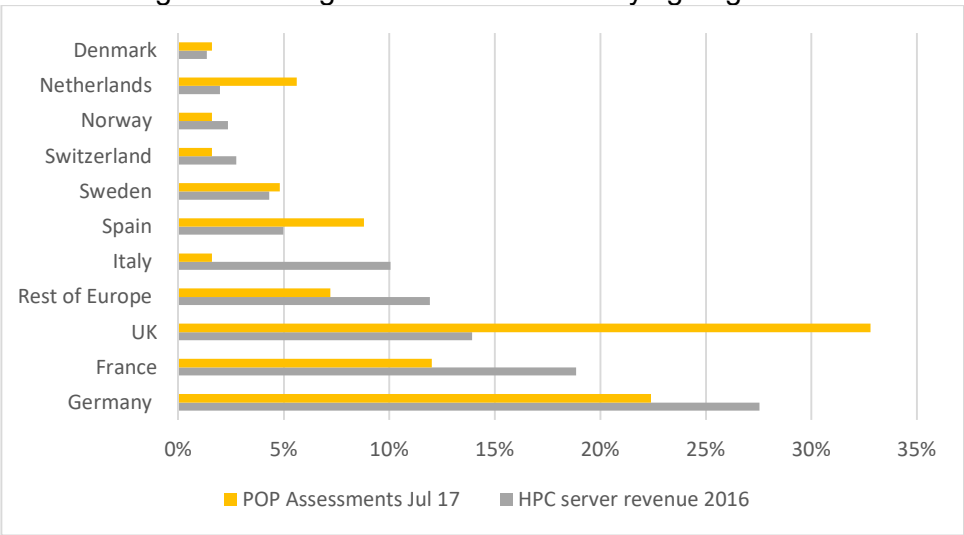

Figure 2 Share of POP assessments by country compared with HPC server revenue from 2016

To put our numbers more into focus, in Figure 2, the percentage share by country is plotted against the HPC server revenue by country in 2016[6]. A direct correlation between server market and software market size is obviously not expected but it is useful to give us a baseline. For example, the French market share will be overrepresented due to Bull Atos' large share of the HPC server market since, according to IDC, in 2016 Atos/Bull captured 3.6% of European HPC server system revenue. This would bring the French percentage share of the assessments down more in line with our figures. In general, we are servicing countries reasonably proportionally to their HPC server revenue except for the UK, Spain and the Netherlands which are a bit high and Italy, which is lower than expected. Recently we have been focussing our efforts more on France and Italy, which has clearly paid off with more French studies, they are now the third largest section. However, we have still had little traction in Italy.

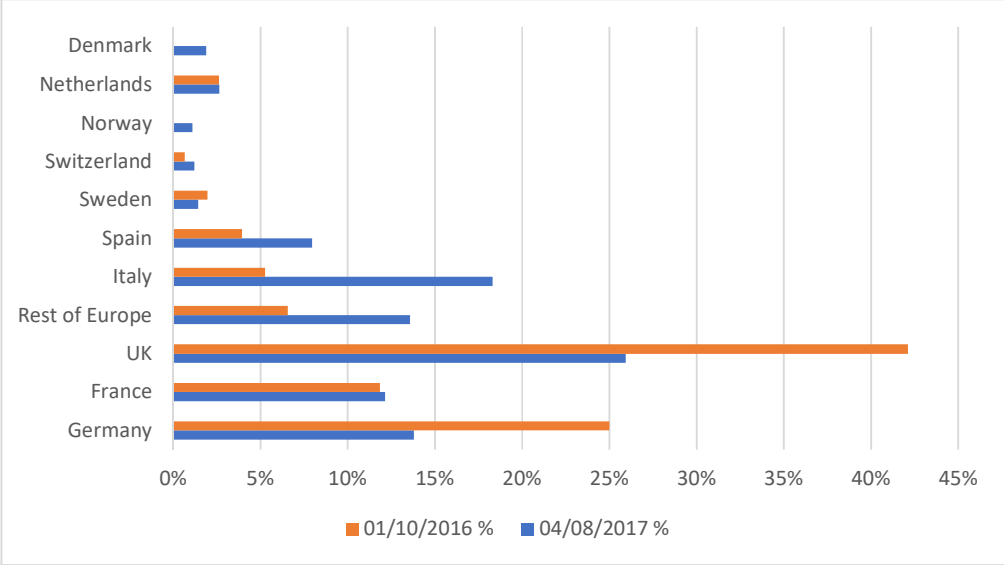

Figure 3 Comparison of the leads entered into the CRM on 1/10/2016 and 4/8/2017 by country.

Looking at the leads added to the CRM (Customer Relationship Management) system used by the POP partners, we have contacted people from 26 EU+ countries and Figure 3 shows the number of leads by country that have been contacted by POP. A lot of Italian leads have been contacted without much conversion to sign ups but in general the number of users is reasonably proportional to the number of leads we have contacted in that country. The figure above compares the proportion of leads from each country in the CRM from the first year of the project and the state as of 04/08/2017. The proportion of leads from the UK has reduced considerably but also leads from Germany have fallen by a similar proportion. The proportion of leads from France is very similar but Italy is much higher than before. This shows that we have made reasonably progress on our goals of reducing the UK's share and building up the other countries that were under represented. We have also doubled the number of countries we have reached out to, as of 1st October

[6] Gathered from the SMART 201/0021 Open Dataset for which the source is IDC 2015

2016 POP had only contacted leads from 13 EU+ but by the 4th August 2017 POP has contacted leads from 26 EU+ countries.

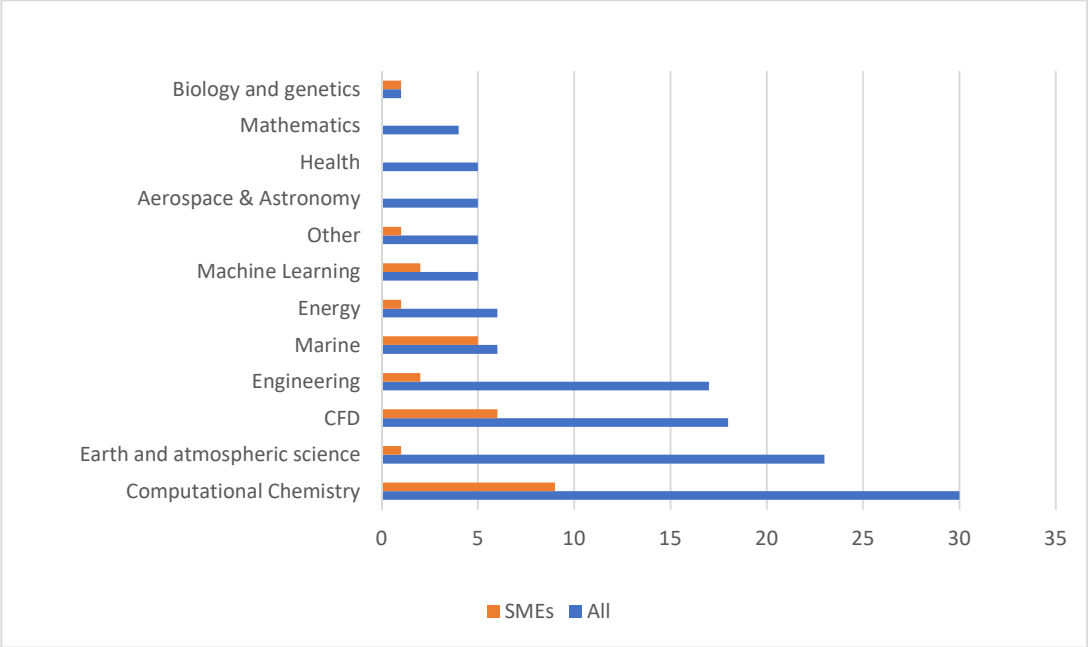

Figure 4 Number of assessments by sector

Looking at our studies by their sector, in Figure 4, we see that there are four main sectors, Computational Chemistry, Earth & Atmospheric science, Computational Fluid Dynamics (CFD) and Engineering, and we are working with SMEs across the board. See Section 5 for the full categorization of the sectors

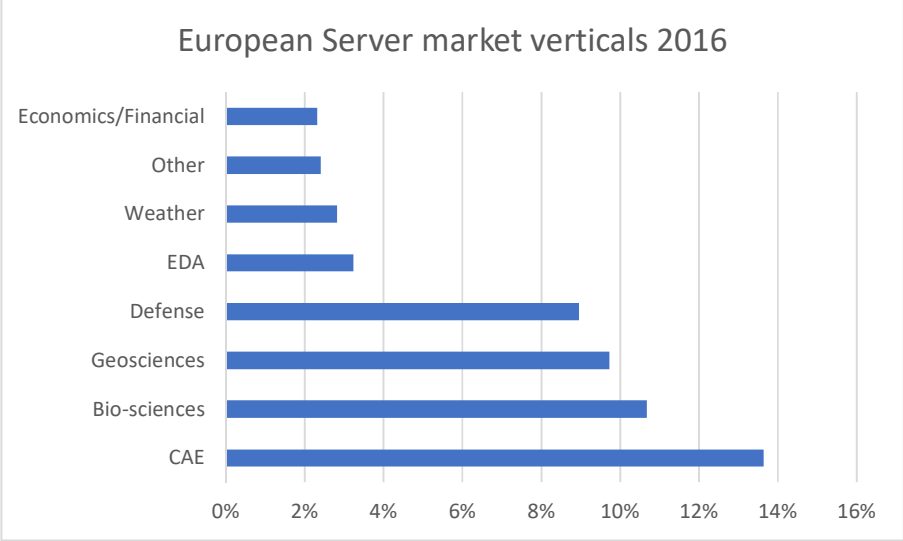

Figure 5 European server vertical markets in 2016, excluding the government and university markets.

As shown in Figure 5, the largest HPC sectors by server system spending in Europe are Computer aided engineering (CAE), Bio-sciences, Geosciences, electronic design and analysis (EDA) and weather. These line up quite well with the application sectors that are most serviced by the POP project. The

sector definitions are slightly different and especially for ISVs it is difficult to pin down the exact sector since the codes may be used for varying applications by end users. For example, one ISV that falls into the Computational Chemistry sector, has an application that is used by pharma and oil and gas. Similarly, a standard CFD package could be used for engineering, health, aerospace and other sectors.

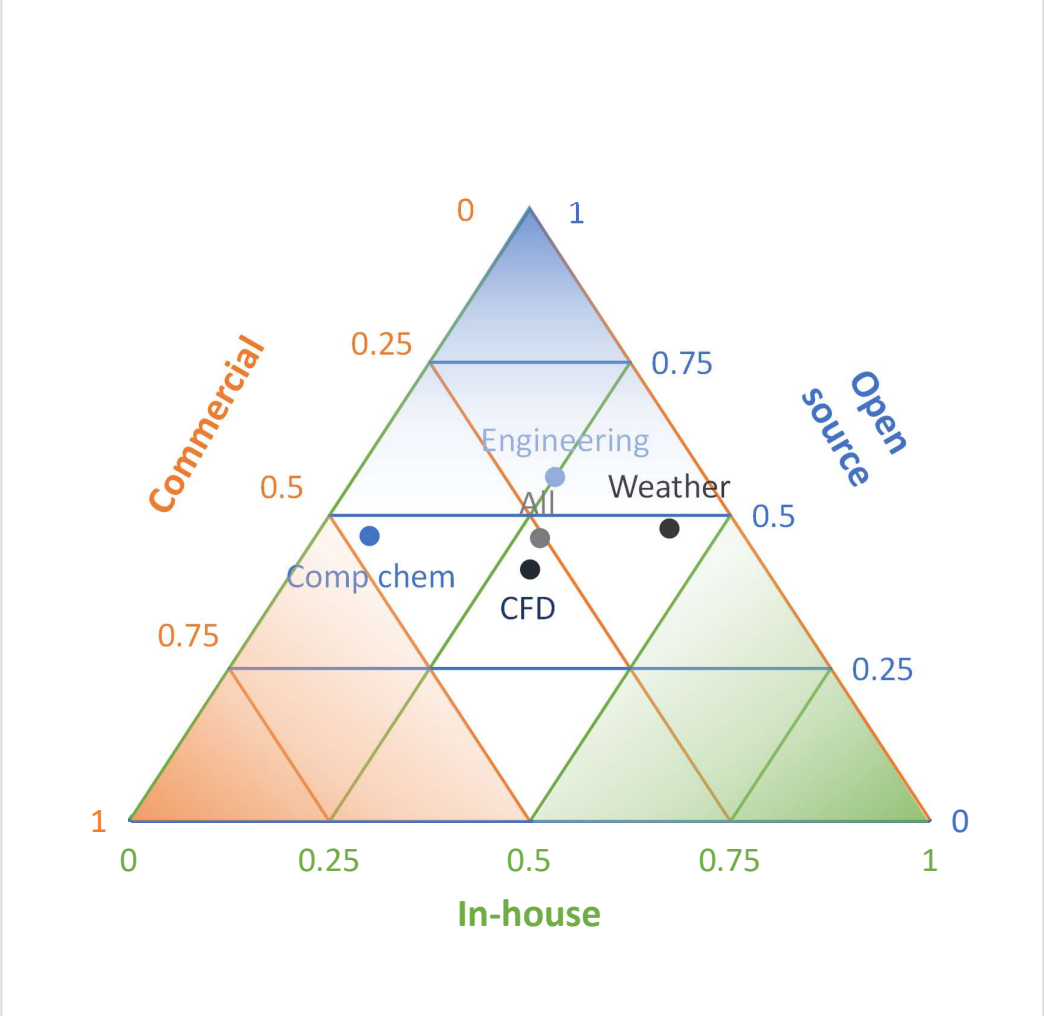

Figure 6 3D plot showing the proportion of assessments that are classified as commercial, open source or inhouse by sector.

Sector	Number of assessments
All	121
Comp chem	30
CFD	17
Engineering	16
Weather	23

Table 1 Absolute number of assessments for sectors

The codes POP has investigated fall into one of three categories: In-House, Open Source or Commercial. Figure 6 shows how the largest four sectors and all code studies are divided between the three categories. The absolute

number of assessments for each sector are shown Table 1. Since the three categorisations are independent they always add up to one for each sector. This diagram gives us a way to visualize different behaviour and characteristics of the main sectors we currently service. The centre of the triangle corresponds to an equal proportion of codes falling into all three of the categories, the background is coloured to indicate the maximal direction for each category. Following the coloured bars to the same coloured origin shows us their value on that axis, for example the 'Engineering' sector lies on a green line which we follow down to the green In-House axis to show that around 25% of 'Engineering' codes are in-house applications.
There is distinct variation amongst the sectors with Computational Chemistry having the largest proportion of commercial software and Earth & Atmospheric sciences having the least commercial but most in-house. The values for 'All' assessments are 28% in house, 46% Open Source and 26% Commercial. Overall the proportion of Open Source software is high for all sectors (over 33%). This means that POP has a wide-reaching impact since the commercial and open source software have a wider pool of end users than in-house applications. In-house was defined to be applications that had limited availability outside of a certain group and does include some large collaborative codes (especially for weather simulations).

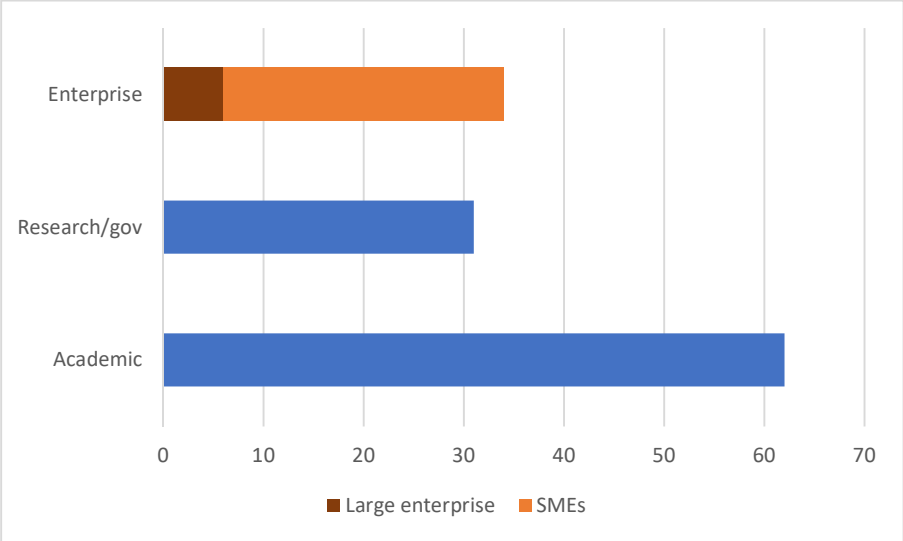

Figure 7 Organisation classification for POP studies

The breakdown of POP studies by customer type is shown in Figure 7. Academia is by far the largest segment, which is unsurprising since academics can be quick moving and can often make the decision to sign-up without management oversight. Commercial organisations make up over a quarter of the total number of assessments performed by POP. This is larger than the percentage share of the market that European ISVs hold and is very positive.

2.2 Current Marketing Approach

Our current marketing approach is paying off well, we are ahead of target of our aim to get 150 studies for POP and those studies are from a broad range of countries, sectors and organisations.
Our approaches are briefly discussed in the following sections.

2.2.1 Events

POP has continued to attend a wide range of HPC events as well as smaller sectorial ones. Locating SMEs at events is challenging because they don't often have large budgets to allow them to travel or attend many events. However, this has still been quite a fruitful source of leads especially for academics and building relationships with other Centres of Excellence, PRACE and large vendors.

2.2.2 Direct Marketing

Direct mailings involve searching for leads using LinkedIn and internet resources then emailing the contact directly. Most recent leads have been located using these means. It has been an important way to reach out to SMEs and potential customers that we wouldn't meet at events.

2.2.3 Webinars

NAG has recently created a POP Webinar series to disseminate our work and knowledge and as a vehicle for marketing and community development to find interested leads and improve participation and engagement with POP. This is a new venture and will be discussed in slightly more detail.
So far two Webinars have been advertised and delivered. The advertisement has been mainly via direct mailings specifically searching for people working in SMEs and more under-represented countries in POP.

For the first webinar, of the emails that were delivered over 30% were opened and of those opened emails over 40% registered for the webinars. Nearly 60% of the registered people actually attended the live webinar. Compared to NAG's normal experience these are very good rates.
The rates for the second webinar very closely match the first one and close to 30% of people registered for the second webinar were also registered for the first one. This shows that we have managed to reach out to a wide range of people with the webinars but also that people continue to engage with POP on a recurring basis.
We plan to continue to produce webinars for the remainder of the project.

3 Lessons Learnt

Lessons learnt from the operation of the first phase of POP should be used to inform the next steps and pave the way for sustainability.
Discussion of options for the sustainability of POP have already occurred in D3.5 (Report and Recommendations for POP Sustainability), D3.2 (First Market Review) and in the POP CoE Proposal a business plan was put forward.
The main uncertainties around the previous discussion has been the costs for POP to provide the services and the need for data from WP2 Customer Advocacy to further inform pricing regarding:

1. How much customers would be willing to pay for the service provided by POP
2. Customers' interest in staying in the service
3. The services customers found the most useful
4. Customers' estimated ROI, excluding the investment made by POP.

D3.5 outlined alternative approaches to ensuring the sustainability of POP. This included discussion on the issue of public funding potentially preventing some groups of users accessing the POP services. It recommended that we work towards sustainability by applying for a further period of CoE funding and partnering more closely with HPC Centres and PRACE to offer service vetting applications and ensuring efficient use of HPC resources.

The business plan put forward in the POP CoE Proposal, in hindsight, is overly ambitious at this stage. Covering 50% of the project costs in an application to a second CoE call that might be as large as €6-8 million without having done any commercial work or tested the viability is overly risky. The necessary number of customers paying fees is also very high, estimated in the proposal to be 70 in the first year after POP 1. Considering the KPI for WP3 is having 50 members of the POP Community at the end of POP 1, this is unrealistic. POP will have serviced around 150 codes via our assessments, a significant number of those are follow up work or work with the same organisation - close to 30% of closed audits have follow up work. This leaves us with 105 unique organisations that might become paying members of the POP Community. We would need a 66% conversion rate to reach the first-year estimate in the proposal, which is extremely high.
To ensure the long-term sustainability of POP we need to minimise risks and slowly build towards sustainability based on lessons learnt during the first phase of POP.

3.1 The Big Picture

POP aims to boost productivity for the EU ecosystem of HPC and parallel software whilst becoming self-sustaining. Understanding how to move to a position where we can meet both these goals is a challenge, and nothing like this has ever been done before. So far POP has been successful and as a result, we now have a much deeper understanding of how to provide these services and of the challenges around sustainability. The consortium has the unique mix of skills and experience necessary to deliver the required services. However, we need continued EC support to move towards sustainability, and to continue providing these important services to the HPC community in Europe.
The Key Challenges that we need to overcome to ensure the success and sustainability of the POP service are:

- How do we ensure an optimal balance between maximising income and maximising benefit to the EU when key organisations may have little or no funds?
- How do we ensure the continuity of income necessary to sustain the POP partners?

- How do we build an EU community if reliant on commercial income from specific organisations?

As a transversal CoE rather than one built around a community, POP's main offering is a service rather than a community so if we sustain ourselves purely on commercial revenue it will put up barriers with regards to who benefits and makes it difficult to offer the service to all of Europe.
For example, academics are used to getting HPC for free and there is a culture averse to funding POP type services for academics. This makes it hard to build a business that also boosts the full range of the European HPC Ecosystem. If academics apply for research council funding for POP services these budgets are often fixed for two or three-year chunks. Hence organisations may have considerable lead times before they get funds to pay for POP, which calls for a more gradual move towards self-sustainability
POP's aim is to improve the full range of codes on HPC from small users to big consortia, including people that aren't necessarily able to get their own funding. Therefore, there is a fine balance to find between commercial sustainability and making sure there are no barriers discriminating against certain types of applicants.
Academics will always require some form of funding to use this service and SMEs are often in a difficult position with respect to resources so these groups might find it especially difficult to pay for POP services. To continue to service academic codes and help advance science in Europe POP partners would be unable to bring in regional funding without certain limitations on the beneficiaries. Having central funding for POP to service academics would also lower the overheads compared to if all participants needed to bring or apply for their own funding separately.
The end goal of POP is to be sustainable and thrive in the long term. We have shown that our service and expertise brings large benefits to Europe, financially in terms of the ROI as well as the impact on science, allowing faster results and greater (i.e. more realistic) complexity. We want to ensure the wide availability and broadest impact of POP for the future.

The following sections will discuss areas of operation of the POP service and identify the scope for improvement.

3.2 Tools

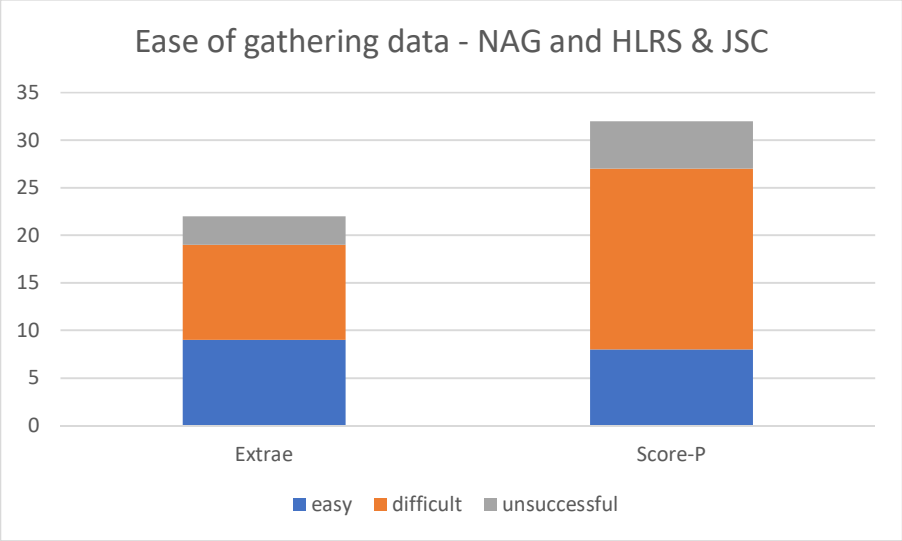

Figure 8 Categorisation of the usability of the performance analysis tools from data gathered by the POP partners NAG, HLRS and JSC.

We assessed the usability of the tools used by the POP analysts, Figure 8 shows the usability of the main tools the consortium uses with regards to gathering data. These most frequently used tools are part of the VI-HPS consortium[7] and are developed in part by members of the POP consortium. The categories were defined as: "easy" - use of the tools needed only the official documentation; "difficult" - required other sources e.g. developers or other specialists; "unsuccessful" - it was impossible to obtain the full trace data.

Figure 8 shows that for more than 50% of all uses of these tools gathering the data was difficult or unsuccessful.
A reasonable proportion of POP analyses that were accepted pushed the boundaries of what is possible with the tools so that we could work out the limitations. This will have an impact on the data, and is an argument for improving the capabilities of the tools as well as the usability, to allow POP to target the broadest definition of HPC applications and have the widest impact.

[7] http://www.vi-hps.org/

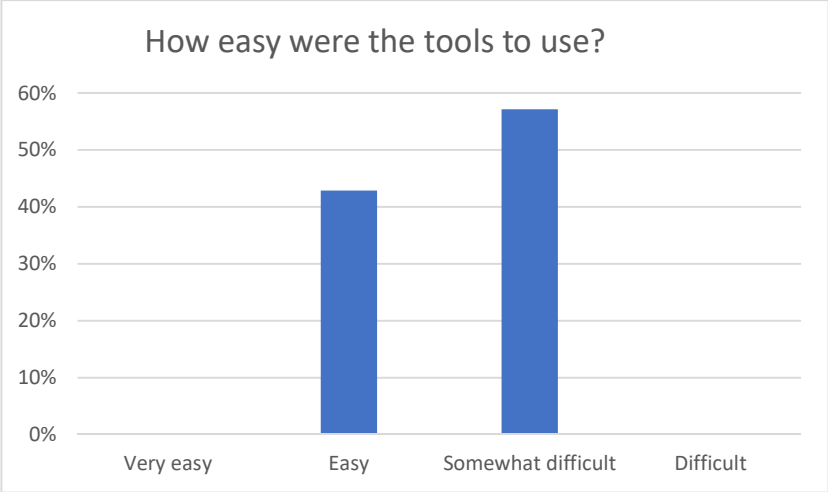

Figure 9 Data from WP2 user questionnaire asking, "How easy were the tools to use?", based on 7 responses.

This is an issue that doesn't just affect the POP team but customers as well, since they often gather the data themselves. When asked how easy the tools, in a questionnaire sent out by WP2, were to use over 50% said they were somewhat difficult, as shown in Figure 9.
With the high failure rate, it is often necessary to change tools entirely to gather the data. Often the new tool will be more specialized but consequently have limited information available to the analysts. All of this has a large impact on the time taken to perform a POP analysis which drives up the cost of the service.
Having difficult to use tools is especially an issue for SME customers because they are more likely to want to gather the data themselves for confidentiality reasons and if we charge them for the service they will not be satisfied with having to spend a lot of time as well as money because our software is not robust enough or doesn't support all possible parallel features found in complex software.
Improving the usability and productivity of the tools POP uses would reduce the time each service takes considerably, as well as boosting the engagement with the tools and boost the inclusion of them in to standard development workflows.

3.3 Industrial Customers

Working with industrial users, especially SMEs, has been a very useful aspect of POP and has highlighted some challenges we will need to face going forward. When these customers engage with the service they find it very useful. However, there have been a couple of problems to overcome. Firstly, the degree and need for confidentiality, POP partners have been happy to sign NDA agreements and a few have been signed. This does initially slow down the process, especially as most partners need to get other staff involved. Being clear about our ability and willingness to sign NDAs in our first mailing to industry leads helps to reassure them but does result in some reports being unavailable to all partners due to these restrictions, which would not be the case if POP itself could sign the NDAs.

Secondly, in D4.3 it is shown that studies for SME customers consistently have the longest time stalled waiting for customer availability. This may not be such an issue if they are paying for the service, but it may negatively influence their decision to do so. SMEs are often working with tight deadlines and limited resources and hence find it difficult to set time aside when it can't clearly be justified. Therefore, we need to make a better case for the benefit of POP, set out what effort is expected from them and think about ways to minimise the effort they need to expend.

3.4 Services

Of the current service offering of POP, the Audit and Performance Plan assessments have worked well and the consortium has defined and set the scope for both pieces of work. This will need to be further refined for commercial services to ensure consistency of work and cost to the consortium.

To date, the Proof-of-Concept work has had less uptake than expected in the CoE proposal. Due to the clear and targeted recommendations made in the Audit and Plan assessments, customers are often able to easily implement the changes themselves and the benefits are clearly demonstrated in the reports. With their expert knowledge of the code base this is often quicker than having the POP team work on it. After discussions at POP user forums, via WP2 customer advocacy and discussions with SesameNET[8] about their experiences, we conclude that SMEs especially are more interested in long-term co-development work to fully implement the recommendations in a style more like the Fortissimo and SHAPE projects that help them build up their technical skills and incorporate performance analysis into their workflow at the same time.

There clearly is still scope for the work, since POP can investigate and demonstrate ideas that would have a large impact on the code and hence are often too risky for the developers to spend much time on.

There is also the potential to offer more light-weight services, for example to HPC Centres that want to assess a wide range of applications. With the current service offering this would be a very large time commitment since one Audit currently takes at least a month. However, we could offer a light-weight service calculating only the POP metrics for applications to give a first impression of the overall efficiency of many applications. This would rely on the usability of the analysis tools and is currently not feasible.

3.4.1 Cost of Services

To become sustainable via commercial revenue we need to understand the costs of the work to form the basis of pricing.

The average cost of a technical full-time employee (FTE) across the partners is €6500 for one person month (PM). The average amount of effort taken to perform one POP assessment based off the data from year two of the project is 1.9 PMs.

This means that the average cost of an assessment is €12,350 without any overheads. When we include overheads of management, finance, marketing and business development which at a minimum would be one third of the

[8] https://sesamenet.eu/

technical effort, this grows to €15,700 cost for each assessment. This is the average of all assessments, audits and plans. Therefore, this is a minimum price that we could offer the POP service without losing money.
The current KPIs for the POP project state that the average amount saved should be €20,000, this is close to our current minimum price and clearly is uneconomical to expect customers to pay such a high amount with a large possibility that they might not re-coup the costs. For POP to build a commercial service, it clearly needs to bring more value to the customer than it costs them.
We need to take advantages of ways to reduce costs that can be highlighted from our current operation to make the service economically viable. The discussions above highlighted a few key areas with potential to improve the commercial viability of the POP service.
The main cost of the POP service is the time of the POP analysts. Reducing the time spent on each service would drastically cut our costs. For example, if we halved the time it takes for a service then this would bring the minimum price down to €8,550, which is a lot more commercially viable.
One key area for improvement would be to improve the usability of the tools and increase the automation of the POP assessments so that analyst time is focussed on the expert interpretation of the metrics and locating underlying issues rather than the base work of installation, gathering the data and getting it into a good state to analyse. This is currently taking up a significant amount of the effort and is not making the best use of POP analysts' time. It is a critical bottleneck for the commercial performance of POP. This also becomes an issue when customers perform some of those steps themselves, which is most common with industry, as it can be a large barrier.

3.5 Action Plan

From the first phase of POP we have gained and achieved a considerable amount of experience and knowledge of how to build towards sustainability but we are not there yet. We therefore, propose that the scope of the next CoE be changed fundamentally to be a Centre of Excellence in Technology and Practice of Performance Analysis.
To become sustainable, the services we offer must be economical and competitive for which we must significantly cut our costs. Similarly, to boost the skills and incorporation of performance analysis into standard workflows, we must reduce the barriers and the effort involved in using the performance analysis tools.
Therefore, we propose a future strand of POP to develop the tools and technology. They are currently great academic projects but not robust or user-friendly enough for wide spread use when charging commercial rates for the service and the data has shown them to have the largest potential to reduce the costs of the project. They are areas which are not regarded as research and hence with no alternative funding channels. A new strand of the POP CoE will be tasked with improving usability, such as user interfaces, documentation, increased automation etc. Investigation of the capability requirements for POP to service a broad range of HPC applications is also necessary to inform future feature development.
This work should be done before we are able to take on commercial work.

4 D3.6 POP 2 Business Plan

The main idea and model for POP has not changed since the initial CoE proposal and hence this deliverable will build upon the text from that proposal, updated in the light of lessons learnt and with data and experience acquired during the operation of the current phase of the CoE.

4.1 Mission

The Mission of POP is to help application developers and users understand issues related to the performance of their applications, together with ways to improve them, thus reducing waste and inefficiencies in their execution. Encouraging the incorporation of performance analysis into the standard development workflow and building a community around this is a key part of the mission of POP.

4.2 Key Expertise

The Project Team represents the European excellence-level expertise in the development of performance tools and programming models for large scale parallel computing, operation of Tier-0 systems as well as in promotion of HPC and its application to the industrial sector.
The Team's expertise is best-in-class and it represents the hands-on, practical attitude needed in this work, combined with an excellent track-record.
POP has standardised and continues to develop a unique and extremely informative methodology to assess the efficiency of parallel applications.

4.3 Customers

The primary customers of POP are code developers and owners including the academic organisations and companies that develop computing codes for a variety of domains.
POP will also target secondary users:

- Academic users – e.g. research projects deploying computing applications
- Infrastructure and Service Centres – e.g. PRACE and other Centres of Excellence, HPC Centres
- Industrial Users – corporate or SME organisations using computing applications
- Standardisation Bodies – feeding back our knowledge

Some of the customers have different motivations and ability to pay for POP services. As discussed in Section 3.1, academic code developers will have to acquire funding to pay for POP services but it is currently culturally difficult to justify funding POP type services. To continue to service large numbers of academic applications, POP will potentially require outside funding to cover these costs.
HPC Centres have a vested interest in optimising all the codes being run on their machines as this increases both value for money and enables better science (i.e. bigger or more complex problems).
Infrastructure Vendors often do benchmarking studies and performance optimisation that is similar to POP's services to help purchasers of machines get the best use out of them. POP has already worked with the Atos/Bull supercomputing presales team.

Industrial code developers and users, especially SMEs, often find it difficult to find the time to use the POP service, with long delays at their end before work can start. However, when the work is carried out, they find it extremely useful and they are much more likely to pay for POP's commercial services.

4.4 Market

The project will target any type of parallel computing applications, regardless of scientific or technical area, from single node to many thousands of processors. Commercial services provided alongside the POP assessments may also target serial codes to get them ready for HPC and the key POP services.

For the main discussion of the market see the previous section on Market Analysis in Section 2.

4.5 SWOT Analysis

Strength	**Weakness**
World-class academic expertise Unique Methodology Complementary tool-sets Existing relationship with customers Proven track record and success stories Geographical coverage	Difficulty showing the value and worth of the service Cost of performing the services Difficult to use the POP tools
Opportunity	**Threat**
Large number of potential codes Expand to new regions outside of Europe Exascale and large number of new technologies require expertise Software is key limiting factor for HPC leadership	Academics unlikely/hard to get funding for POP services SMEs hard to convince and hard to locate 'can do it themselves' mentality

The POP team has made great headway so far and is on target to meet its KPIs. We are continually standardising our methodology and adapting it for new programming models. A lot of person hours go into performing a POP assessment and because it is so unique we should do more work educating people about what it entails and provides before we can build on our success and charge commercially for it. We need more time to establish the benefits for end users to market it effectively.

Improving the process and usability of the tools is also where we could have a significant impact in reducing our own costs and improving the viability of the commercial services.

4.6 Value Added and Competitive Advantage

POP will deliver savings in the operation of software applications. It will help eliminate waste, delays and inefficiencies from applications and related processes. POP is in a unique position to offer academic world-class expertise in this area. The future marketing strategy of POP will emphasise the worth of the POP tools in a wider business context, i.e. in the context of saving generation, lean manufacturing, zero-defect and Six Sigma initiatives.

The Added Value of POP is based on the following premise: the cost of improving a code is significantly lower than the losses incurred by its operating below its optimal level (and the loss incurred in the processes related to the code). In other words, the ROI offered by POP results from the difference between the cost of fixing an application and the cost of running it and related processes.

An example from an application assessed by POP that went through the whole suite of services highlights the savings from the pre-audited version to after the Proof-of-Concept changes were implemented. The code has between 100-200 users and is widely distributed as a commonly installed package on supercomputers. The specific example is just from the users that run on ARCHER the UK national academic supercomputer, who save €15.58 per run as well as being 72% faster-to-solution. Using the monthly usage data of the machine, the savings for a year were calculated to be around €56,000.

Another example from an application assessed by POP for a smaller code is that, one year of operating the code costs €20,000 and implementing the recommendations for POP costs €2,000. The achieved code speed improvement of 62% gives a yearly saving of €12,400 in the compute costs and provides a ROI of 620%.

Apart from savings in the cost of running an application, POP will also deliver further value in terms of fast and more efficient processes. This can facilitate the use of more detailed and complex models in simulations or increasing the size or scope of the problem solved.

POP will also deliver value in terms of the optimal operation of critical processes and related operations. For example, if an application is used to perform car crash tests, its improved speed and performance will improve the ROI of the entire manufacturing process, the factory and the company itself.

POP's main competitive advantage is its combination of world-class academic and commercial expertise together with the customer advocacy process and user forums that enable customers to participate in the governance of the project.

As of writing this document, POP's average ROI gathered from customer feedback by WP2 is 585% which is well above the target of 100%.

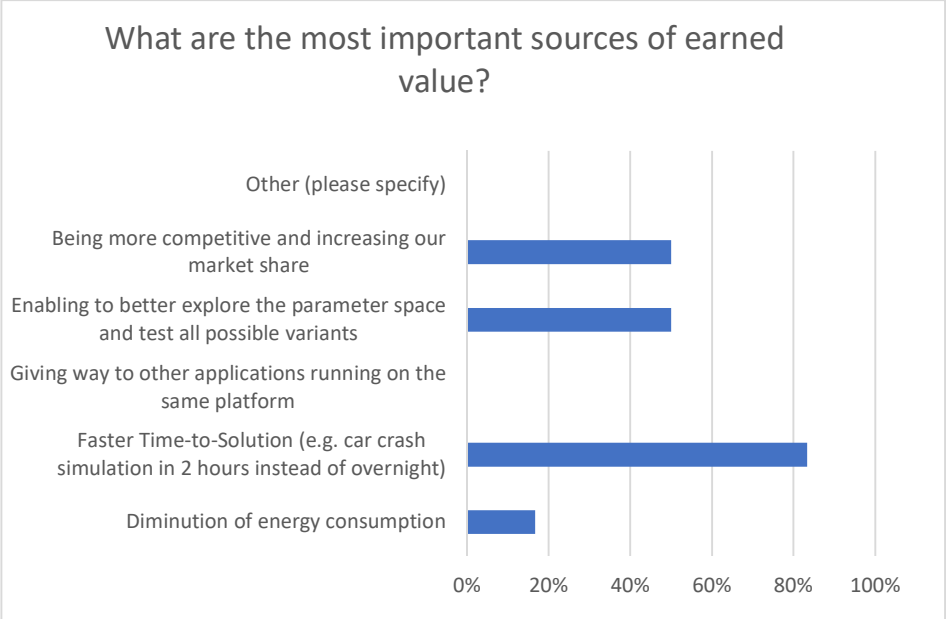

Figure 10 Data gathered from WP2 questionnaire asking, "What are the most important sources of earned value?", based on 6 responses.

It can be seen from Figure 10 that the most important source of earned value to existing POP customers is the faster time-to-solution they gain by implementing POP's recommendations. However, there is also a lot of value added by the improved code speed allowing more simulations to be run. Competitiveness is also an important source of added value from POP.

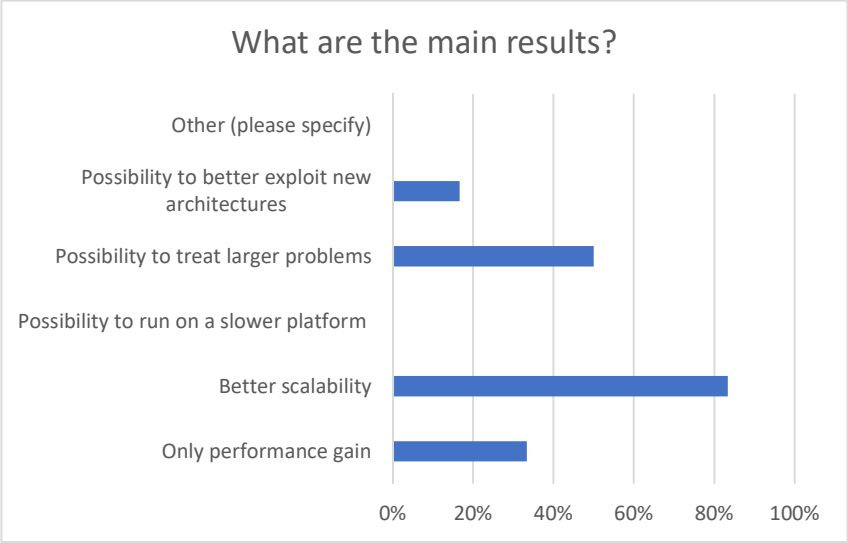

Figure 11 WP2 data gathered by questionnaire asking, "What are the main results?", based on 6 responses.

The main results gained from the performance improvements made using the POP recommendations, shown in Figure 11, are the improved scalability and being able to take better advantage of the HPC resources as well as being able to study larger problems.

So far, the average code speed up gained after implementing the POP recommendations is 43%, from data gathered in WP2 questionnaires.

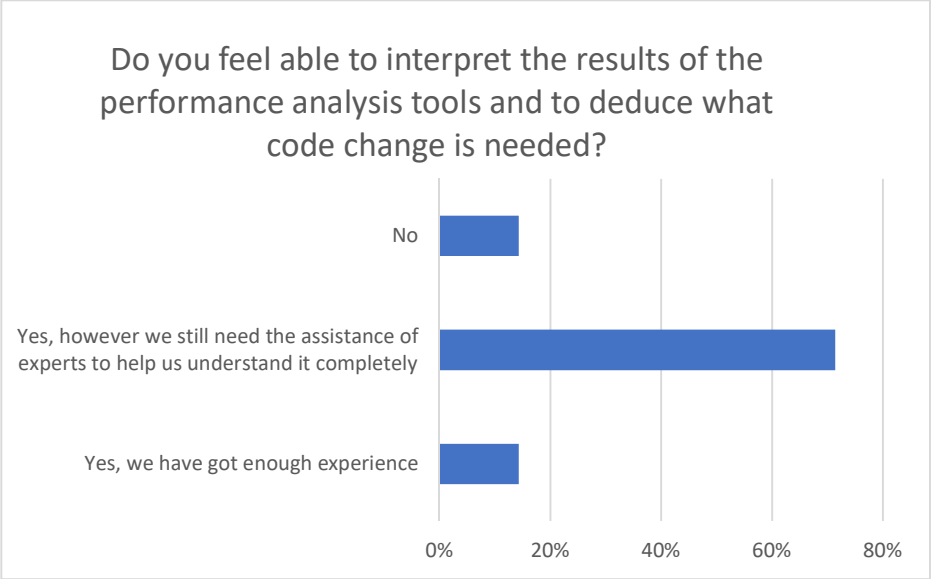

Figure 12 WP2 data gathered by questionnaire asking, "Do you feel able to interpret the results of the performance analysis tools and to deduce what code change is needed?", based on 7 responses.

Customers were also asked if they would be able to perform the work themselves, as shown in Figure 12, and whilst the customers might have the software development expertise most people need the specialist POP expertise to fully understand and interpret the results. The unique expertise of POP is clearly a very strong selling point and where we can add value to European organisations.

4.7 Marketing Mix

The Marketing Mix of POP includes the detailed definition of its Product, Price and Promotion.

4.7.1 Product

POP will carry out performance assessments and related services on all types of computing applications from server to Exascale level. The POP products and collateral are detailed below:

- **Application Performance Audit** – Making the customer aware of the inefficiencies in their application. Use of tools and reporting metrics to highlight the nature and magnitude of the issues. Up to 140 consultancy hours included in the service. This is the primary service of POP – it will identify the issues with a customer's code and provide the right directions to improve it in the most efficient way.
- **Application Performance Plan** – Follow-up to the service above. This identifies the root causes of the inefficiencies located. It qualifies and quantifies approaches to address the issues and provides an estimate of the resulting improvement. Up to 420 consultancy hours included in this service (three person months of effort).

- **Training** – Targeted training events and material on tools, programming models and parallelisation approaches and especially on the POP methodology for performance analysis. For example, training for system operators, support staff at HPC centres or development teams.
- **Webinar series** – We will use a webinar series to boost awareness of the project and performance analysis whilst reaching a wider audience for training.
- **Proof-of-Concept** – Experiments and mock-up test-bases for customer codes, testing and demonstrating potential ways to improve performance; includes elements of code development. Up to 6 person months of effort.
- **Full Implementation** – Longer term co-development with the customer to fully implement POP recommendations.
- **Tools** – The performance analysis tools and the improvements made will be marketed as one of the main collaterals of POP. Encouraging developers and teams to include performance analysis as part of their workflow by making them more user friendly and robust.

Code developers are the main target of the CoE project and can benefit from all types of service offered. They will benefit from applying the insight gained via the recommendations and utilising the knowledge acquired to improve their codes not only in the short term but also in the medium and long term. Academic and industrial users of codes not developed by them but used in production mode can also benefit from the services. Detecting issues relating to the environment, data sets and target platforms will enable these customers to address their code providers with precise evidence and requests for potential improvements. Since a lot of the codes that POP works on are Open Source, code users could easily engage POP to fully implement the requested changes.
Infrastructure and service Centres are interested in optimising the throughput delivered by their resources. Efficiency or scalability measurements of the applications are often used in the resource application and allocation process, but real productions runs may later expose different behaviour. The Centres can certainly benefit from precise performance assessments on applications they may suspect as having some performance issue or are taking up the largest proportion of cycles on the machine. The support staff are also an important target for the training activities as they can in the future provide a first level support for the users in the Centre.

4.7.2 Price

Currently customers receive the POP services for free. We will gradually move from this funded model to full sustainability by first reaching a point of commercial viability due to reducing our costs and potentially selling services alongside the POP Performance Audits. The current average cost of one assessment is €15,700 and the current target for the average savings from the POP service is €20,000. To gauge a viable price, we have investigated the pricing of the closest available options.

HPC Centres often offer some services that are similar to POP, aiming at improving the performance of applications. However, selling ourselves as training and collaborating with them is possible and beneficial to both. For example, POP has done work passed on by the Research Software Engineers Computing Support team at Sheffield, provided training for other CoEs on performance analysis, been advertised through other University HPC Centres (Manchester Research IT, Warwick CSC, Birmingham) and collaborating with Sesame NET and Arctur HPC Centre in Slovenia. We have worked with Atos/Bull providing POP assessments for their supercomputing presales team.
We are also in contact with PRACE to offer our POP skills to PRACE users. The first collaborations have been as part of a SHAPE project from a POP user and defining the policies and alternatives to include POP services in the PRACE portfolio

The Fortissimo Market Place is a great place to find HPC services, the closest service available on it is HPC benchmarking[9] priced at €7,000 – up to 155 consulting hours "Get a report about the performance of your HPC application". This is much lower than the level of detail than in a POP assessment as there is no profiling analysis or concrete recommendations. This is potentially a lower bound on the pricing of POP Audit assessments. Alternatively, organisations could do the performance analysis themselves. This would involve using a performance analysis tool to get the same level of information. Taking as an example the Intel tools, an upper limit on the price allowing customer to use the tools for a wide range of parallel software on a cluster would be a 2 seat Intel Parallel Studio XE Cluster Edition (which includes VTune and ITAC) for around €12000 (for the first year). As well as the license, time for training, learning and to perform the analysis will cost the company considerably.

If they don't have the tools, the experience, the hardware or the personnel free to do the work POP can add a lot of value.
NAG or TERATEC will subcontract commercial POP work to the university partners since they have the appropriate processes and insurances in place to deal with commercial work. Any commercial revenue gained under POP whilst we are in receipt of EC funding will be offset against the funding.
The price put forward in the POP CoE proposal was €5000 but, this is clearly not feasibly as it is well below the cost per audit for the consortium.
When we have reduced our costs, we will offer both subscription and one-off purchase models which will suit different clients. For example, HPC Centres will benefit from a subscription model as they have many codes running on their machine that need assessing whereas a small SME ISV that sells a single code may prefer the occasional one-off cost of assessing the efficiency of their code. Section 3.1 discussed the difficulties that Academics, Government Labs and SMEs would have to gain benefit from the POP service as it would be a difficult and long process to apply for funding. If we reduce

[9] https://www.fortissimo-project.eu/buy-services/hpc-benchmarking

our costs and charge slightly more, we could use the profit from commercial services to fund some free of charge services to organisations that would not be able to use them otherwise.

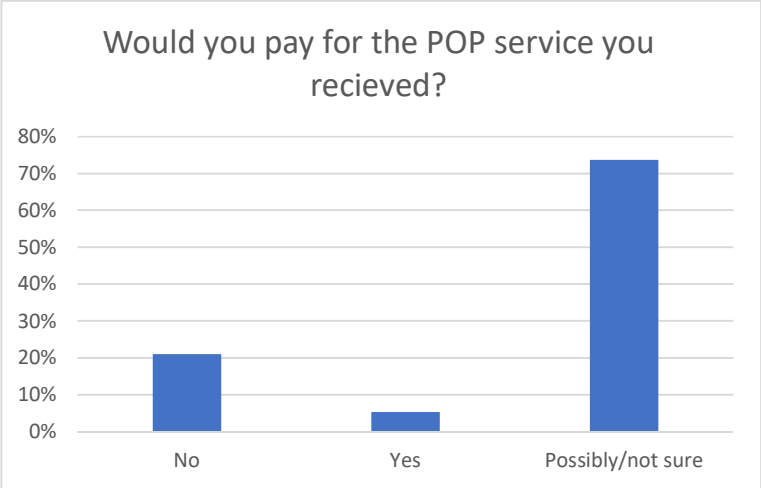

Figure 13 Cumulative data from WP2 questionnaires asking a variation of, "Would you pay for the POP service you received?", based on 19 responses.

A further difficultly in setting a viable price is a lack of robust data to inform potential pricing of the services from our current customers as, understandably, existing users of the service don't want to make any commitments to paying for something that they currently get for free. However, the data we do have, shown in Figure 13, is reasonably positive. Not many respondents completely dismissed the idea of paying for the services but clearly not many wanted to commit themselves to paying. We only had 19 responses to this question about the services, which makes it very difficult to currently put forward a clear price for the service.

For the future, when we have reduced the time effort per study, our pricing model will be to set the effort (person hours) available for each service and charge the customers the average rate amongst the partners. A viable amount to charge will depend on the market at the time, however, if we halve the person effort this becomes a realistic proposition and in line with the pricing of the closest available options.

4.7.3 Promotion

The promotion activities and current marketing strategy have already been discussed in some detail in Section 2.2. The main promotion of the service will be based on:

- **Success stories** – case studies from existing customers discussing the benefits and gains they achieve from the POP service
- **World-Class Expertise** – POP delivers a unique combination of world-class academic and commercial skills in the area
- **Long-term Relationship** – POP will continue to support its customers in the long-term by building a POP Community and enabling users to influence the organisation through the Customer Advocacy process.
- **Unique Methodology** – The POP partners have developed and honed a unique methodology for performance analysis.

4.8 Work Volume

The current work volume is sustainable for the current project size. To increase this, we would need more technical personnel and more marketing and business development. There is a strong correlation so far in the project between the number of hours devoted to market development and the number of customers signing up for the service. Therefore, to increase the rate of work more technical and marketing personnel are necessary. Once POP becomes fully sustainable, the volume of work will relate to the number of FTE technical experts and the success and time put into the Marketing and Business Development strategies. The full details of this will be left to the proposal for POP 2 CoE.

4.9 Revenues

Since we need more time to improve the commercial viability of the service and will be applying for the next round of CoE funding in part to allow us time to reduce our costs, it would be premature to provide revenue projections. However, we will highlight the potential for reducing the person effort of the services and hence the cost.

We will assume we have reduced our cost to such a point where it takes a Full-Time Employee (FTE) half a month of effort to complete a POP Audit. Using the current average technical FTE cost for the project of €6,500 and the current rate of work of six new studies per month. This will need three technical FTEs, two Marketing and Business Development FTEs and one Management and Finance FTE, in total six FTEs. Therefore, each Audit can be priced at €6,500 as a minimum.

If we increase the price of the Audits to €8,700, the six studies in a month would bring in €52,200 revenue per month and cost us €39,000, leaving us with €13,200 which is enough to offer two free POP services per month to customers without the means to pay. This model can easily be scaled up with more technical and business development FTEs as POP grows commercially. NAG or TERATEC will sub-contract work amongst the POP partners which will allow us this flexibility.

4.10 Operational Strategy

The potential customers that have applied for the service will be selected by the Project Executive Board based on proposals of the marketing and technical teams and whether they fit our criteria. If the potential customer is outside our criteria, for example based outside of Europe or has received multiple funded assessments already, then discussion around pricing will also take place.
The technical work will be carried out as close to the customer's real environment as possible which may require POP's tools to be installed there. Depending on the individual circumstances, e.g. confidentiality issues, the customer or the POP team may gather the required data. The analysis will be performed by POP followed by discussion/review with the customer and the presentation of the final report.

We will continue to build a community around POP, encouraging people to incorporate performance analysis into their standard workflows. Our webinar series, along with other communication, will guide and inform the community about key concepts and issues with respect to performance analysis and the associated tools and technology.
The process of the assessments during the current phase of POP will effectively stay the same:

1. Discuss requirements with customer
2. Gather data for representative test case (by POP or by Customer)
3. Analyse the data
4. Write report
5. Report to customer
6. Customer advocacy follow-up

Currently, the duration and effort of the assessments on average is longer than estimated in the proposal for the POP CoE. We aim to bring this down by around 50% via improved usage of the tools and their automation. Further work is also necessary to ensure the services are ready for commercialisation with respect to standardisation within the Consortium. We need to offer the same amount of work and analysis to customers paying the same amount otherwise customer satisfaction will be impacted.

4.11 Operational Goals

To achieve our goal of boosting productivity for the EU ecosystem of HPC and parallel software whilst becoming self-sustaining we have set the following Operational Goals.

- Make our services economically viable by reducing our costs via streamlining of the Audit process and reducing the personnel effort required
  - Reduce the Person Hours to complete an Audit by 50%
- Improve the usability and robustness of the POP tools
  - Aim for <25% of uses being classified as “difficult” or worse.
- Continue to build a community which uses performance analysis as part of their standard workflow
  - Produce POP webinars every three months
- Sustain excellent levels of customer satisfaction
  - >90% satisfied with the service in feedback survey
- Apply for the second round of CoE funding
- Aim to provide customers a combined ROI of at least 100% including the investment made by POP

5 Sector categorisation

We have categorised the codes analysed by POP as follows:

- **Aerospace**: Applications to solve scientific problems associated with the movement of airborne vehicles. Although the scientific methods these codes implement may be found in codes under other sectors (typically CFD), these codes are tailored to the aerospace domain e.g. through the choice of problem formulation or solver.
- **Biology and genetics**: Applications from the life sciences domain. POP audit examples include codes for protein structure and function prediction, and for genomics analysis.
- **Computational Fluid Dynamics**: General-purpose CFD applications that can be used in a variety of domains. For POP, this also includes a code that performs fluid modelling for computer graphics.
- **Computational chemistry**: Codes that model materials on the atomistic level, these are used in a wide range of domains, for example, Pharma, Oil and Gas, and Catalysis.
- **Earth and atmospheric science**: Codes that model physical processes in the environment with applications including meteorology, climate modelling and seismic activity.
- **Energy**: Applications to solve scientific problems specific to the generation and distribution of energy. POP has audited code for plasma physics (used in the simulation of fusion reactors) and wind turbine modelling.
- **Engineering**: Applications for the computer-aided design and analysis of physical structures. POP has audited several finite-element based structural analysis codes, as well as a fire dynamics simulator.
- **Health**: Applications that focus on aiding the diagnosis and treatment of illness. For example, POP has audited a lung-modelling code.
- **Machine learning**: Codes that implement techniques such as neural networks or Bayesian statistics.
- **Marine**: Applications to solve scientific problems specific to the maritime domain, including particle transport and the action of waves on structures.
- **Mathematics**: Codes that implement generic mathematical kernels that can be applied in a wide range of scientific domains. POP has audited two codes that implement the Fast Multipole Method for n-body simulation.
- **Other**: Everything not covered in the previous sectors. Topics covered by POP audits that have been placed in this category include acoustics, a code parallelisation framework, and text processing.

6 Acronyms and Abbreviations

- BSC – Barcelona Supercomputing Center
- CA – Consortium Agreement
- CFD – Computational Fluid Dynamics
- CoE – Centre of Excellence
- D – deliverable
- DoA – Description of Action (Annex 1 of the Grant Agreement)
- EC – European Commission
- EU – European Union
- FTE – Full-Time Employee
- GA – General Assembly / Grant Agreement
- HLRS – High Performance Computing Centre (University of Stuttgart)
- HPC – High Performance Computing
- ISV – Independent Software Vendor
- Juelich – Forschungszentrum Juelich GmbH
- KPI – Key Performance Indicator
- M – Month
- NAG – Numerical Algorithms Group Ltd
- NDAs – Non-Disclosure Agreements
- PM – Person month / Project manager
- POP – Performance Optimization and Productivity
- PoC – Proof-of-Concept
- ROI – Return on Investment
- RWTH Aachen – Rheinisch-Westfaelische Technische Hochschule Aachen
- SME - Small and medium-sized enterprises
- USTUTT (HLRS) – University of Stuttgart
- WP – Work Package
- WPL – Work Package Leader