



D4.1 Analysis Processes Version 1.0

Document Information

Contract Number	676553
Project Website	www.pop-coe.eu
Contractual Deadline	M3, December 2015
Dissemination Level	Public
Nature	Report
Author	Judit Gimenez (BSC)
Contributor(s)	
Reviewer	Cristoph Niethammer (HLRS)
Keywords	performance assessments, operational procedure



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	J. Gimenez	Initial Draft
V0.2	C.Niethammer	Initial Review and minor text changes
V0.3	J. Gimenez	Improvements as suggested by internal review
V1.0	J. Gimenez	Final version



Table of Contents

Executive Summary	4
1. Introduction	4
2. Performance Assessments	4
2.1 Performance Audit	5
2.1.1 Service identification and assignment	5
2.1.2 Performance data collection	5
2.1.3 Analysis and reporting	6
2.2 Performance Plan	7
2.2.1 Service identification and assignment	7
2.2.2 Performance data collection	7
2.2.3 Analysis and reporting	7
3. Implementation	7
4. Forms, documents and templates	8
4.1 Request Service Form	8
4.2 POP Service Terms and Conditions	9
4.3 Questionnaire on User Needs	10
4.4 Assessment Report Template	12
Acronyms and Abbreviations	14



Executive Summary

Two of the main services provided by the POP Centre of Excellence (Performance Audit and Performance plan) are implemented in the Analysis work package (WP4). During the first months of the project we have defined the operational procedures to carry out these services based on the partners past experience and the first studies started within the POP framework.

This deliverable describes in detail the standardized operational procedures to be followed in the assessments of Performance Audit and Performance Plan services. The report also includes the associated forms and templates defined.

1. Introduction

Two of the main services provided by the POP Centre of Excellence are covered by work package 4. During the first months of the project we have defined the operational procedure to implement these services, identifying the requirements and main steps based on the partners previous experiences.

This deliverable describes the operational procedures defined to implement the Performance Audit and the Performance Plan services. They were already partially used with the initial assessments carried out during the first months of the POP project. The procedures could certainly not be fully applied there because they were being defined at the same time. But these early experiences allowed validating the proposed procedure for the Performance Audit (we expect to start the first Performance Plans early next year). Despite what is defined in this document, future POP services may identify small improvements on the mechanism that will be adapted as soon as they are identified.

2. Performance Assessments

POP offers as service two types of performance assessments:

The Performance Audit is the primary service and may be considered a kind of health check for the codes diagnosing with a set of defined metrics the efficiency achieved on different aspects (e.g. parallelization, load balance, IPC, data transfer) and recommending areas for improvement. Even there is a commonality on all the Performance Audit; the customer can indicate specific topics of interest like serial code performance, scalability, and communications. See the Request Service Form for the full list of the predefined kind of studies.



The Performance Plan is a follow-up of the previous service targeting a more detailed analysis that identifies the root causes of the issues found on the Performance Audit and quantifies and qualifies approaches to address them. As a result of this second service, the POP CoE will identify potential cases where to apply the Proof of Concept service implemented in WP5.

We have defined the operational procedures described in the next subsections having in mind the sequence is to start with a Performance Audit and for some of the studies continue with the Performance Plan.

2.1 Performance Audit

2.1.1 Service identification and assignment

The Performance Audit has to be requested filling the Request Service Form available on the POP website (<https://pop-coe.eu/request-service-form>) where the user accepts the terms and conditions of the service (https://pop-coe.eu/sites/default/files/public/popw1/POP_Terms_and_Conditions.pdf). This form has been kept as simple as possible avoiding that potential POP customers do not request our service due to a large list of questions to be answered.

This first questionnaire acts as a filter to have certain guarantee that POP can serve to the customer. Nevertheless, more information is required before assigning the service to one of the POP partners from WP4. The information required is defined in the Questionnaire on User Needs (<https://pop-coe.eu/form-on-user-needs>) that can be filled directly by the user or by WP3 with the data collected in the initial interview.

Once all the data is collected and available on the ticketing system, WP4 will assign the partner and tool to be used based on the information provided by the user and the current workload of each partner. Initially we will try to discuss the assignment by email but if multiple requests have to be assigned we will organise calls to make the decision. The POP partner responsible for the study will then contact the user to start the Performance Audit.

2.1.2 Performance data collection

Four different scenarios are foreseen for this phase depending on the customer requirements and constrains.

- 1) The user obtains the data on his machine
- 2) POP obtains the data in the user machine
- 3) The user obtains the data on POP or PRACE machines
- 4) POP obtains the data on POP or PRACE machines

We will promote to use the customer machine so the performance data reflects the behaviour on a system typically used for the executions. POP or



PRACE machines may be used only in few cases for executions at large scale or when the user has a problem to assign resources for the run.

In case the user machine is used, an additional first step can be to install the instrumentation tool on the target platform. Depending on the user restrictions and preferences it can be done by the user or by the POP team.

In all the cases the user will select and specify the input cases to be analysed.

There may be differences on approach promoted for the data collection depending on the tool that it is going to be used. For instance, if the selected tool is Paraver or VTune, we will promote the user obtains the traces so he/she gets familiar with the procedure. In case the selected tool is based on Score-P the preferred option is that the POP team obtain the traces because the filtering step requires some expertise. In both cases, the objective is that after the study the user knows how to instrument his/her application.

2.1.3 Analysis and reporting

The POP experts will analyse the performance data collected in the previous phase documenting the results and recommendations in the Assessment Report following the defined Assessment Report Template (https://wiki.pop-coe.eu/img_auth.php/7/73/Skeleton_basic_report.doc). The template defines a typical structure with the goal to uniform the reports but the final content will be flexible depending on the study.

The report is structured to fulfil three main goals. The first one is to identify the application structure and to measure the efficiency for the relevant region(s) of interest on the different components (load balance, synchronization, data transfer, IPC). The second goal is to describe the main characteristics of the behaviour for each of these components. Finally, the last section of the report summarises the results of the Audit. In this section we identify the areas where the user can work to improve the performance. The recommendations may suggest some aspects where a deeper study may be beneficial (Performance Plan) and even include suggestions for specific actions or code refactoring (Proof of Concept).

The audits carried out during the first months of POP suggest that it would be beneficial to define this phase as an iterative process where we share the results of the analysis with the user while writing the report so we can complement the report with additional input from the user. The majority of the POP Audits would be blind analysis of unknown codes so to involucrate the user on this phase would significantly improve the reports. At least one meeting/call should be organised previous to the report finalization to advance the results of the analysis and maybe gather additional information to make the report more focused to the user needs and concerns.



2.2 Performance Plan

2.2.1 Service identification and assignment

After a previous Performance Audit, the user can request to POP an extended assessment. As the user already accepted the POP Service Terms and Conditions there is no need to repeat the process and the starting point will be a meeting/call to discuss which of the suggested actions the analysis should focus. This will be a decision obtained in a consensus between the user and the POP expert.

In most cases the partner that provided the initial assessment will be in charge to extend the study. There may be cases where the collaboration of other expert is beneficial for instance because of his expertise on a given tool. In all cases we will try to maintain as first contact the person responsible of the Performance Audit.

2.2.2 Performance data collection

The extended analysis will require collecting new performance data. In this case the same approach of the Performance Audit will be followed.

2.2.3 Analysis and reporting

As in the Performance Audit, the POP experts will analyse the performance data collected and document the results in the Assessment Report. We have not defined a template for this report because the sections may depend on the targets for the Analysis Plan. Nevertheless in the next months and once we work on the first Performance Plan services we will try to identify commonalities in order to provide a uniform report.

This phase would be implemented as an iterative process where we comment the results of the analysis with the user while writing the report so we can complement the report with additional input from the user. Multiple meetings/calls are foreseen previous to the report finalization.

3. Implementation

We have decided to use a ticketing system to manage the POP services. After some discussions about the requirements (for instance free, integration with POP web site, previous knowledge by some partners) for this system and evaluation of available solutions, we selected TRAC as the POP ticketing system. While this deliverable is being written we have started to deploy TRAC and connect it to the POP web site.



When a user fills the Request Service Form, the web site will connect to the TRAC system opening a new ticket. The system will generate an identifier that would be used to connect this ticket with the additional information provided in the second questionnaire. TRAC should facilitate us a better control on the services status as well as obtaining global statistics with data from all the services.

As it was mentioned before, we have already started to partially use the procedure defined in this deliverable with the first assessments carried out during the first months of the POP project. One of the initial feedbacks from the current studies is that the elapsed time required for the analysis is somewhat larger than the planned schedule based on the average duration of the work, mainly due to delays obtaining the data or getting the required resources assigned (including human resources).

A second feedback is the benefit to discuss the identified results with the user during the analysis process and before the report is completed. This interaction with the users may also extend the duration of the service but allows gathering additional information to make the report more focused to the user needs and concerns. As most of the POP Audits would be blind analysis of unknown codes we believe that the report would be significantly improved with this interaction.

These two facts will not have an impact on the number of assessments targeted by POP but stresses the need to open more studies at the same time to guarantee we reach the desired throughput.

4. Forms, documents and templates

This section includes the forms, documents and templates referenced in the previous section with a short description justifying their contents and/or structure.

4.1 Request Service Form

This form has to be filled by the customer as a first step to request the service of POP. It includes the minimal questions identified to enable a certain guarantee that the proposed work can be served by POP. Since end of November it is available in the POP web site (<https://pop-coe.eu/request-service-form>).

Request Service Form

Contact Details

- Applicant's Name *
- Institution *
- e-mail *

Code



- Name of the code *
 - Scientific/technical area and class of problems it solves * (Aerospace, Archeology....)
 - Contribution * (Core developer / Module developer / User)
 - Access to sources *
 - Programming languages * (C, C++, Java, Fortran, Python, Others)
 - Parallel programming models * (MPI, OpenMP, OmpSs, Pthreads, CUDA, OpenCL, Others)
- Performance Service**
- Service request * (Code health check / Serial code performance / Parallel efficiency / Scalability / Communications / Identify areas of improvement / Others)
 - Describe your perception of the performance problem
- Terms and Conditions ***
- I accept the Terms and Conditions.*

4.2 POP Service Terms and Conditions

The POP Service Terms and Conditions have been minimized to cover the most relevant aspects of the service:

- POP is a European project offering free service to organizations based on the European Union. We are not a commercial service provider.
- The user agrees to give feedback on the results and grants POP with permissions for anonymous dissemination.
- Any publication originated from POP will acknowledge the project and the EC.

POP Service Terms and Conditions

POP is a European project financed by the European Union's Horizon 2020 research and innovation programme under grant agreement No 676553. POP is the Performance Optimisation and Productivity Centre of Excellence in Computing Applications and provides performance assessment, optimisation and productivity services. POP services are delivered free of charge to organisations based within the European Union.

POP is not a commercial service provider so there is not a commercial relationship between your organisation and the POP project partners, and commercial or civil rules between provider and customer are therefore not applicable.

By accepting the POP Service Terms and Conditions you apply to become a POP user. Applying does not guarantee that the service will be carried out, but POP commits to reply to all requests. If you become a POP user you agree to give us feedback on the quality of the service you obtain from POP, and grant us permission to publish statistical information on the percentage of potential performance improvement, the potential savings and other equivalent metrics that measure the results and impact of the POP service.



POP undertakes that all such data will be anonymised before publication, and that we will not publish any other information concerning the service that we are providing to you without your explicit permission.

Any publication originated from the POP services must include the following acknowledgment of the support received from the project and from the European Commission: "We gratefully acknowledge the support of the POP project which has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 676553, and has partially contributed to this work."

NO WARRANTY. THE POP SERVICES ARE MADE AVAILABLE "AS IS", WITHOUT WARRANTY OF ANY KIND. POP USER SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF SATISFACTORY QUALITY AND FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY AGAINST INFRINGEMENT OF ANY INTELLECTUAL PROPERTY RIGHT OF ANY THIRD PARTY OR ANY WARRANTY ARISING OUT OF ANY PROPOSAL OR SPECIFICATION. ALL THESE WARRANTIES ARE HEREBY DISCLAIMED.

4.3 Questionnaire on User Needs

The questionnaire on user needs collects more information about the code. It is divided in two sections; the first one includes additional questions required for a better assignment of the service to the POP partners. The second part includes additional details that are not mandatory for the assignment but that will help to characterize the study. There is no need to fill this second part before the assignment and the idea is that the information is completed during the study.

Questionnaire on user needs

Introduction

The objective of this questionnaire is to get a first overall description of the practices, perceived performance issues and needs. With this information, we will prepare an action plan with the suggested first steps and methodology of the service that we propose to you towards an in-depth analysis quantifying the causes of possible performance issues and suggesting ways to address them.

If you have different operation conditions for the same code (used in a fundamentally different physical problem, platforms, execution mode or setups), please fill one questionnaire for each of them.

Questions in Section I are optional.

Questions (additional to previous questionnaire)

- Select your preferred approach:
 - You will install the tools and obtain the required data for the POP staff
 - You will allow POP staff to access your system and you will provide the input data to run there the performance tests



- You would like to get access to PRACE /POP systems to install your code and you will obtain the required data for the POP staff
- You would like to get access to PRACE /POP systems to install your code and input data but you want that POP staff will run the performance tests
- Type of current system:
 - Processor type (Intel, Sparc, PowerPC, other)
 - System type (Cray, Bluegene, Generic Linux cluster, other)
 - Accelerators (GPU, MIC, other)
- Describe current development and production runs

	Number of cores/nodes	Elapsed time	#iterations
Production run			
Development test case			

- Code was analysed before (Yes/No)
If Yes:
 - Tools used for analysis: _____
 - Identified performance issues: _____
 - Can existing analyses be provided to POP staff (Yes/No)
- Familiarity with tools (Yes/No).
If Yes, Please select which one(s): Paraver, Scalasca, Vampir, VTune, Other
- Anything else that we should be aware of for considering your application

Section I. Additional Information:

Here we can collect specific details about the user, the program, the current (or target) platform used (to be used) to run the code, set up for typical production runs and additional information about previous analyses.

This information is not mandatory to decide who will do the analysis. It can be filled offline by the user or by the responsible of the analysis.

Code

- Is the code (standalone, coupled application, framework, library, etc.)
- License (GPL, BSD, commercial, other) [boxes?]:

Platform

- Typical system location (local cluster, on resources of an external provider (e.g. PRACE, Cloud...))
- Additional characteristics of the system: Number of cores, Memory per core, Network interconnect, I/O system type (Lustre, NFS), Operating system
- Available compilers in the system _____
- Available performance tools _____
- Working PAPI installation (Yes/No)
 - Version: _____



Production runs

- Scaling Target
 - Strong (fixed problem size, increase core count)
 - Weak (increase problem size and core count proportionally)
- Do you execute concurrently many instances of the same application (e.g. parametric studies)(Yes/No)
- Up to which numbers of cores are you satisfied with the performance of the program?

4.4 Assessment Report Template

The objective is that no matter which performance tool was used for the analysis, the kind of information provided to a user is structured in a very similar way. We defined a template for the report and the first experiences filling may identify potential improvements. We include here the current description of the template structure.

Performance analysis report

Application name:

Input data:

Performance study:

Application structure

This section will describe the application structure detected by the tools identifying the main temporal phases (initialization, iterative loop, etc.). If possible should include some views or tools captures. We may include brief details on how the data was obtained in case it was done by POP partners.

ROI (Region of interest)

It would be good to focus the analysis on a selected region to avoid that the initialization phase perturbs the metrics for the execution in cases where only few iterations are executed. This section will describe the region selected and its internal structure (number of iterations, phases within the iteration). The rest of the report will focus the results on the selected region.

Efficiency

The objective of this section is to report on the achieved efficiencies. Ideally we would like to report numbers between 0 and 1, the higher the better. BSC has a model that we can use as reference if the partners agree (part of it requires traces and Dimemas, but this part can be optional). For a pure MPI program the parallel efficiency is the average time outside MPI. In our model the parallel efficiency is decomposed in three main factors: global load balance, serialization/dependencies and transfer. Optionally we can include as measurement of the efficiency the percentage of time in MPI or the user function profile.



Load Balance

This section will report on the global load balance of the application. Even if the efficiency is good, we can indicate if there is a good balance with respect to the completed instructions, if there is some variability with respect to the IPC. The objective is to identify if we should recommend improving the load balance.

Serial performance

This section will report on the serial performance achieved. It may include the IPC of the main computing regions, correlation with other hardware counters to detect potential reasons for a low IPC. The objective is to identify if we should recommend to improve the serial performance and for which regions.

Communications

This section will report on the impact that the communications have on the application. Even if the efficiency is good, we can indicate if there are some MPI calls that achieve bad performance, the sensitivity of the application to the network characteristics, if asynchronous communications may improve the performance. The objective is to identify if we should recommend improving the communications.

Summary

This section will summarize the analysis results and will make some recommendations based on the previous sections. The recommendations may be related to code modifications or deeper analysis with the tools.



Acronyms and Abbreviations

- BSD – Berkeley Software Distribution
- CoE – Centre of Excellence
- CUDA - Compute Unified Device Architecture
- HPC – High Performance Computing
- GPL – GNU General Public License
- IPC – Instructions per Cycle
- MPI – Message Passing Interface
- NFS – Network File System
- OmpSs – OpenMP Superscalar
- OpenCL – Open Computing Language
- OpenMP – Open Multi-Processing
- PAPI – Performance Application Programming Interface
- POP – Performance Optimization and Productivity
- PRACE – Partnership for Advanced Computing in Europe
- Pthreads – POSIX Threads
- ROI – Region of interest
- WP – Work Package
- WPL – Work Package Leader