



Helix Nebula – The Science Cloud

Deliverable Title: D6.1 Best Practices Report

Partner Responsible: EGI.eu

Work Package: WP6

Submission Due Date: 30.09.2018

Actual Submission Date: 20.12.2018

Distribution: Public

Nature: Report



Abstract: This document outlines how the second part of the pilot phase of the HNSciCloud Pre-Commercial Procurement project has been organized. A set of dissemination and marketing activities to report the current status of the pilot services and promote their uptake in the scientific and business worlds have been also reported. Before the end of the project, several iterations between the procurers and contractors were necessary to develop accurate and meaningful TCO results to better determine the direct and indirect costs for supporting the PanCancer and the Alice reconstruction and analysis trains use cases. Additional lessons were learned during this phase based on the feedback retrieved from the Buyers Group. The document concludes with a set of recommendations and suggestions for possible future projects.

Document Information Summary

Deliverable number:	<i>D6.1</i>
Deliverable title:	<i>Best Practices Report</i>
Editor:	<i>Giuseppe La Rocca (EGL.eu)</i>
Contributing Authors:	<i>Bob Jones (CERN), Joao Fernandez (CERN)</i>
Reviewer(s):	
Work package no.:	<i>WP6</i>
Work package title:	<i>Pilots Evaluation and Recommendation</i>
Work package leader:	<i>EGL.eu</i>
Work package participants:	<i>CERN, CNRS, EMBL/EBI, ESRF, DESY, INFN, KIT, IFAE, SURFsara</i>
Distribution:	<i>Public</i>
Nature:	<i>Report</i>
Version/Revision:	<i>0.8</i>
Draft/Final:	<i>Draft</i>
Keywords:	<i>HNSciCloud, pilots evaluation, best practices</i>

Disclaimer

Helix Nebula – The Science Cloud (HNSciCloud) with Grant Agreement number 687614 is a Pre-Commercial Procurement Action funded by the EU Framework Programme for Research and Innovation Horizon 2020.

This document contains information on the HNSciCloud core activities, findings, and outcomes, and it may also contain contributions from distinguished experts who contribute to HNSciCloud. Any reference to content in this document should clearly indicate the authors, source, organisation, and publication date. This document has been produced with co-funding from the European Commission. The content of this publication is the sole responsibility of the HNSciCloud consortium and cannot be considered to reflect the views of the European Commission.

Grant Agreement Number: 687614

Start Date: 01 January 2016

Duration: 36 Months

Log Table

Issue	Date	Description	Author/Partner
V0.1	05/09/2018	First draft version available for partners' input and specific contributions	Giuseppe La Rocca/EGI.eu
V0.2	09/09/2018	First internal review	Bob Jones/CERN
V0.3	20/09/2018	Added more contributions after the M-PIL-3.3 event	Giuseppe La Rocca/EGI.eu
V0.4	20/10/2018	Second internal review	Bob Jones/CERN
V0.5	13/11/2018	Third internal review	Bob Jones/CERN
V0.6	12/12/2018	Added more contributions after the M-PIL-3.4 event	Giuseppe La Rocca/EGI.eu
V0.7	13/12/2018	Fourth internal review	Bob Jones/CERN
V0.8	20/12/2018	Extended abstract and executive summary	Giuseppe La Rocca/EGI.eu
V1.0		Final Version	

Executive Summary

During the last five months of the HNSciCloud Pre Commercial Procurement (PCP) project pilot phase (from June to Nov. 2018), the RHEA Group and T-Systems contractors were requested to progress, against the list of outstanding R&D activities, in order to deliver stable pilot platforms that could be used by the Buyers Group to deploy data-driven use-cases. After introducing the context of the HNSciCloud pilot phase in section 1, the report outlines how the phase was executed (section 2), with its main dissemination and marketing events (section 3) During the M-PIL-3.3 event in September the Buyers Group reviewed the status of the pilot platforms. The final status of the R&D activities was reviewed at the M-PIL-3.4 face to face event hosted by CERN on 28-30 November which marks the end of the pilot phase. In section 4, the Total Cost of Ownership (TCO) study is introduced to help the Buyers Group and contractors to determine the direct and indirect costs for supporting the PanCancer and the Alice reconstruction and analysis trains use cases. In section 6 this report presents the lessons learned, based on the feedback received from the Buyers Group, during the second part of the pilot phase of the HNSciCloud project. As part of the R&D activities, in section 5, is documented the procedure to register the initial pilot services developed by the project in the eInfraCentral service catalogue to improve the visibility of the commercial providers. The report concludes with a set of recommendations for future projects that are summarised below:

- Weekly meetings between Buyers Groups and contractors proved to be very fruitful to address the issues raised during testing of use-cases.
- The exchange of quotas between the Buyers Group on voluntary basis, proved to be essential to perform larger-scale tests.
- It is recommended to set-up a testing environment and procedures to validate software components in advance.
- The registration of the pilot services in the unified European Open Science Cloud on-line service catalogue potentially improves the visibility of the commercial cloud providers.
- Voucher schemes can promote the adoption of the new services to end-users for limited scale usage.
- Widely recognised performance metrics should be included as part of the assessment of cloud services.
- In future PCP projects, it is suggested that Total Cost of Ownership (TCO) studies should be a deliverable of each phase (design, prototype and pilot) with increasing accuracy as chosen use-cases progress towards deployment.

- The Buyers Group and contractors should monitor regularly the consumption of cloud resources and make adjustments accordingly to ensure all the objectives can be achieved.
- The Pre Commercial Procurement instrument does not adequately support the pay-as-you-go model for resource consumption.
- The stakeholders should reassess the total amount of resources required for each phase of the PCP project.
- Allocating sufficient resources during the design and pilot phase will simplify the execution of the pilot phase.
- The network requirements for data intensive use-cases should be established in-detail during the tender preparation phase.
- PCP projects should include the option to purchase the resulting pilots through a license to use the final developed solution.

Table of Contents

1. Introduction.....	8
2. Execution of the pilot phase (second part)	8
2.1. Kick-start of WP6 activities.....	9
2.2. Weekly WP6 teleconferences	10
2.3. Weekly teleconferences with the contractors.....	10
3. Dissemination events.....	11
3.1. CHEP 2018 Conference.....	11
3.2. GridKa School 2018.....	11
3.3. M-PIL-3.3 event	12
3.4. DI4R 2018	12
3.5. DESY administrators event.....	12
4. Total Cost of Ownership (TCO).....	13
5. Registration of the HNSciCloud pilot services	14
6. Lessons learned during the second part of the pilot phase	15
6.1. Objectives	15
6.2. Roll-out of Onedata releases.....	16
6.2.1. Release rc.10	17
6.2.2. Release rc.11	18
6.2.3. Release rc.12.....	18
6.2.4. Release rc.13.....	18
6.3. The HNSciCloud vouchers for service adoption.....	18
6.4. The HNSciCloud vouchers for the long tail of science.....	19
6.5. Resources capacity in the pilot phase	20
6.6. Exchanges of resources quotas.....	21
6.7. Prioritization of the R&D activities	21
7. Summary.....	21

1. Introduction

Stimulated by a Pre-Commercial Procurement commitment of leading research organisations from 7 countries, the HNSciCloud project pulled together commercial cloud service providers, publicly funded e-Infrastructures and resources in-house of 10 Procurers organisations - the Buyers Group - to build a hybrid cloud platform on top of which a competitive marketplace of European cloud players can develop new services. The project work-plan is broken-down in three successive and highly competitive phases – only contractors that successfully completed the previous phases were admitted to bid in the next one.

During the period January to June 2018 the project entered the pilot phase - the final step - in the implementation of the hybrid cloud platform proposed by the selected contractors. The first part of the pilot phase was coordinated by Work Package 5 (WP5) and led by INFN. The second part of the pilot phase was coordinated by Work Package 6 (WP6) and led by EGI.eu from June through to November 2018.

WP6 is composed of members of the following organisations belonging to the Buyers Group: CERN, CNRS, DESY, EMBL-EBI, ESRF, IFAE, INFN, KIT, STFC and SURFsara. At the end of the pilot phase, WP6 produced three deliverables:

- The present document (D6.1): Best Practices Report based on an evaluation of the results from the PCP and the best practices assessment from earlier phases: (D3.2 – Summary report of the design stage and lessons learned, D4.2 – Summary report of the prototype stage and lessons learned, and D5.2 – Summary report of the pilot stage and lessons learned).
- Demonstration of the resulting pilot services (D6.2).
- The roadmap for the implementation of a full-scale European Open Science Cloud (D6.3). This report aims to produce recommendations on how commercial services can be integrated and contribute to support the nascent EOSC. Inputs from the EOSC-hub Technical Architecture and standards roadmap (v1) will feed into this report.

2. Execution of the pilot phase (second part)

The figure below shows the timeline of the execution of the second part of the pilot phase, including the presentation of the HNSciCloud pilot services at CHEP (D-PIL-3.7) and the Total Cost of Ownership (TCO) deliverable for the two use-cases selected by the Buyers Group. Additional training events, not listed in the timeline, hosted at the HNSciCloud procurer’s premises are also reported in this document.

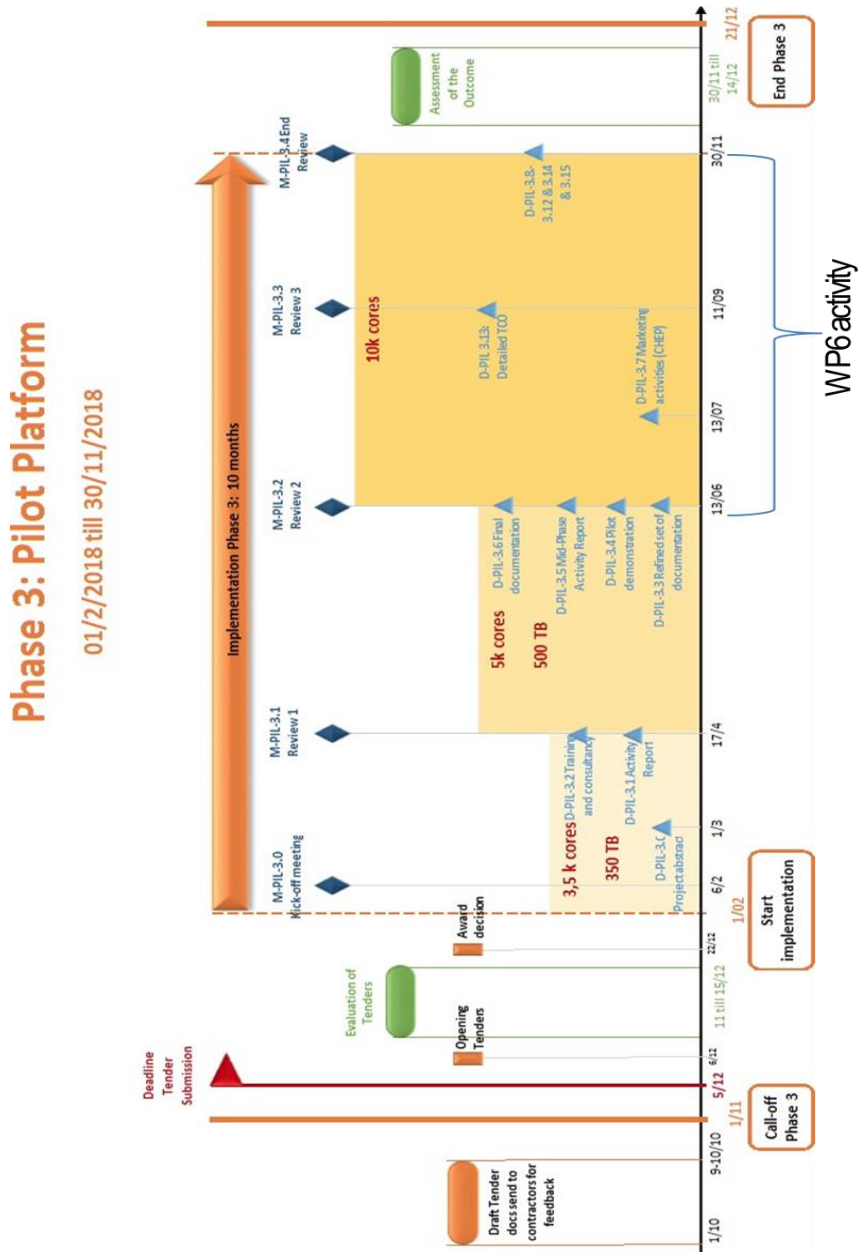


Figure 1: Timeline of the execution of the pilot phase (second part)

2.1. Kick-start of WP6 activities

The second part of the pilot phase started with:

- a face-to-face kick-off meeting and the second mid-term progress review (M-PIL-3.2) hosted at CERN (13-14 June),
- followed by marketing activities (D-PIL-3.7) to report the current status of the HNSciCloud pilot services at CHEP 2018 in Sofia (09-13 June),
- a training event hosted by KIT (28 August),
- the third progress review (M-PIL-3.3) hosted by SURFsara and EGI.eu (10-11 September),
- the DI4R2018 conference (9-11 October),
- a training for system administrators organized by DESY in Hamburg (24 October),
- and finished with the pilot phase review (M-PIL-3.4) hosted at CERN in November 2018.

During the M-PIL-3.3 event an assessment against the list of R&D activities documented in the M-PIL-3.2 feedback report and the Buyers Group test suite, including the addendum on Onedata tests, has been performed. The third progress review was also the deadline for the delivery of the deliverable: D-PIL-3.13 – Total Cost of Ownership study using the two use-cases proposed by the Buyers Group.

2.2. Weekly WP6 teleconferences

From the end of June till the end of the pilot phase, a weekly teleconference, meeting internal to WP6, was used to discuss project progress, coordinate testing activities, collect the quota allocation of resources for the next weeks, exchange experiences, prepare questions for the contractors and agree how to answer contractors' questions. Minutes of each meeting were produced and made available to all WP6 members.

2.3. Weekly teleconferences with the contractors

WP6 decided to follow the best practices identified during the previous phases of the project. Every week, immediately after the teleconference meetings internal to WP6, we had dedicated teleconferences with each contractor involved in the pilot phase.

These meetings usually started with a weekly report from the contractors' perspective, highlighting the status of progress activities to be accomplished in order to meet the upcoming deadlines and progress reviews. Afterwards, the Buyers Group were invited to report any additional feedback and issues encountered during the previous week.

Even if a dedicated mailing-list to interact with each contractor was already set-up, these weekly meetings offered the opportunity to establish a direct and fast link between contractors and the key contacts – members of the Buyers Group – involved with the testing

of the scientific applications. Overall, **weekly meetings with the contractors proved to be very fruitful to address the open issues and requirements coming from the Buyers Group.**

3. Dissemination events

3.1. CHEP 2018 Conference

The HNSciCloud pilot services have been presented to High Energy and Nuclear Physics experts and scientists attending the CHEP 2018 conference¹ in Sofia (09-13 July 2018).

The RHEA Group and T-Systems contributed to the event with two keynotes to showcase how the hybrid cloud developed by the HNSciCloud Pre-Commercial Procurement (PCP) project can support the high-performance and data-intensive use-cases.

The RHEA Group's keynote focused on the HNSciCloud Nuvla multi-cloud solution. Through this platform users can deploy Virtual Machines and/or Containers to multiple clouds and monitor their usage and cloud performance. To facilitate the deployment of VMs across multi-clouds the platform provides a brokering system which allow to choose which cloud to use based on price, performance, location or other factors which are important to them. Users may have their own orchestration tools and can optionally deploy directly to the clouds of their choice by using their native APIs. The platform supports the eduGAIN and Elixir AAI (SAML 2.0) identity federations, allowing users to access cloud resources via a web browser, Application Programming Interface (API) or Command Line Interface (CLI) – with access rights accorded by their unique identity. The Nuvla platform uses Onedata Data Management solution to allow data to be shared across multiple clouds as well as with local infrastructures.

The T-Systems' keynote, started with an overview of the performance and scale of use-cases that have been successfully deployed. Afterwards, it addressed how large-scale data can be processed in an intelligent way by pre-fetching the data or leaving the data remote at the existing infrastructures, making use of the state-of-the-art Onedata Data Management solution from Cyfronet. Lastly, the results of the new high level of transparency and budget control dashboard developed for the project were also demonstrated.

3.2. GridKa School 2018

¹ <http://chep2018.org/>

The HNSciCloud project organized a training session during the GridKa School 2018² hosted by the Karlsruhe Institute of Technology (KIT) (28 August 2018). During the training session Paulo Alexandre Canilho from CERN provided an overview of the current status of the project, its current status and business model and the voucher scheme for service adoption that the project is promoting to encourage the uptake of the pilot platform services for limited scale usage. The two contractors of the project attended the event contributing with hands-on sessions for their pilot platforms.

3.3. M-PIL-3.3 event

The event organized in Amsterdam at SURFsara's premises³ (10-11 September) coincided with the second pilot phase review. The event was a good opportunity for the two contractors to meet the Buyers Group and update them about the current status of the outstanding R&D activities and the pilot services. The status of the pilot services were further discussed by the Buyers Group during the WP6 session. The feedback collected has been used to assess the progress reported and drive the attention on the activities, with high priority, that have to be finalized before the end of the pilot phase.

On Sept. 11 the project has also organized a public session to provide all users, and potential newcomers, with an overview of the Onedata transparent data access solution.

3.4. DI4R 2018

The project contributed to the annual Distributed Infrastructure for Research conference⁴ with an oral presentation from T-Systems and a poster from the RHEA Group. The oral presentation focused on the latest project achievements and reported how the hybrid cloud platform can support the Buyers Group high-performance data-intensive scientific use-cases and the research sector at large.

3.5. DESY administrators event

The last procurer hosted event, before the end of the pilot phase, was the administrators event⁵ hosted by DESY in Hamburg (24 October). During the full day event, the HNSciCloud pilot services have been presented to system admins & DevOPS experts. Some use-cases developed at DESY were also presented at the event. During the event, the latest development activity with the Onedata software stack, with a particular focus on the new

² <http://www.kit.edu/english/index.php>

³ <https://www.surf.nl/en/about-surf/subsidiaries/surfsara/>

⁴ <https://www.digitalinfrastructures.eu/>

⁵ <https://indico.desy.de/indico/event/21675/>

functionalities to target the DESY use-cases, has been reported by Cyfronet. An additional presentation on the SLURM set-up, used to support the SURFsara use-case, has been reported by T-Systems, while RHEA Group presented how to run functions as a service and Terraform to run services.

4. Total Cost of Ownership (TCO)

Comparing the costs of cloud solutions vs. on-premises solutions is a complex and challenging task. To facilitate this assessment, the ECAR working group has created the Total Cost of Ownership (TCO) framework. The TCO framework addresses the following three main areas:

- **Foundational Risks:** These drive many of the considerations made in the TCO.
 - Data Sensitivity – how securely must the data be held and protected?
 - Business Criticality – how critical is the functionality to the business of the organization/project?
- **Quantitative Factors:** These are measurable costs that can be readily identified.
 - On-Going Costs.
 - One-Time Costs.
 - Hidden Costs and Subsidies (on-going, one-time).
- **Qualitative Factors:** These are factors that are hard to quantify in terms of Euros but can represent significant advantages or disadvantages for a solution.

The two contractors in the pilot phase have been tasked to produce the TCO study for the following two use-cases: PanCancer (supported by EMBL-EBI) and ALICE reconstruction and analysis trains (supported by CERN, CNRS, INFN, STFC and SURFsara). This study is a deliverable (D-PIL-3.3) scheduled for September 2018. For both use-cases the Buyers Group have sent to the contractors an initial list of requirements. With this initial list of requirements, and the follow-up discussions initiated with the Buyers Group, from June to August, the status of the TCO study has been further developed. Additional contributions/clarifications have been also provided to address the points raised by the contractors/buyers in order to allow the finalization of the study in the due time.

Before the M-PIL-3.3 event the two contractors produced an estimate of costs and a list of resources needs for the two use-cases. Overall, the TCO study focused on the costs of the solution, in terms of scaling of volumes (scale resources) and usage (scale of the usage profile), for supporting each use-case. Moreover, two variations of the PanCancer use-case to compare procurer on-premises hosted data and data hosted by the contractor using Onedata have been considered. During the M-PIL-3.3 event the Buyers Group had the

opportunity to discuss the details of the TCO studies during the closed sessions with both contractors and during the WP6 closed session.

The feedback from this event has been used by the two contractors to further refine the TCO study. Several iterations between the procurers and contractors were necessary to develop accurate and meaningful TCO results. The details of the TCO study using the two Buyers Group use-cases are described in D-PIL-3.13.

Overall, after all the iterations with the contractors and the Buyers Group, we concluded that **the network requirements for data intensive use-cases should be established in-detail during the tender preparation phase** and provisions for network connectivity and data ingress/egress taken into account. For this reason, future **PCP projects should include the option to purchase the resulting pilots through a license to use the final developed solution, including network access, after the project end.**

The TCO studies proved to be valuable to the procurers as input to their future IT strategies. **In future PCP projects, it is suggested that TCO studies should be a deliverable of each phase (design, prototype, pilot) with increasing accuracy as chosen use-cases progress to deployment.**

5. Registration of the HNSciCloud pilot services

The project has demonstrated how the PCP instrument can incite public and commercial providers to develop innovative services that can satisfy the needs of Europe's research communities. Thanks to these outstanding results, the HNSciCloud project has been highlighted by the EC High Level Expert Group as a concrete example of EOSC in practice, providing an innovative vision of how to develop capacity necessary to support the nascent EOSC intended to create a single digital research space for Europe's 1.8 million researchers. In order to explore how commercial cloud services can be integrated into the EOSC marketplace, the HNSciCloud contractors were requested to register their services in the service catalogue being developed by the eInfraCentral project. The overall goal of the eInfraCentral H2020 project⁶ is to structure open discussions between different e-Infrastructures aiming at defining a common catalogue for EOSC services. The new platform developed by the project acts as a gateway for end-users. Through this gateway users can browse the extensive catalogue of services and identify the provider matching their needs. As such, the eInfraCentral service catalogue offers a potential commercialisation channel and route into EOSC for the HNSciCloud pilot platform services.

The workflow to register a new service in the eInfraCentral Catalogue is the following:

⁶ <http://www.einfracentral.eu/>

- Use the eInfraCentral's Service Description Template (SDT)⁷ to collect information about the pilot services to be published.

The eInfraCentral project, through a consultation process, will:

- Help the service provider to achieve quality service descriptions.
- Provide an assessment report that may be used during project reviews and communications toward stakeholders.
- Incorporate the pilot services in the eInfraCentral Gateway and in the EOSC portal.

HNSciCloud contractors already produced the SDTs to describe the basic cloud compute services to be registered in the eInfraCentral catalogue. During a second iteration with the eInfraCentral team, the SDTs about the service pilots have been further improved before to be officially made visible in the eInfraCentral catalogue as follows:

- Nuvla Multi-cloud Application Management Platform⁸;
- Open Telekom Cloud (OTC)⁹.

The registration of the pilot services in the unified on-line service catalogue potentially improves the visibility of the commercial cloud providers.

6. Lessons learned during the second part of the pilot phase

This section describes different aspects of the execution of the second part of the pilot phase and identifies a number of lessons learned.

6.1. Objectives

The main objectives of the pilot phase were to:

- Assess the expanded prototypes deployed by the selected contractors;
- Open the pilot deployments to end-users so they can perform trials with their own applications.
- Provide the platform on which the final demonstrations can be performed.

⁷ <https://www.dropbox.com/s/dnrdrw5lnhlq1ip2/eInfraCentral-JNP-ServiceDescriptionTemplate.xlsx?dl=0>

⁸ http://catalogue.eosc-portal.eu/service/SixSq.nuvla_multi-cloud_application_management_platform

⁹ http://catalogue.eosc-portal.eu/service/OTC.open_telekom_cloud

6.2. Roll-out of Onedata releases

As documented in D5.2 – Summary report of the pilot stage: lessons learned, during the prototype phase, two of the three contractors adopted the same solution, built on top of Onedata software, to address the “transparent data access” PCP challenge. During the pilot phase, this software was considered to be not ready for large-scale and high-performance production usage raising reservations about the possibilities of this challenge being successfully fulfilled during the lifetime of the project.

Although the progress report registered in the first part of the pilot phase was promising in terms of computing resources utilization, the Buyers Group still were not able to use the full allocation of storage resources.

To address this issue, from the second mid-term review (M-PIL-3.2) onwards, additional effort has been put in place, by the two contractors, to improve the performance of the data management solution and mitigate the limitations identified by the Buyers Group during the WP6 weekly meetings. During the second progress review it was agreed to identify a third party, responsible for testing Onedata releases before they are rolled-out to the Buyers Group for deployment.

Starting from July the two contractors, in collaboration with Cyfronet, allocated additional resources to set-up a dedicated testing environment to validate the performance of Onedata. The testing environment used to simulate a hybrid cloud environment was composed by two Oneprovider environments and a Kubernetes cluster.

The Key Performance Indicators (KPIs) used to validate the new releases were discussed together with Cyfronet and members of the Buyers Group. To support the validation of new Onedata releases, INFN and DESY provided test-suites to be used during the testing phase. Since the management of HDF-files imposes several additional functional and performance challenges, particular attention was dedicated to test the DESY use-case.

In August, the two testing environments have been used to test the release rc.10 of Onedata taking into consideration the Buyers Group test-suites.

In September, rc.10 has been officially made available to the Buyers Group by the two contractors.

Before the end of the pilot phase, two additional releases: rc.11 and rc.13 have been made available to the Buyers Group. The reports about the Onedata releases are documented in the next sections.

Future PCP projects should foresee testing performed by a third party to validate the software components before they are made available to end-users. In addition,

allocating more resources to the design and prototype phase will simplify the pilot phase deployments.

6.2.1. Release rc.10

Following the new validation process, before the M-PIL-3.3 event (10-11 September 2018), T-Systems made available the new release of Onedata to the Buyers Group. This process included the production of a benchmark report with all the tests that T-Systems used during the testing and validation of the Onedata software.

During the testing and validation phase the performance of the software stack has been further improved with new features aiming to address issues with data caching and with the data import raised by the Buyers Group. The features included also a new set of APIs to monitor the data replication of files from the Buyers Group premises to remote cloud providers. From a technical perspective, these APIs notify the user when the file is correctly replicated. This API is particularly relevant for the DESY use-case since hundreds of files have to be processed by the application. Unfortunately, scalability tests were not included in the benchmark report provided by T-Systems and this prevented DESY from executing data-intensive applications.

The new release included a fix to prevent the locking issue with the ceph library. This issue prevented the PanCancer application (supported by EMBL-EBI) to read the reference file by all clients simultaneously. To validate the rc.10 release, the following two tests have been used by EMBL-EBI:

- Reading job-specific data files, via Oneprovider, with 1200 parallel clients;
- Configured clients for reading their common reference data files directly from Oneprovider, rather than having it installed locally on the node.

In both cases, a significant improvement of network performance (~10% aggregate) has been reported. Thanks to these improvements, EMBL-EBI completed the Onedata wave 2 tests.

INFN tests were postponed until the release rc.11 became available.

The RHEA Group reported several issues during the testing phase that delayed the roll out of the rc.10 release after the M-PIL-3.3 event. The new release, officially announced on 17 September, included new features and fixes to improve the overall performance and stability of the software stack as a whole. In more detail, the release:

- Improved the number of IOPS for data delivery.
- Replaced the ceph storage driver with the new RADOS driver.
- Added a fix to avoid performance degradation when multiple jobs were processing the same file.

- Improved the oneclient in order to log all its low level calls.
- Added a fix to force the update the size of the files during the process of syncing.

Unfortunately there were still several problems reported as known issues during the execution of scalability tests preventing a full exploitation (especially from DESY).

6.2.2. Release rc.11

Release rc.11 has been made available to the Buyers Group by the two contractors. With rc.11, DESY was able to verify all the functional tests (e.g.: data replication from the Onedata component to the DESY's premises), even if only with few files. A notable improvement of the stability and robustness of the Onedata software was reported by DESY in this phase. Additional scalability tests have been also prepared by DESY to evaluate whether this release could be used to run production jobs. The execution of these scalability tests indicated DESY was not able to run production jobs.

This release confirmed a notable performance increase and additional functionality, for the INFN scenario.

EMBL-EBI expressed interested to test this candidate release in an in-house deployment.

6.2.3. Release rc.12

This release was not made available on the pilot platform since it contained minor updates compared to the previous candidate release.

6.2.4. Release rc.13

This release was made available to the Buyers Group on November 30. This new release further improved the performance and includes WebDAV support. This feature is particularly important for Buyers Group members using dCache storage.

6.3. The HNSciCloud vouchers for service adoption

During the pilot phase the two contractors proposed a scheme by which vouchers can be distributed to end-users selected by the Buyers Group to encourage the uptake of the pilot platform services for limited scale usage.

The proposed voucher scheme should address the following points:

1. Each voucher should have a validity of one year from the date of issue and can be redeemed against any service included in the contractor's pilot platform.
2. A total of 100 vouchers of a value of 250 euros each would be sufficient for the Buyers Group tests.
3. The Buyers Group will identify the users entitled to access the pilot services.
4. It is essential that consumption is automatically blocked as soon as the voucher credit is exhausted.
5. It must be possible to add additional credit to an account after a voucher has been exhausted.
6. Data hosted by the cloud provider should be available for download even if the credit for an account has been exhausted.

The proposed voucher schemes were presented by the contractors at CERN during the M-PIL-3.2 event in June.

The Buyers Group considered the voucher scheme proposed by the RHEA Group to be flexible and with potential for wide use. The available documentation¹⁰ describes how the user can either create a new account using the voucher provided by Exoscale, or add credits to a voucher in case of an existing account. While the sign-up process for a new user and redemption of a voucher were clear, the Buyers Group raised some additional comments which have been promptly addressed by the contractor with a separate document.

Early adopters from SURFsara and EGI long tail of science started to test the vouchers distributed by the contractors during the M-PIL-3.3 event.

6.4. The HNSciCloud vouchers for the long tail of science

CERN and STFC decided to sponsor the long tail of science users consuming vouchers through the EGI Applications on Demand (AoD) service¹¹.

To facilitate the uptake of the HNSciCloud pilot services by the long tail of science users, EGI.eu, in collaboration with the Universitat Politècnica de València (UPV)¹², a partner in the EOSC-hub¹³ project responsible for the Elastic Cloud Computing Cluster (EC3)¹⁴ portal and

¹⁰ <http://hn-docs.readthedocs.io/en/latest/getting-started/exoscale.html?highlight=vouchers>

¹¹ <https://marketplace.egi.eu/42-applications-on-demand-beta>

¹² <https://www.upv.es/index-en.html>

¹³ <https://eosc-hub.eu/>

¹⁴ <http://servproject.i3m.upv.es/ec3/>

the Infrastructure Manager (IM)¹⁵ framework, extended the IM software stack developing two new modules to access the commercial cloud providers involved in the HNSciCloud PCP pilot phase. The activity started after the M-PIL-3.2 event has produced the development of additional IM modules to access the Open Telekom Cloud (OTC)¹⁶ provided by T-Systems based on OpenStack cloud middleware.

Thanks to these new IM modules, users of the EGI long tail of science users can use the pilot services to create elastic virtual clusters with some pre-configured applications. The list of application libraries/tools installed in the front-node can be exported via NFS in all the compute nodes.

6.5. Resources capacity in the pilot phase

T-Systems has progressively expanded the amount of resources offered during the pilot phase. The full capacity declared by the contractor includes: 10K cores, 1PB of storage and 40Gbps of network connectivity with Géant as stated in the original work order. Starting from July 2018 RHEA Group completed the ramp-up of the pilot capacity and a 40Gbps of network connectivity with Géant (aggregate) has been achieved. In addition, 18 GPUs were provided from July onwards. At the end of the pilot phase the Buyers Group were able to consume most of the allotted computing capacity in both contractors. The quota of storage capacity consumed was limited compared to the computing resources. This was in part justified by the delay in delivering a stable release of Onedata by the contractors.

In addition, to compensate for service outages, RHEA Group agreed to provide 4000 cores and 400 TB of storage over 3 months (1st Dec. 2018 – 28 Feb. 2019), while T-Systems agreed to provide additional support until 21 December 2018 for those Buyers Group members that were not able to complete their tests during the pilot phase.

Starting from the M-PIL-3.3 event, the two contractors enabled the Buyers Group to access their financial dashboards to monitor the amount of resources consumed. Several discrepancies have been reported by the Buyers Group during WP6 weekly meetings, and the contractors have promptly addressed.

The Buyers Group and contractors should monitor regularly the consumption of cloud resources and make adjustments accordingly to ensure all the objectives can be achieved.

The Pre Commercial Procurement instrument does not adequately support the pay-as-you-go model for resource consumption.

¹⁵ <http://www.grycap.upv.es/im/index.php>

¹⁶ <https://imdocs.readthedocs.io/en/latest/client.html#open-telekom-cloud>

6.6. Exchanges of resources quotas

During the second part of the pilot phase the Buyers Group agreed to temporarily exchange, on a voluntary basis, a fraction of their assigned IaaS resources. **This best practice offered the opportunity to perform larger-scale tests than would have been possible using their own quotas.** Adjustments between the procurers were agreed on a weekly basis and implemented by the two contractors.

6.7. Prioritization of the R&D activities

The Buyers Group agreed on a set of R&D activities to be implemented by the two contractors before the end of the pilot phase. This list of R&D includes:

- Further support of SLURM clusters at T-Systems.
- Implement full quota management for tenant/sub-tenant (RHEA Group and T-Systems).
 - It is essential that each procurer can assign quotas to individual user groups they sponsor.
- Release of Onedata rc.11 to target the technical requirements of the INFN and DESY procurers (RHEA Group and T-Systems).
- Address the comments from the Buyers Group and provide more accurate TCO studies (RHEA Group and T-Systems).
- Improve the set-up of the HPCaaS to address the issues reported by ESFR during the execution of MPI-based applications on both contractors (RHEA Group and T-Systems).
- The RHEA Group Financial dashboard: the financial dashboard has just been briefly demonstrated.
- GPUs: a Docker has been produced by the CERN team in order to be able to work-around the problem of the mpi_learn module. Action now on RHEA to achieve progress.
- The vouchers scheme provided by T-Systems is not yet mature and well defined for broader future use.

7. Summary

This section summarizes the lessons learned during the execution of the second part of the HNSciCloud pilot platform deployments:

- Weekly meetings between Buyers Groups and contractors proved to be very fruitful to address the issues raised during testing of use-cases.
- The exchange of quotas between the Buyers Group on voluntary basis, proved to be essential to perform larger-scale tests.
- It is recommended to set-up a testing environment and procedures to validate software components in advance.
- The registration of the pilot services in the unified European Open Science Cloud on-line service catalogue potentially improves the visibility of the commercial cloud providers.
- Voucher schemes can promote the adoption of the new services to end-users for limited scale usage.
- Widely recognised performance metrics should be included as part of the assessment of cloud services.
- In future PCP projects, it is suggested that Total Cost of Ownership (TCO) studies should be a deliverable of each phase (design, prototype and pilot) with increasing accuracy as chosen use-cases progress towards deployment.
- The Buyers Group and contractors should monitor regularly the consumption of cloud resources and make adjustments accordingly to ensure all the objectives can be achieved.
- The Pre Commercial Procurement instrument does not adequately support the pay-as-you-go model for resource consumption.
- The stakeholders should reassess the total amount of resources required for each phase of the PCP project.
- Allocating sufficient resources during the design and pilot phase will simplify the execution of the pilot phase.
- The network requirements for data intensive use-cases should be established in-detail during the tender preparation phase.

PCP projects should include the option to purchase the resulting pilots through a license to use the final developed solution.