Deliverable D4.5

Project Title:	World-wide E-infrastructure for structural biology
Project Acronym:	West-Life
Grant agreement no.:	**675858**
Deliverable title:	Report on the progress of the deployment of consolidated platform and its interactions with infrastructure
WP No.	4
Lead Beneficiary:	4: Masarykova Univerzita
WP Title	Operation and maintenance of the computing and data infrastructure
Contractual delivery date:	Month 26
Actual delivery date:	Month 26
WP leader:	Name: Ales Krenek / Partner: MU
Contributing partners:	STFC, EMBL-HH, MU, CSIC, CIRMMP, UU, Luna, INFN

Deliverable written by Marco Verlato, Robbie P. Joosten, Alexander Bonvin, Laura del Cano, Chris Morris, Antonio Rosato, Ales Krenek, Daniel Kouril

Contents

1 Executive summary

This document reports on the activities of the Work Package 4 (WP4). Building on a considerable legacy of existing operational web portals developed and used by the Structural Biology / WeNMR user community, WP4 aims at provisioning a consistent e-infrastructure to allow gradual integration of the existing isolated solutions to a single computing and data processing environment, based on the state of the art grid and cloud open source software tools and frameworks.

This report follows the document D4.3, delivered at project month 15, and the document M4.4 delivered at month 24, so that only the progress achieved until project month 26 not already described previously will be reported here.

The document starts with an updated description of the resources potentially available for the project from the EGI e-infrastructure, on top of which the consolidated West-Life platform is being built. It then presents a detailed view of resource usage and their geographical distribution in the second year of the project, as obtained from the EGI Accounting Portal. The remaining of the document reports the latest experiences and updates the activity plans on the consolidated job management mechanism, the programmatic access to datasets and the unified security and accounting model.

2 Project objectives

With this deliverable, the project has reached or the deliverable has contributed to the following objectives:

No.	Objective	Yes	No
1	**Provide analysis solutions for the different Structural Biology approaches**		X
2	**Provide automated pipelines to handle multi-technique datasets in an integrative manner**		X
3	**Provide integrated data management for single and multi-technique projects, based on existing e-infrastructure**	X	
4	**Foster best practices, collaboration and training of end users**		X

3 Detailed report on the deliverable

3.1 The HTC and Cloud production e-infrastructure

The consolidated platform of West-Life leverages the resources provided by the EGI. EGI is a federation of computing and storage resource providers united by a mission to support research and development. The EGI federated e-infrastructure is publicly funded (it has been supported by the EU project EGI-Engage until August 2017 and will be further supported within the next EOSC-Hub project starting in January 2018) and provides compute and storage resources to support research and innovation. As of November 2017, the EGI federation comprises over 300 data centres, located mostly in European countries, and a number of integrated resource providers in Canada, USA, Latin America, Africa and the Asia-Pacific region. A total of 730,000 CPU-cores are available for the High Throughput Computing (HTC) platform, while 7,000 CPU-cores are available for the Cloud platform (known as the EGI Federated Cloud). Furthermore, 300 PB and 385 PB are respectively available as online and archive storage.

In the following sub-sections we summarize the use made by West-Life project of both HTC and Cloud e-infrastructures, in term of the metrics which were selected as KPI for the entire Work Package.

3.1.1 HTC e-infrastructure usage

West-Life users get access to the EGI HTC e-infrastructure (also known as EGI Grid) via the enmr.eu Virtual Organisation (VO). The enmr.eu VO was established in 2007 in the context of the EU project e-NMR and further supported in the follow-up EU project WeNMR. Out of the 730,000 CPU-cores currently provided in total by the EGI grid, around 100,000 (211,000 at month 15) belong to the 30 (38 at month 15) resource centres that support the enmr.eu VO. This support is generally shared with many other VOs, so that the effective availability of resources to the project is difficult to estimate, but can be argued by the EGI accounting system. The 30 resource centres are distributed in 10 (13 at project month 15) countries around the world, and the figures 1 and 2 below show the total number of jobs and the normalized CPU time (in HEPSPEC06 hours) provided by each country in the second year of the West-Life project, from November 2016 until the end of October 2017. During 2017 EGI.eu and the MoBrain Competence Centre (a consortium represented by the Faculty of Science – Department of Chemistry of Utrecht University) have renewed the Service Level Agreement (SLA) signed in 2016, granting to enmr.eu VO for the period 1/1/2016 – 31/12/2017 an amount of opportunistic computing time up to 60 Million of normalized CPU hours and opportunistic storage capacity up to 265 TB [1]. A total of eight resource centres signed the SLA: INFN-PADOVA (Italy), IFCA-LCG2 (Spain), RAL-LCG2 (UK), TW-NCHC (Taiwan), SURFSara and NIKHEF (The Netherlands), NCG-INGRID-PT (Portugal), CESNET-MetaCloud (Czech Republic).

By signing the SLA the above resource providers committed to: operate 24x7x365 excluding planned maintenance windows or service interruptions which have to be notified via e-mail in timely manner (i.e. 24 hours before the outage); provide support through the EGI Service Helpdesk (http://helpdesk.egi.eu/), from Monday to Friday and from 9:00 to 17:00 in the time zone of the relevant Resource Centres;

VO Admin — Total number of jobs by Country and Half-year

Country	Nov 2016 — Apr 2017	May 2017 — Oct 2017	Total	Percent
China	18,975	39,312	58,287	1.67%
Germany	228	0	228	0.01%
Italy	42,604	95,899	138,503	3.97%
Netherlands	1,230,185	832,644	2,062,829	59.18%
Poland	51,940	47,734	99,674	2.86%
Portugal	45,033	41,122	86,155	2.47%
South Africa	6	0	6	0%
Taiwan	61,598	6,310	67,908	1.95%
United Kingdom	275,729	696,400	972,129	27.89%
Total	1,726,298	1,759,421	3,485,719	
Percent	49.52%	50.48%		

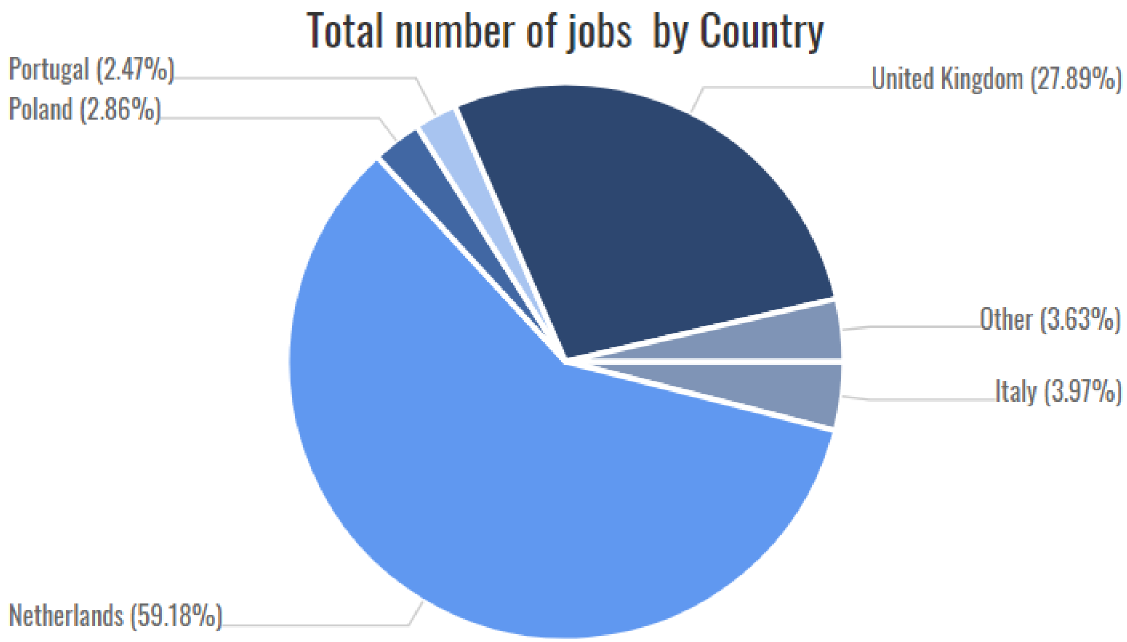

Figure 1: number of jobs by Country

VO Admin — Normalized CPU time (hours) by Country and Half-year

Country	Nov 2016 — Apr 2017	May 2017 — Oct 2017	Total	Percent
China	54,542	108,686	163,228	0.89%
Germany	939	0	939	0.01%
Italy	192,903	331,138	524,041	2.86%
Netherlands	5,922,216	4,660,663	10,582,879	57.69%
Poland	505,692	521,859	1,027,551	5.6%
Portugal	154,552	127,446	281,998	1.54%
South Africa	5	0	5	0%
Taiwan	108,469	31,684	140,153	0.76%
United Kingdom	1,509,450	4,112,842	5,622,291	30.65%
Total	8,448,767	9,894,317	18,343,084	
Percent	46.06%	53.94%		

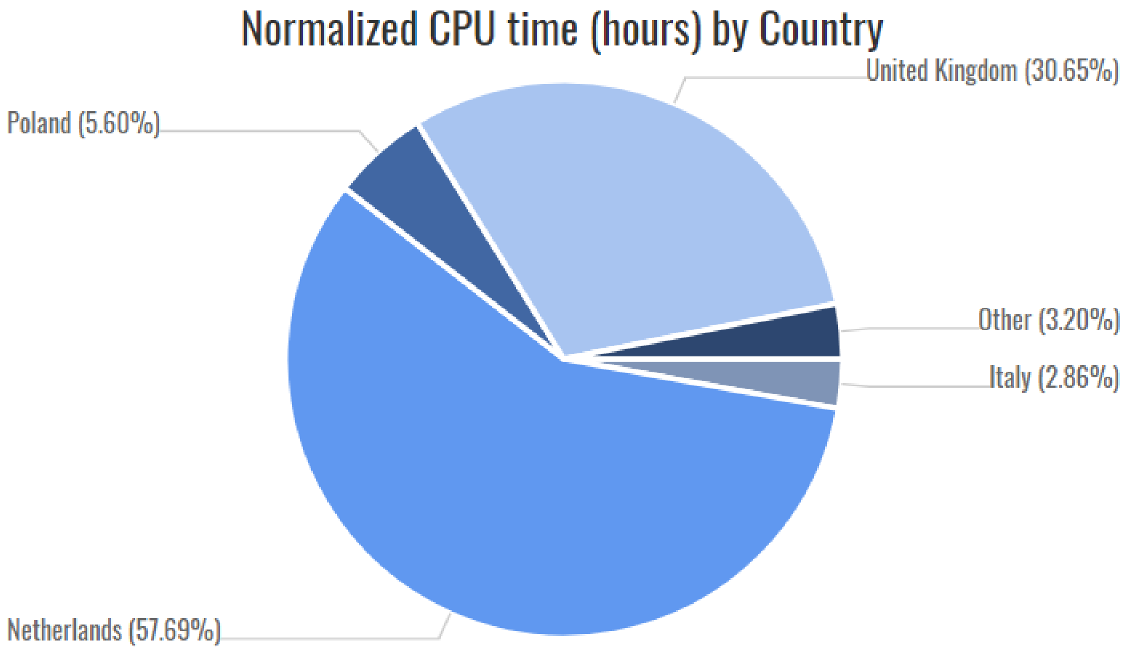

Figure 2: normalized CPU hours by Country

handle incidents according to the Quality of Support level that is estimated according to the impact of the outage or service quality degradation; target a minimum monthly availability (defined as the ability of a service or service component to fulfil its intended function at a specific time or over a calendar month) of 85% and monthly reliability (i.e. availability excluding scheduled maintenance periods) of 90%.

The above targets are constantly monitored by the EGI operations team who publishes the actual achieved values every six months. The figure 3 below shows the latest available values computed in the first six months of 2017. The only site significantly below the targets (values with orange background) for most of the period was IFCA-LCG2. In particular, although the Spanish site did sign the SLA, it was not possible to successfully submit jobs there due to configuration issues that were promptly reported to its site managers and were fixed only in December 2017. This explains why Spain is not yet reported in the accounting pictures above.

INFN-PADOVA	Service target	2017-01	2017-02	2017-03	2017-04	2017-05	2017-06
Availability	85%	100,00%	94,78%	100,00%	100,00%	93,40%	100,00%
Reliability	90%	100,00%	94,78%	100,00%	100,00%	93,40%	100,00%
IFCA-LCG2	**Service target**	**2017-01**	**2017-02**	**2017-03**	**2017-04**	**2017-05**	**2017-06**
Availability	85%	79,33%	86,74%	76,29%	73,54%	98,24%	49,61%
Reliability	90%	79,33%	86,74%	80,61%	80,45%	98,24%	50,31%
NCG-INGRID-PT	**Service target**	**2017-01**	**2017-02**	**2017-03**	**2017-04**	**2017-05**	**2017-06**
Availability	85%	100,00%	99,49%	97,07%	100,00%	100,00%	99,80%
Reliability	90%	100,00%	99,49%	97,07%	100,00%	100,00%	99,80%
NIKHEF-ELPROD	**Service target**	**2017-01**	**2017-02**	**2017-03**	**2017-04**	**2017-05**	**2017-06**
Availability	85%	98,13%	93,06%	99,20%	95,79%	100,00%	98,46%
Reliability	90%	98,13%	93,06%	99,20%	95,79%	100,00%	98,46%
RAL-LCG2	**Service target**	**2017-01**	**2017-02**	**2017-03**	**2017-04**	**2017-05**	**2017-06**
Availability	85%	99,04%	100,00%	100,00%	98,48%	98,43%	100,00%
Reliability	90%	100,00%	100,00%	100,00%	98,48%	99,41%	100,00%
SARA-MATRIX	**Service target**	**2017-01**	**2017-02**	**2017-03**	**2017-04**	**2017-05**	**2017-06**
Availability	85%	100,00%	98,62%	99,62%	100,00%	100,00%	96,17%
Reliability	90%	100,00%	100,00%	100,00%	100,00%	100,00%	100,00%
TW-NCHC	**Service target**	**2017-01**	**2017-02**	**2017-03**	**2017-04**	**2017-05**	**2017-06**
Availability	90%	100,00%	100,00%	100,00%	89,19%	95,18%	92,72%
Reliability	90%	100,00%	100,00%	100,00%	92,17%	95,18%	92,72%
CESNET-MetaCloud	**Service target**	**2017-01**	**2017-02**	**2017-03**	**2017-04**	**2017-05**	**2017-06**
Availability	85%	100,00%	98,93%	96,66%	100,00%	94,11%	98,44%
Reliability	90%	100,00%	98,93%	96,66%	100,00%	94,63%	98,44%

Figure 3: site availability/reliability summary

According to the EGI accounting system more than 3.4 Million jobs (3.3 Million in the first year) have run, translating into 18.3 Million (19.6 Million in the first year) normalized CPU hours consumed by the West-Life applications. In both metrics we can notice a quite stable usage when comparing with the previous statistics published in D4.3, despite the reduction in the number of opportunistic sites supporting the project. Most of the jobs in fact now run on the resource centres who signed the SLA, with a foremost contribution of the ones operated in Netherlands (~60%) and UK (~30%). We should note here that jobs are not being targeted to specific sites and the distribution of jobs is left to the workload management system used (WMS and DIRAC4EGI). This explains why a fraction of the jobs still are executed e.g. by resource centres in Poland and China.

As pointed out already in D4.3, last year two resource centres, CIRMMP in Italy and Queen Mary University of London in UK, have made a total of 6 NVIDIA Tesla K20m, 4 NVIDIA Tesla K80, and 1 NVIDIA Tesla K40c GPGPU cards available. These nodes could be included in the HTC e-infrastructure through a prototype developed by INFN (in collaboration with EGI-Engage project) to enable GPGPU support in CREAM-CE, and allowed AMPS-NMR, DisVis, PowerFit applications to greatly enhance their performance. In December 2017 a new version of CREAM-CE with this GPGPU support has been released in production for UMD (Unified Middleware Distribution) 4.6.0 for CentOS7, in coordination with the EGI Software Provisioning team. The release has been announced at the EGI Operation Management Board, so that a wider adoption by the EGI resource centres interested in exposing their GPGPU resources on the HTC e-infrastructure is now possible. This work has been presented at the ISCG 2017 Conference, held in Taipei, Taiwan, the 5-10 of March [2] and at the EGI 2017 Conference / INDIGO 2017 Summit held in Catania, Italy, the 9-12 of May [3].

3.1.2 Cloud e-infrastructure usage

The EGI Federated Cloud as of November 2017 provides a total of 7,000 CPU-cores. 624 of them belong to the 3 resource centres that support the enmr.eu VO: these are CESNET-MetaCloud in Czech Republic, INFN-PADOVA-STACK in Italy and IISAS-GPUCloud in Slovakia. The figures 4 and 5 below show the total number of Virtual Machines (VM) and their elapsed time (in hours) provided by each country in the second year of the West-Life project, from November 2016 until the end of October 2017. More than 850 VMs (1,300 in the first year) and 107,000 CPU wall time hours (151,000 in the first year) have been consumed by the West-Life applications (mainly by Scipion, Haddock, DisVis, PowerFit and Gromacs).

The cloud instances hosted at CESNET-MetaCloud and IISAS-GPUCloud sites are equipped respectively with 4 NVIDIA Tesla M2090 and 4 NVIDIA Tesla K20m GPU cards, allowing to instantiate VMs with one or more GPU devices attached via PCI passthrough virtualisation, a solution implemented by IISAS laboratory in the context of the EGI-Engage project. In the last year the Scipion application has been successfully deployed on these two cloud sites (see section 3.2.2.2), with the goal of using GPUs for computation and volume rendering. The results of the initial tests, and a performance comparison with an analogous deployment on the AWS EC2 commercial cloud were presented at the joint EGI Conference 2017 / INDIGO Summit 2017 [4]. At time of writing this document, Masaryk University started the installation of a GPU-enabled cluster (8 nodes with 2 NVIDIA 1080Ti cards each) which will be available (though not entirely exclusively) to West-Life services.

Cloud VO Admin — Total number of VM run by Country and Half-year

Country	Nov 2016 — Apr 2017	May 2017 — Oct 2017	Total	Percent
Czech Republic	234	382	616	72.05%
Italy	133	7	140	16.37%
Slovakia	22	77	99	11.58%
Total	389	466	855	
Percent	45.50%	54.50%		

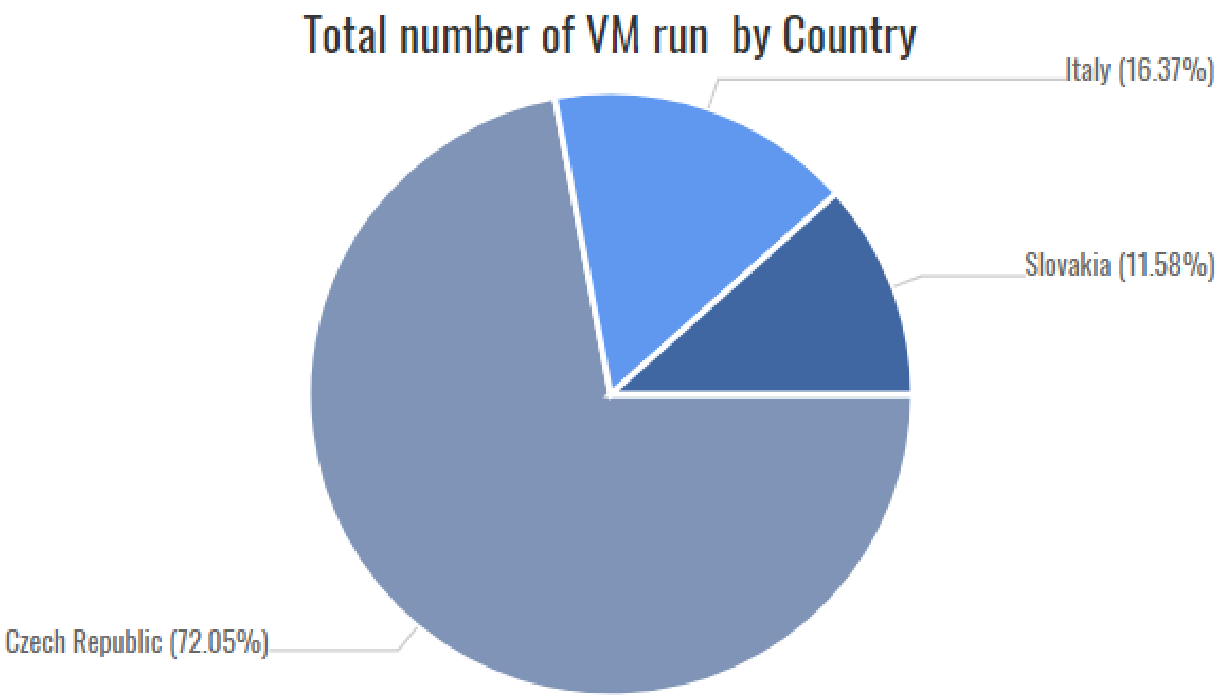

Figure 4: number of VMs by Country

Cloud VO Admin — Total elapsed time (hours) by Country and Half-year

Country	Nov 2016 — Apr 2017	May 2017 — Oct 2017	Total	Percent
Czech Republic	35,375	51,606	86,981	80.85%
Italy	5,259	216	5,475	5.09%
Slovakia	2,530	12,601	15,131	14.06%
Total	43,165	64,423	107,587	
Percent	40.12%	59.88%		

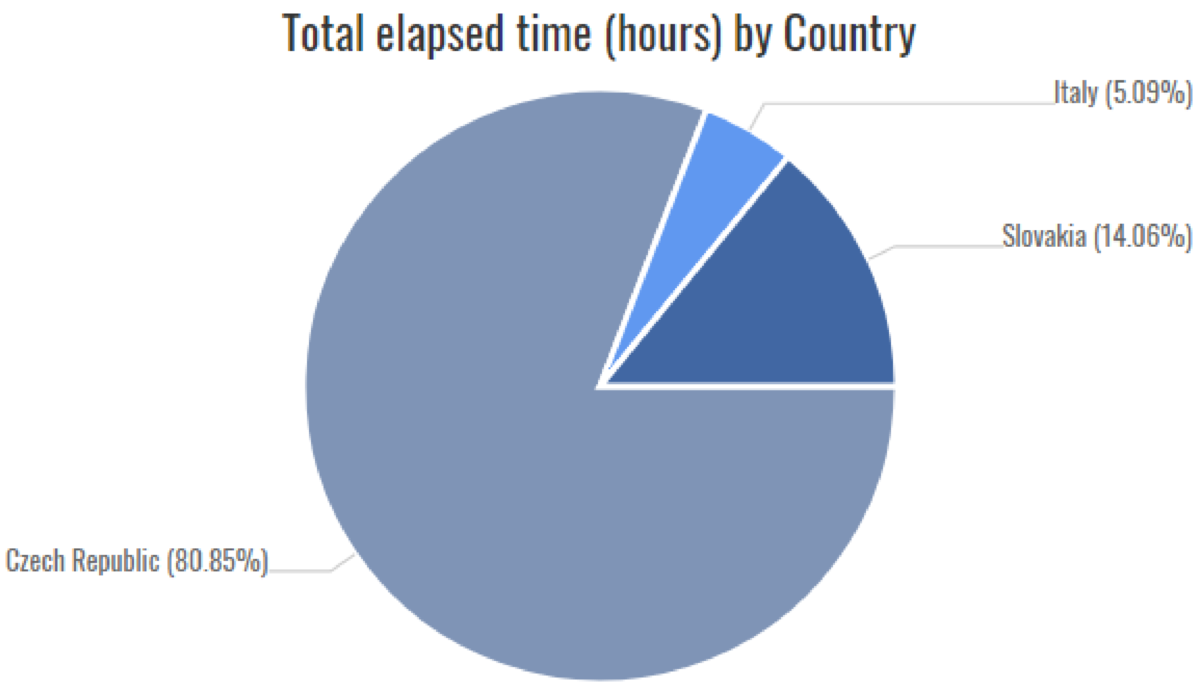

Figure 5: cloud elapsed time by Country

3.1.3 KPIs summary

The table below summarizes the time evolution of the KPIs related to resource capacity and usage.

Metric	Project year 1	Project year 2	Difference (%)
CPU-cores (best effort)	211,000	100,000	-53%
CPU-cores (SLA)	6280	6850	+9%
Storage (SLA)	215 TB	265 TB	+23%
No. of HTC jobs	3,385,457	3,485,719	+3%
HTC Norm. CPU hours	19,667,725	18,343,084	-7%
No. of Cloud VMs	1325	855	-35%
Cloud elapsed hours	151,306	107,587	-29%

Concerning HTC metrics, we can notice a quite stable usage (slightly more jobs executed but less CPU time consumed overall) despite the strong reduction in the number of CPU-cores provided by the opportunistic sites supporting the project. More than 90% of the jobs nowadays run in fact in the resource centres that signed the SLA and committed a fixed amount of resources. Therefore, the visible drop of best-effort resources does not indicate a problem. On the contrary, it is an evidence of consolidation of the EGI infrastructure and its usage pattern. Cloud metrics show instead a significant reduction, but we have to consider here that small absolute numbers are involved, e.g. the total elapsed hours in the project year 2 correspond to 3 VMs with 4 CPU-cores continuously running for one year. It also corresponds to about the 5% only of the total HTC usage. Moreover, unlike the HTC, the cloud statistics include idle time of the allocated resources. Therefore, the decrease of absolute resource usage indicates more careful approach to resources while testing as well as higher maturity of the cloud deployment mechanisms, which prevents resource wasting. There is currently one production portal running on cloud resources, Scipion Web Tools, while other Scipion services deployed on the EGI Federated Cloud and the new generation of the Gromacs portal are currently being used more for testing / development purposes and for workshops, as pointed out in the next section 3.2.2.2.

3.2 The consolidated platform on top of the e-infrastructure

3.2.1 Consolidation of job management mechanism

The common job management mechanism behind the application portals, which distributes the highly demanding payloads across the HTC e-infrastructure, has been widely described both in deliverables 4.1 and 4.3, and in milestones 4.2 and 4.4. Therefore we will mainly report here the latest updates, while adding some more details to complement the shorter milestone documents.

3.2.1.1 DIRAC4EGI service

DIRAC (Distributed Infrastructure with Remote Agent Control, http://diracgrid.org/) was the software chosen to implement the West-Life job dispatcher, and the DIRAC4EGI service is a cluster of instances, operated by EGI e-Infrastructure, to allow any user belonging to an EGI supported VO to distribute computational tasks among the available EGI resources, monitor their status and retrieve the results. The migration to DIRAC has become more important because in June the EGI Operations team has announced that the gLite/EMI WMS will be decommissioned at the end of 2017, and all the instances still active around the world will be removed from production after that time. HADDOCK portal was pioneering this migration in the past, but now all other West-Life grid-enabled application portals still relying on the gLite/EMI WMS should migrate to DIRAC as soon as possible. To achieve this goal, a one-day workshop was organized at the West-Life All Partners Meeting in January for the portal developers. After that, the current status about the adoption of DIRAC as underlying job submission engine is the following:

- The CS-Rosetta3 portal, which was originally making use of job submission via gLite/EMI WMS has recently been migrated to DIRAC and put back into production.
- The migration of CIRMMP portals to DIRAC4EGI will take place in January 2018.

3.2.1.2 Application Software Deployment

No relevant news with respect to last D4.3 concerned the software deployment mechanism. Two West-Life Virtual Appliances already published in EGI Application Database were updated in May:

- ScipionCloud was updated to version 1.1_beta [5] to include the GPGPU support.
- The West-Life VM Appliance, containing the Virtual Folder consolidating multiple storage providers and software ready for structural biology by supporting WEBDAV, EUDAT, B2DROP and Dropbox, was updated to version 17.05 [6]. It comes with two different images: one supporting VirtualBox hypervisor and suited for OpenNebula based clouds, and the other one supporting KVM hypervisor an suited for OpenStack based clouds.

Other West-Life software distributed via CVMFS system at the repositories /cvmfs/wenmr.egi.eu, /cvmfs/west-life.egi.eu and /cvmfs/facilities.gridpp.ac.uk was aldo updated when necessary.

In September the INDIGO-DataCloud project ended. However, most of its software products will continue to be supported in the context of the next H2020 EOSC-Hub project starting in January 2018. Docker images for some of the West-life applications (AmberTools, DisVis and PowerFit) will then continue to be hosted and kept updated at the INDIGO-DataCloud dockerhub repository [7].

Based on the docker containers used for DisVis and PowerFit, INDIGO and UU team implemented a continuous integration pipeline as proof-of-concept [8] . The implementation provides a test and validation suite that ensures the viability of any change in the application's source code, by the automatic and sequential execution of the 1) code inspection, 2) software packaging, and 3) application validation tasks.

The *source code inspection* tries to detect issues at early stages in the development phase, otherwise tougher to resolve once the software has been released. Any type of source code testing (style, unit, functional, integration) can fit at this stage. The quality assurance checks executed for the prototype are devoted to guarantee the readability and sanity of the code by making it compliant with Python's PEP8 standard. The *application packaging* phase deploys the application, packages it as a container image and publishes it to the INDIGO-DataCloud's online repository. The uploaded image is tagged with the current code version, so it does not overwrite any previous one. The recently created version is then fetched and tested as the final step in the pipeline. The *application validation* phase executes the container with a set of defined inputs and the results obtained are programmatically compared with reference results.

3.2.2 Deployment of applications in the Cloud e-infrastructure

Several steps forwards have been made during the last months regarding the portal virtualization activity and the experimental work to set up Gromacs and Scipion applications directly in the EGI Federated Cloud via cloud orchestration services. Some of these were already described in short on the M4.4 document published at month 24. We will report them here with more details and with the very last updates.

3.2.2.1 Gromacs

The grid-based Gromacs service hosted by the Utrecht partner is currently suspended since April following a hardware failure. It is being replaced by the cloud version with fully automated deployment. The portal code was redesigned to submit jobs to a virtual cluster created by the cloud deployment, and the code implementing the molecular dynamics protocol was upgraded to work with Gromacs version 5.x. The implementation also uses GPU transparently, when it is available. User authentication was switched to the new common West-Life AAI scheme. The architecture was described before at the wiki page [9]. At time of writing this document, the code is fully functional, design of the interface is being adapted to West-Life skinning. The instance can be accessed at the link [10]. Currently it is set up with access to limited resources only, after installation of the new MU GPU cluster (beginning of 2018) the service will be redeployed there.

3.2.2.2 Scipion Web Tools

As reported in M5.3, Scipion Web Tools portal is now deployed and running on the EGI Federated Cloud, at CESNET-MetaCloud site. This service is accessible from the West-Life portal under Services/Electron Microscopy/Scipion Web Tools.

Besides, two full Scipion testbeds are deployed on the EGI Federated Cloud, one GPU enabled at IISAS-GPUCloud site and one without GPU at CESNET-MetaCloud site. These installations are being successfully used for demonstration and training purposes and are based on the architecture developed by MU on its Cloud Orchestration Service, that leverage the use of a remote GPU for 3D volume rendering, as described on M4.4.

Along this line MU and Scipion team are currently working on a full automatic deployment of the Scipion application using this service. This work could be used to set up on-demand machines assigned to a CryoEM project. Snapshots of the portal and its use on GPU resources for 3D rendering are shown in figures 6 and 7 below.

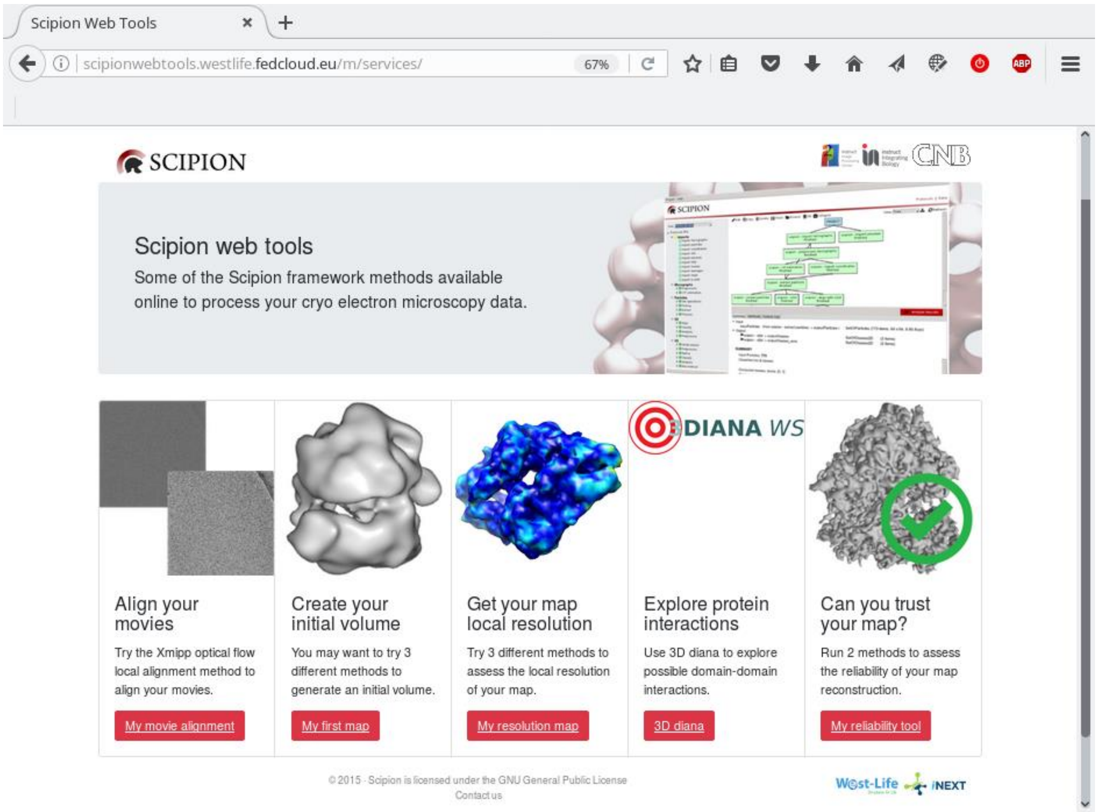

Figure 6: Scipion Web Tools portal deployed at CESNET-MetaCloud site

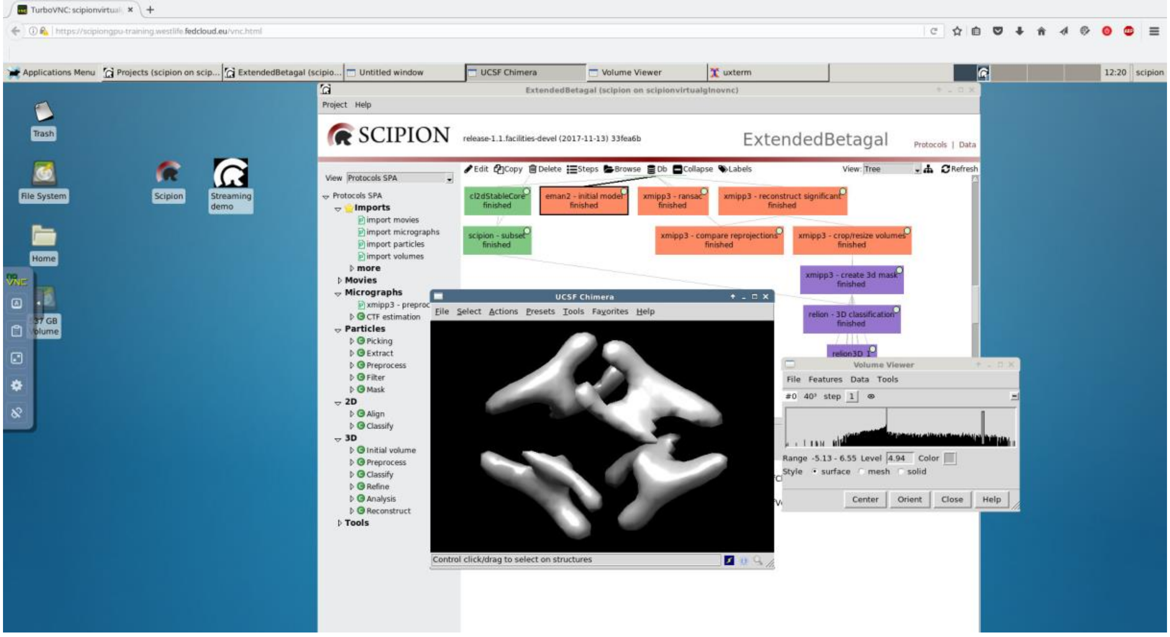

Figure 7: Scipion GPU testbed at IISAS-GPUcloud site. 3D rendering making use of remote GPU

3.2.2.3 AMPS-NMR

We developed a cloud-based service to analyze the trajectories output by AMPS-NMR. The service is currently using the Onedata storage developed within the INDIGO-DataCloud project. However, longer trajectory files are sometimes corrupted. We are now working to link this to the VirtualFolder instead. The prototype of the service runs properly on the local OpenStack cloud at CIRMMP and has been demonstrated in the context of INDIGO.

3.2.3 Programmatic access to datasets

The activity of the task 4.3 of the Work Package 4 is aiming at defining an architecture and the appropriate interfaces to access the relevant biological datasets, including strategies to cache copies of the data across the distributed e-infrastructure. This activity has been reported in deliverable D4.4: *Overview of external datasets, strategy of access methods, and implications on the portal architecture*, and in the milestone document M4.3: *Prototype access to selected dataset* published respectively at month 18 and month 22, so it will not be further reported here.

However, as pointed out in deliverable D4.3, in November 2016 INFN started to deploy Onedata, an INDIGO-DataCloud software component designed to support efficient, high performance and scalable distributed data caching mechanisms, integration with existing storage infrastructures, and transparent client-side import/export of distributed cloud data.

After having procured a Dell MD3600 storage server with 12 x 2TB disks, INFN installed Onedata software at INFN-PADOVA-STACK site of the EGI Federated Cloud and made it available to West-Life users in July. West-Life users can obtain a storage space by login at the Onezone server hosted at INFN-CNAF data centre. In the initial implementation, West-Life users can at first register with the IAM service instance provided by the INDIGO infrastructure, and then use it as authentication method for automatically obtaining an account on the Onedata system. The use of the West-Life SSO mechanism described in the next section as additional authentication method will be implemented early 2018. The documentation with the instructions for West-Life users on how to use the Onedata storage is available at the West-Life main web site.

3.3 Unified security and accounting model

Progress in AAI activity achieved in the last months focused on proper support of identity and group management for the West-Life community, and were shortly described in the milestone document M4.4. A training session tailored for service operators presenting the main principles of integration and giving information about how to enable support for federated authentication and SSO using SAML was held at the All Partner's meeting in Amsterdam at the end of August. The overall design and implementation was also presented at the DI4R Conference held in Bruxelles the 30 November - 1 December [11].

The AAI architecture follows recommendations of AARC2 project "Blueprint" architecture. Its major building blocks and their interactions are shown in Figure 8. The implementation is based on components that are used in AAI of Elixir and BBMRI infrastructures, and which are expected to form the emerging common "life science AAI".

In this scenario, user credentials come from various sources: institutional identities associated in EduGAIN are preferred, existing ARIA and legacy WeNMR accounts can be used. Social identities (Facebook, Google, Linkedin etc.) are supported as well. The system provides *identity consolidation*, i.e. a user may use either his/her university identity as well as Facebook to access the same services, after the identities are confirmed to be merged explicitly.

Access control is enforced via membership in groups, which are authorized to use individual services (at differing levels eventually). For management of groups we use Perun IAM system, developed jointly at Masaryk University and CESNET [12] . Legacy WeNMR groups are inherited for the transition period.

The central component of the architecture is the *IdP-SP-Proxy* developed before. This component simplifies the relationships between services and identity providers from the "many to many" model to "one to many" only. Moreover, Proxy was enhanced to support the Perun system so that attributes about the user can be made available to end services.

Following user's consent to release them, the Proxy sends a set of basic attributes (user's name, email address, unique identifier) that are assigned to the user. A list of groups that the user is member of is also released as part of the attribute set, which allows the services to enforce authorization. Due to the identity consolidation, the user is still recognized as a single person by the end service, regardless of the credentials he/she chooses.

The utilization of the Proxy made it possible to perform basic checks even before a user accesses an end service. When the user is accessing a service, the Proxy first checks that the user is already registered with the West-Life community maintained by the Perun IdM service. Users who are not known to the system are requested to either register or link their current identity with an existing user record. Users who are already registered with Instruct will be provided with a straightforward procedure to get members of the West-Life user base that will not require additional approval by community maintainers.

The WP4 team provided a guide for users [13] describing typical operations. There is also documentation intended for service operators [14] with information about how to enable support for federated authentication and SSO using SAML.

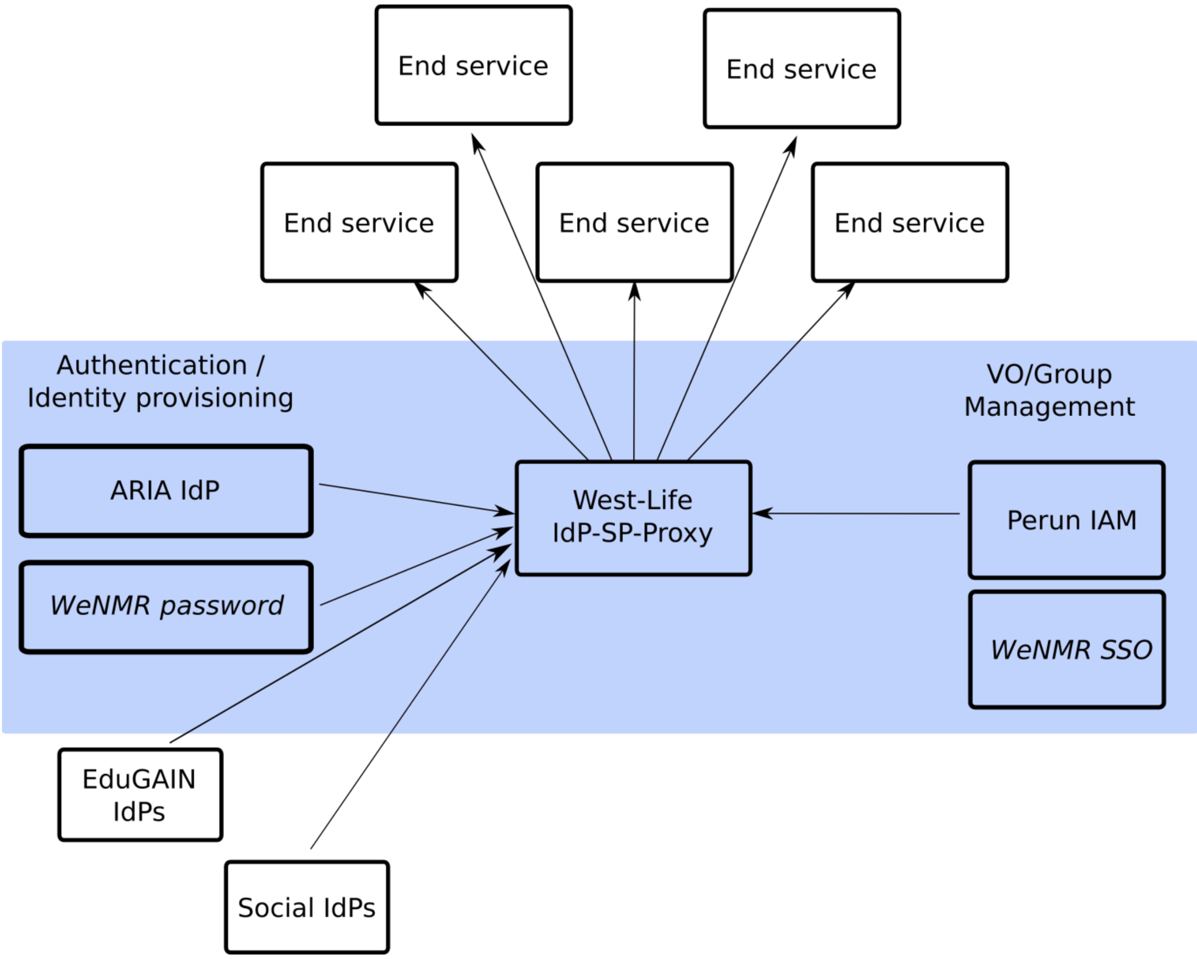

Figure 8: block diagram of West-Life SSO

As a first result, the INFN-PADOVA-STACK OpenStack based cloud instance of the EGI Federated Cloud implemented the West-Life SSO as additional authentication mechanism. West-Life users with a minimum of familiarity with the OpenStack Horizon Dashboard can now login with their West-Life credentials and create their own Virtual Machines. The figures 9 and 10 below show a snapshot of the authentication box where the user can select the West-Life SSO method, and the "Instances" view of the Horizon Dashboard with the mapped West-Life user ID on top right of the screen.

The Gromacs portal described above also uses West-Life SSO, migration of DisVis and PoweFit is on the way.

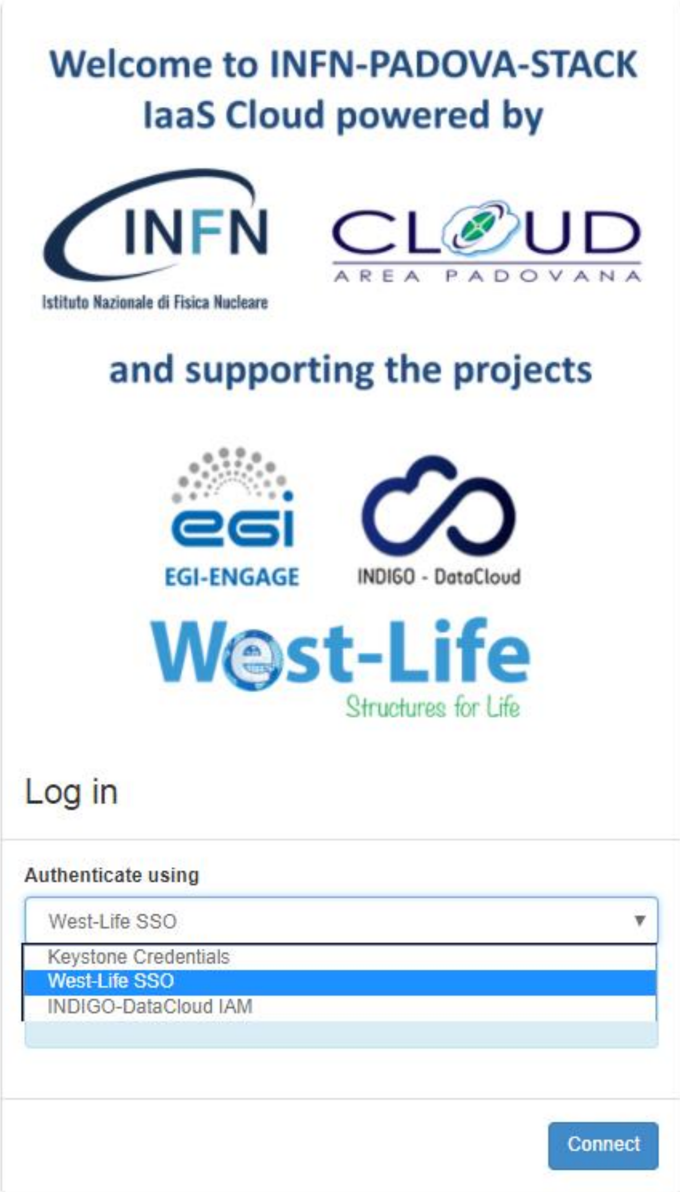

Figure 9: snapshot of authentication box for accessing INFN cloud

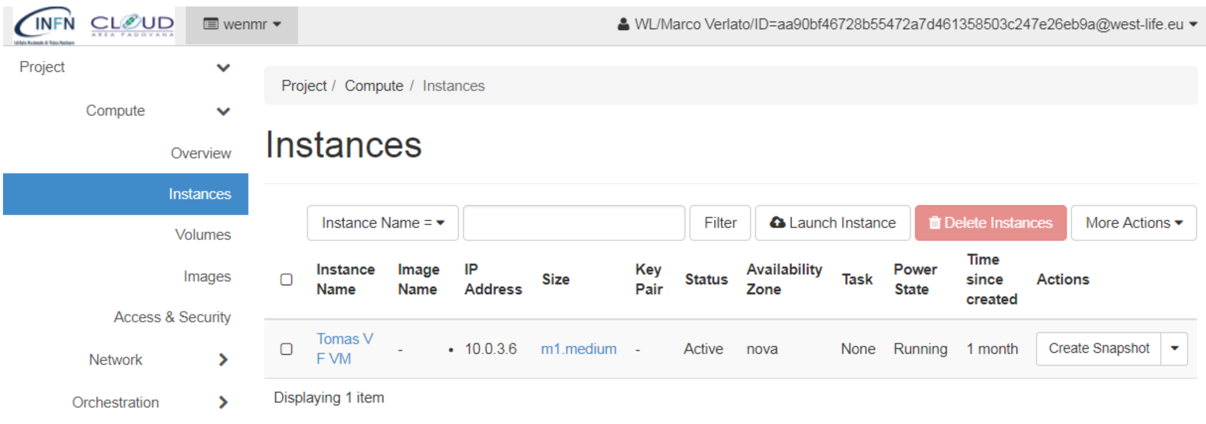

Figure 10: view of the Horizon dashboard for managing VMs on the INFN cloud

References cited

[1] Updated VO SLA MoBrain and VO OLAs MoBrain, https://documents.egi.eu/document/2751

[2] EGI federated platforms supporting accelerated computing, PoS(ISGC2017)020, https://pos.sissa.it/293/020/pdf

[3] GPGPU computing support on HTC, , talk at EGI Conference 2017/INDIGO Summit 2017, https://indico.egi.eu/indico/event/3249/session/30/contribution/16

[4] ScipionCloud: Large scale cryo electron microscopy image processing on commercial and academic clouds, talk at EGI Conference 2017/INDIGO Summit 2017, https://indico.egi.eu/indico/event/3249/session/30/contribution/179

[5] Scipion Appliance: https://appdb.egi.eu/store/vappliance/scipion.v1.0/vaversion/latest

[6] West-Life VM Appliance for VirtualBox: https://appdb.egi.eu/store/vappliance/d6.1.virtualfoldervm/vaversion/latest,
West-Life VM Appliance for KVM: https://appdb.egi.eu/store/vappliance/west.life.vm/vaversion/latest

[7] Dockerhub repository: https://hub.docker.com/u/indigodatacloudapps

[8] https://jenkins.indigo-datacloud.eu:8080/view/pipe-apps-disvis/

[9] http://internal-wiki.west-life.eu/w/index.php?title=WP4_Portal_Virtualization#Gromacs

[10] http://gromacs.westlife.dyn.cerit-sc.cz/

[11] User authentication, authorization, and identity management for services in structural biology, talk at Digital Infrastructures for Research 2017, https://indico.egi.eu/indico/event/3455/session/12/contribution/115

[12] Perun IAM system: http://perun.cesnet.cz/

[13] AAI guide for West-Life users: http://internal-wiki.west-life.eu/w/index.php?title=Registering_new_account_and_linking_identities

[14] AAI guide for West-Life service operators: http://internal-wiki.west-life.eu/w/index.php?title=Enabling_SAML2_for_end_services

Background information

This deliverable relates to WP4; background information on this WP as originally indicated in the description of work (DOW) is included below.

WP4 **Title:** Operation and maintenance of the computing and data infrastructure
Lead: MU
Participants: STFC, EMBL, MU, CSIC, CIRMMP, UU, Luna, INFN

Work package number	4		**Start date or starting event:**	1
Work package title	Operation and maintenance of the computing and data infrastructure			
Activity Type	OTHER			
Participant number				
Person-months per participant:				

Objectives
The principal objectives of this work package are:
- **O4.1:** Setup the project testbed, define interfaces used to provision hardware resources, and negotiate provisioning with the resource providers at the technical level
- **O4.2:** Define, implement, and deploy consolidated architecture for job submission and data access.
- **O4.3:** Review existing security frameworks and define consolidated solution
- **O4.4:** Ensure smooth migration of the legacy portals to the consolidated architecture.

Description of work and role of participants
Task 4.1: Consolidation and operation of the infrastructure
Leader: INFN
Participants: MU, STFC, EMBL-HH, UU, CIRMMP, LUNA
The task is responsible to continue and improve the operations of previously developed computational platforms and services, including those developed within the previous WeNMR project as well as the ones already offered by the different partners addressing issues in X-ray crystallography and cryo-EM. We will achieve a better integration among them and simplify user interaction over a large set of experimental techniques in Structural Biology.
The West-Life project will leverage for resource provisioning both the EGI High-Throughput Data Analysis platform (the grid) and the EGI Federated Cloud, and possibly other public commercial or private cloud providers. The inventory of the computing and storage resources supporting the project and the selection of services available from the existing production e-infrastructures and/or EU projects to access, manage and operate these resources in a coherent manner will be done at the beginning of the project (M4.1), including a testbed available for rapid prototyping. The testbed will become the basis for the West-Life production platform. Because the introduction of new features and migration to new interfaces must not jeopardise the scientific work of the current user communities, we expect this platform to evolve in several generations that will co-exist in time in the phases of

unstable integration, stable production, and eventual phase-out. Times of switching an integrated next generation of the platform to production are the principal milestones of the work package (M4.2, M4.4, and M4.5).

All the partners are either directly involved or they collaborate closely with their National Grid Infrastructures (NGIs), hence having direct connection to EGI. Those synergies will be leveraged during operation of the project infrastructure rather than building it independently. In particular, established operational mechanisms for infrastructure monitoring, reliability evaluation, security alert and incident management etc. will be reused.

Task 4.2: Consolidation of job management mechanisms

Leader: LUNA

Participants: MU, INFN, UU, CIRMMP

An important point of the entire architecture is the interaction with the underlying cloud e-Infrastructure. Due to the legacy of current solutions, many different dispatchers are used submitting jobs to various infrastructures. This approach brings non-negligible overhead of maintaining these components, and it is not sustainable in the long term. On the contrary, uniform dispatcher technologies compatible with cloud interfaces have emerged over the past few years (DIRAC, HTCondor, Mesos, and Docker -container technology-). In addition, standardized interfaces to cloud resources running various cloud-management middleware (OpenStack, OpenNebula, ...) and even commercial clouds (Amazon EC2, Microsoft Azure) will be considered. In collaboration with T4.1, gradual migration from legacy interfaces to the standardized ones will be negotiated with resource providers so that legacy interfaces can be phased out smoothly, without affecting the user community.

Specific scientific code is expected to be wrapped in virtual machines in a standardized way so that further applications can be added easily. This task will develop guidelines for such wrapping in order to integrate with the dispatcher smoothly (preparation of specific VM images is done elsewhere, WP5 and WP7 in particular). For this, the project will select a consistent solution in its initial phase after a detailed analysis, and a consolidated architecture will be proposed in D4.1.

The interfaces among the components of the entire architecture are critical for its stable operation. While the work of WP5 is development of the Web Portal/VRE, it’s the responsibility of WP4 to specify the interfaces/API precisely, so that the VRE can connect to the dispatcher in a seamless and reliable way. The interface specifications are included in the architecture description deliverables D4.1, and they will be revised according to deployment experience in D4.5 and D4.6.

Further, the present task will keep supporting deployment of the portal frontend and applications using the consolidated job management mechanism. Based on experience with the deployment and according to incoming detailed requirements, the job submission architecture will be updated and revised eventually (thus contributing to deliverables D4.3, D4.5, and D4.6).

Task 4.3: Programmatic Access to datasets

Leader: STFC

Participants: LUNA, INFN

Many services rely on existing external data sets. The infrastructure will have to be capable of leveraging this external data upon users’ request. The task will review the relevant datasets to be made available, and it will define architecture and appropriate interfaces to access them, including eventual strategies to make “caching” copies of the data. It will be built on the metadata services to be offered in WP6.

Together with T4.4 security issues (authorization and user identity delegation in particular) will be addressed.

Unlike the other tasks of WP4, T4.3 brings functionality that is not widely present in current portal solutions. Therefore, it starts later (year 2) in order to build on experience and intermediate outcomes of WP6.

Task 4.4: Unified security and accounting model

Leader: MU

Participants: INFN, LUNA
e-Infrastructures require the users of the precious resources (computing power and storage capacity) to be reliably authenticated. Typically, rather heavyweight mechanisms like X.509 certificates are used. On the other hand, when the user interacts with the infrastructure through an application portal, lightweight mechanisms (username + password) are strongly preferred. Identity federations, which allow the user to authenticate with his/her home institute credentials, have gained popularity recently.
Due to various technical reasons, the primary user credentials are normally not passed or mapped directly to the computing and storage resources. Instead, portals often use “robot” certificates – the identity of the portal itself, which is accepted by the resources. Users are identified with a lightweight mechanism with the portal, and the portal must maintain mapping of the users credentials to specific use of the infrastructure with the robot credential in order to ensure accountability and traceability.
The task will review security mechanisms used by the various community portals, and it will design a unified mechanism for the project, balancing the impact on the existing users with clean design and maintainability of the resulting solution. The design will be specified in D4.2, and its gradual implementation will be part of the consolidated deployments – milestones M4.2, M4.4, and M4.5.

Deliverables

No.	Name	Due month
D4.1	Consolidated architecture of job submission and interaction with infrastructure	9
D4.2	Common security model design	9
D4.3	Report on experience with deployment of consolidated platform and its interaction with infrastructure	15
D4.4	Overview of external datasets, strategy of access methods, and implications on the portal architecture	18
D4.5	Report on progress of the deployment of consolidated platform and its interaction with infrastructure	26
D4.6	Final report on deployment of consolidated platform and the overall architecture	36