

Deliverable D4.6

Project Title:	World-wide E-infrastructure for structural biology	
Project Acronym:	West-Life	
Grant agreement no.:	675858	
Deliverable title:	Final report on deployment of consolidated platform and the overall architecture	
WP No.	4	
Lead Beneficiary:	4: Masarykova Univerzita	
WP Title	Operation and maintenance of the computing and data infrastructure	
Contractual delivery date:	Month 36	
Actual delivery date:	Month 36	
WP leader:	Name: Ales Krenek	Partner: MU
Contributing partners:	STFC, EMBL-HH, MU, CSIC, CIRMMP, UU, Luna, INFN	

Deliverable written by Marco Verlato, Aleš Křenek, Daniel Kouřil, Tomáš Kulhánek, Laura Del Cano, Mikael Trellet, with contributions by many others.

Contents

Contents	2
1 Executive summary.....	3
2 Project objectives.....	3
3 Detailed report on the deliverable	4
3.1 The HTC and Cloud production e-infrastructure	4
3.1.1 HTC e-infrastructure usage.....	4
3.1.2 Cloud e-infrastructure usage.....	8
3.1.3 KPIs summary	11
3.2 The consolidated platform on top of the e-infrastructure	12
3.2.1 Consolidation of job management mechanism.....	12
3.2.2 Deployment of applications in the Cloud e-infrastructure	14
3.2.3 Programmatic access to datasets.....	20
3.3 Unified security and accounting model	23
3.3.1 The West-Life security model	23
3.3.2 Data Protection and Privacy	25
3.3.3 West-Life Accounting.....	25
References cited	27
Background information	28

1 Executive summary

This document is the final report about the activities of the Work Package 4 (WP4), aiming at provisioning a consistent e-infrastructure gradually integrating the existing isolated software solutions in the structural biology field into 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 documents D4.3, MS14, D4.5 and MS15, respectively delivered at project month 15, 24, 26, 34, so that mostly the progress achieved until project month 36 not already described in the previous D4.5 ten months ago will be reported here, with references to MS15 when possible.

The document starts with an updated description of the resources potentially available for the project from the EGI e-infrastructure, on top of which we built the consolidated West-Life platform. It then presents a detailed view of resource usage and their geographical distribution in the third and last year of the project, as obtained from the EGI Accounting Portal. The remaining of the document reports in details the final achievements about the three main aspects of the platform: 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. It is currently supported with public European funds through EOSC-Hub, a three-year Horizon 2020 project started in January 2018. As of January 2018, 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. 850,000 CPU-cores are available for the High Throughput Computing (HTC) platform, while 21 cloud providers are available for the Cloud platform (known as the EGI Federated Cloud). Furthermore, 300 PB and 350 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 that 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 850,000 CPU-cores currently provided in total by the EGI grid, around 79,000 (100,000 at month 26) belong to the 17 (30 at month 26) resource centres that nowadays 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 17 resource centres are distributed in 9 (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 third year of the West-Life project, from September 2017 until the end of August 2018, being the HTC accounting data for September and October 2018 not yet available at the time of writing. During 2018 EGI.eu and the enmr.eu VO (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 until 31/12/2020 an amount of opportunistic computing time up to 53 Million of normalized CPU hours and opportunistic storage capacity up to 54 TB [1].

Five resource centres signed this last version of the SLA: INFN-PADOVA (Italy), TW-NCHC (Taiwan), SURFSara and NIKHEF (The Netherlands), NCG-INGRID-PT (Portugal).

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	Sep 2017 — Feb 2018	Mar 2018 — Aug 2018	Total	Percent
Netherlands	1,012,017	1,026,290	2,038,307	56.86%
United Kingdom	420,187	518,833	939,020	26.2%
Italy	133,090	136,866	269,956	7.53%
Portugal	47,976	55,599	103,575	2.89%
China	47,770	50,858	98,628	2.75%
Belgium	3,033	53,522	56,555	1.58%
Poland	9,597	45,119	54,716	1.53%
Taiwan	12,598	11,271	23,869	0.67%
Spain	65	2	67	0%
Total	1,686,333	1,898,360	3,584,693	
Percent	47.04%	52.96%		

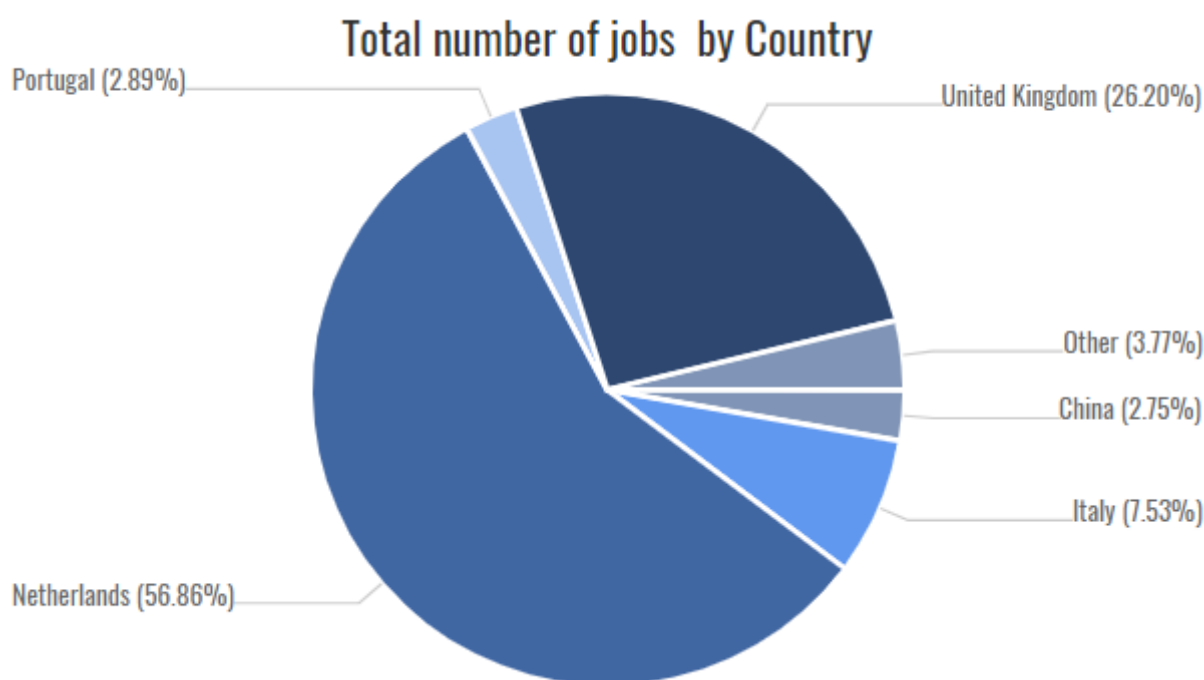


Figure 1: number of jobs by Country

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

Country	Sep 2017 — Feb 2018	Mar 2018 — Aug 2018	Total	Percent
Netherlands	4,265,108	9,325,515	13,590,624	65.08%
United Kingdom	1,859,519	1,932,729	3,792,249	18.16%
Italy	456,858	699,594	1,156,452	5.54%
Poland	191,820	842,311	1,034,131	4.95%
Portugal	178,779	221,791	400,570	1.92%
China	123,109	240,692	363,800	1.74%
Belgium	23,483	305,084	328,568	1.57%
Taiwan	93,882	123,162	217,044	1.04%
Spain	3	0	3	0%
Total	7,192,562	13,690,879	20,883,440	
Percent	34.44%	65.56%		




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 last report available, with values computed for 9 months until June 2018. For those months when a site is below the targets (values with orange background) an explanation is given. It turned out that in most cases the overall site lower performance relies on failures of the Storage Resource Manager (SRM), a component actually not used by West-Life jobs. Therefore, this underperformance did not affect the workload of the West-Life grid applications.

INFN-PADOVA		Previous period				Reporting period				
	Service target	2017-10	2017-11	2017-12	2018-01	2018-02	2018-03	2018-04	2018-05	2018-06
Availability	85%	100,00%	95,30%	84,29%	89,07%	100,00%	98,06%	99,99%	92,74%	100,00%
Reliability	90%	100,00%	95,30%	99,56%	94,44%	100,00%	98,06%	99,99%	92,74%	100,00%

NCG-INGRID-PT		Previous period				Reporting period				
	Service target	2017-10	2017-11	2017-12	2018-01	2018-02	2018-03	2018-04	2018-05	2018-06
Availability	85%	97,17%	98,77%	99,84%	90,74%	76,48%	100,00%	95,06%	99,56%	89,24%
Reliability	90%	97,17%	98,77%	99,84%	90,74%	76,48%	100,00%	95,06%	99,56%	89,24%

Explanation	2018-02	SRM authentication failures
	2018-06	CE host certificate expired, SRM timeout failures

NIKHEF-ELPROD		Previous period				Reporting period				
	Service target	2017-10	2017-11	2017-12	2018-01	2018-02	2018-03	2018-04	2018-05	2018-06
Availability	85%	99,37%	98,41%	93,61%	96,55%	77,33%	100,00%	94,08%	97,61%	93,90%
Reliability	90%	99,37%	98,41%	93,61%	96,55%	77,58%	100,00%	94,08%	97,61%	93,90%

Explanation	2018-02	SRM timeout failures
-------------	---------	----------------------

SARA-MATRIX		Previous period				Reporting period				
	Service target	2017-10	2017-11	2017-12	2018-01	2018-02	2018-03	2018-04	2018-05	2018-06
Availability	85%	94,18%	97,49%	37,53%	42,93%	98,62%	98,69%	99,65%	99,57%	100,00%
Reliability	90%	98,34%	97,49%	37,53%	43,52%	98,62%	99,91%	99,65%	99,57%	100,00%

Explanation	2018-01	SRM not published in the BDII
-------------	---------	-------------------------------

TW-NCHC		Previous period				Reporting period				
	Service target	2017-10	2017-11	2017-12	2018-01	2018-02	2018-03	2018-04	2018-05	2018-06
Availability	85%	97,15%	33,83%	0,00%	0,00%	65,09%	99,73%	99,66%	99,55%	99,98%
Reliability	90%	97,15%	33,83%	0,00%	0,00%	65,09%	99,73%	99,66%	99,55%	99,98%

Explanation	2017-11 / 2018-02	SRM not published in the BDII (already investigated in the ticket mentioned below, no further actions are necessary)								
		SLA Violation: A/R under the target for 4 consecutive months, from 2017-11 to 2018-02						Legend Underperforming On Target		
		GGUS ticket: https://ggus.eu/index.php?mode=ticket_info&ticket_id=133282								

Figure 3: HTC site availability/reliability over time

According to the EGI accounting system more than 3.5 Million jobs (3.4 Million in the second year) have run, translating into 20.9 Million (18.3 Million in the second 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 and D4.5, despite the

reduction in the number of opportunistic sites supporting the project. Most of the jobs (~70%) in fact still run on the resource centres who renewed the SLA, with a foremost contribution of the ones operated in Netherlands. 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 (DIRAC4EGI). This explains why a fraction of the jobs still are executed e.g. by resource centres in UK (which did not renew the SLA), Belgium, Poland and China. Also notice that, due to the pilot job mechanism implemented by DIRAC system, the number of jobs measured by the EGI accounting system does not correspond to the number of individual grid jobs submitted by the users through the application portals, being the latter significant higher (~7.5 Million/year, as accounted by the portals themselves and reported in D5.8).

As pointed out already in D4.5, two resource centres supporting the enmr.eu VO, CIRMMP in Italy and Queen Mary University of London in UK, provide a total of 6 NVIDIA Tesla K20m, 4 NVIDIA Tesla K80, and 1 NVIDIA Tesla K40c GPGPU cards. These nodes were initially included in the HTC e-infrastructure through a prototype for Scientific Linux 6 operating system developed by INFN to enable GPGPU support in CREAM-CE, and allowed AMPS-NMR, DisVis, PowerFit applications greatly enhancing their performance. In coordination with the EGI Software Provisioning team INFN released a production version of CREAM-CE with this GPGPU support for CentOS7 operating system in UMD (Unified Middleware Distribution) 4.6.0 in December 2017, and further update in UMD 4.7.1 in July 2018. The production version of CREAM-CE was installed at CIRMMP during summer.

3.1.2 Cloud e-infrastructure usage

The EGI Federated Cloud as of January 2018 hosts 21 resource centres. Four of them support the enmr.eu VO: these are CESNET-MetaCloud in Czech Republic, INFN-PADOVA-STACK in Italy, IFCA-LCG2 in Spain and IISAS-GPUCloud in Slovakia. The first two also signed the SLA with the VO, so that the EGI operations team keeps monitored their performance as done for the HTC resources. The figure 4 below shows their last availability/reliability report.

Legend

Underperforming

On Target

CESNET-MetaCloud		Previous period			Reporting period					
	Service target	2017-10	2017-11	2017-12	2018-01	2018-02	2018-03	2018-04	2018-05	2018-06
Availability	85%	99,85%	100,00%	100,00%	98,20%	97,64%	98,20%	99,75%	97,52%	100,00%
Reliability	90%	99,85%	100,00%	100,00%	99,55%	97,64%	98,20%	99,75%	98,07%	100,00%

INFN-PADOVA-STACK		Previous period			Reporting period					
	Service target				2018-01	2018-02	2018-03	2018-04	2018-05	2018-06
Availability	85%				91,82%	100,00%	98,12%	99,75%	99,18%	99,04%
Reliability	90%				99,90%	100,00%	98,12%	99,75%	99,18%	99,04%

Figure 4: Cloud site availability/reliability over time

The figures 5 and 6 show instead the total number of Virtual Machines (VM) and their elapsed time (in hours) provided by each site in the third year of the West-Life project, from September 2017 until the end of August 2018. West-Life applications (mainly by Scipion and Gromacs) executed 1369 VMs (855 in the second year) and 64,523 CPU wall time hours (107,587 in the second year).

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

Resource Centre	Sep 2017 — Feb 2018	Mar 2018 — Aug 2018	Total	Percent
CESNET-MetaCloud	385	408	793	57.93%
IISAS-GPUCLOUD	259	285	544	39.74%
INFN-PADOVA-STACK	13	6	19	1.39%
IFCA-LCG2	0	13	13	0.95%
Total	657	712	1,369	
Percent	47.99%	52.01%		

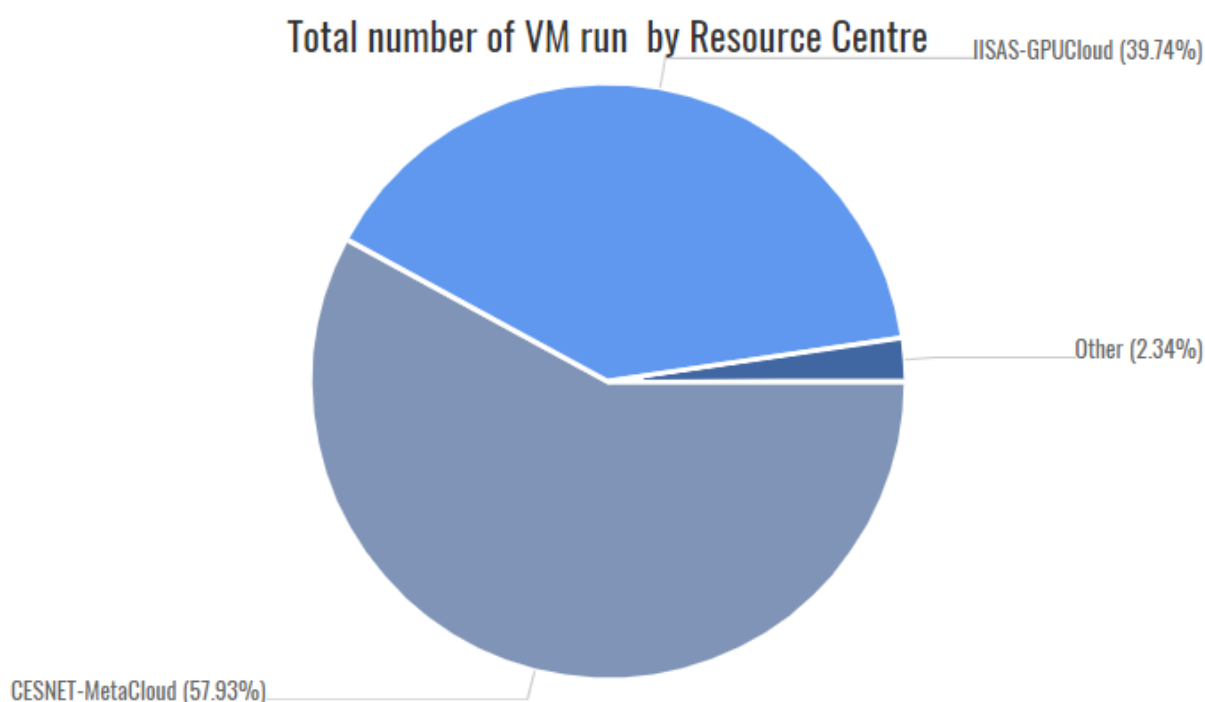


Figure 5: number of VMs site distribution

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. During 2018 Masaryk University has installed a GPU-enabled cluster (6138 Intel CPUs NVidia 1080Ti GPU cards) open (though not entirely exclusively) to West-Life services. This cloud instance was used for the developments of GromacsCloud application described in Sect. 3.2.2, but it is not part of EGI Federated Cloud for temporary technical reasons (transition from OpenNebula to OpenStack). Therefore, the pictures obtained through the EGI Accounting Portal do not show its usage statistics.

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

Resource Centre	Sep 2017 — Feb 2018	Mar 2018 — Aug 2018	Total	Percent
CESNET-MetaCloud	35,471	10,798	46,269	71.7%
IISAS-GPUCloud	12,816	2,325	15,141	23.46%
INFN-PADOVA-STACK	2,826	2	2,828	4.38%
IFCA-LCG2	0	295	295	0.46%
Total	51,112	13,420	64,532	
Percent	79.20%	20.80%		

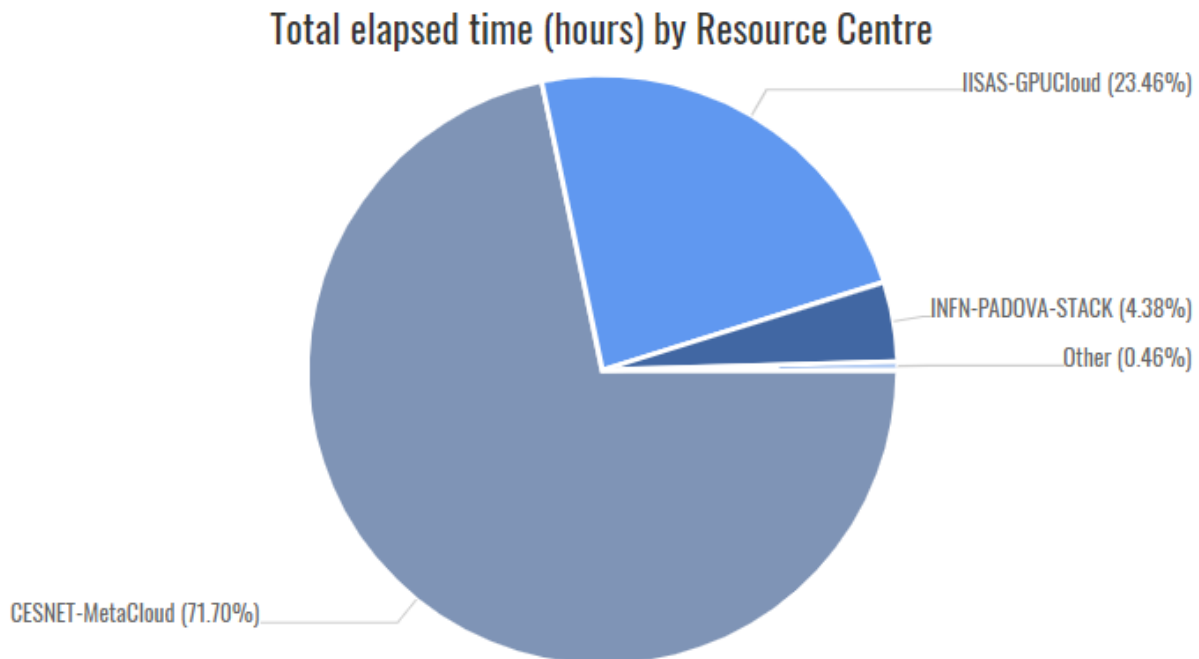


Figure 6: cloud elapsed time site distribution

3.1.3 KPIs summary

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

Metric	Project year 2	Project year 3	Difference (%)
CPU-cores (best effort)	100,000	79,000	-21%
CPU-cores (SLA)	6,850	6,050	-12%
Storage (SLA)	265 TB	54 TB	-80%
No. of HTC jobs	3,485,719	3,584,693	+3%
HTC Norm. CPU hours	18,343,084	20,883,440	+14%
No. of Cloud VMs	855	1,369	+60%
Cloud elapsed hours	107,587	64,532	-40%

Concerning HTC metrics, we can notice a slightly increasing usage despite the strong reduction in the number of CPU-cores provided by the opportunistic sites supporting the project. Around 70% of the jobs nowadays run in fact in the resource centres that renewed 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. However, 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 Section 3.2.2.

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 in the previous WP4 deliverable and milestone documents. 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 previous solution, i.e. the gLite/EMI WMS has been decommissioned at the end of 2017. While HADDOCK portal already migrated to DIRAC before that date, all other West-Life grid-enabled application portals hosted at UU (CS-Rosetta, SPOTON) and CIRMMP (AMPS-NMR) data centers still based on gLite/EMI WMS have migrated during 2018.

3.2.1.2 Application Software Deployment

No relevant news with respect to last D4.5 concerned the software deployment mechanism. Two West-Life Virtual Appliances already published in EGI Application Database were updated:

- ScipionCloud was updated to version 1.2 in June 2018. This appliance is intended to be run on a single node and is packaged with the remote desktop noVNC installed and configured to access the Virtual Machine through a Web Browser. It comes with two different images: ScipionCloud and ScipionCloud-GPU. The second one is ready to use GPGPU's for EM processing and 3D remote visualization.
- The West-Life VM Appliance, containing the Virtual Folder consolidating multiple storage providers and West-Life related software (CCP4, Scipion, etc.), was updated to version 18.04 in April 2018. It comes with two different images: one supporting VirtualBox hypervisor and suited for OpenNebula based clouds, and the other one supporting KVM hypervisor and suited for OpenStack based clouds. The virtual machine starting from these images takes the operating system (Scientific Linux 7.2) from cernvm.ch and software from west-life.egi.eu sites.

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 also updated when necessary.

The INDIGO-DataCloud project ended in September 2017. However, most of its software products has been taken over by the EOSC-Hub project started 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.

In particular, DisVis and PowerFit applications were integrated into a CI/CD pipeline, as already reported in deliverable D4.5. The overall architecture is shown in figure 7, taken from a paper published by INDIGO and UU developers in September 2018 [2], where a more detailed description can be found.



Cloudify is the cloud orchestration framework we evaluated in D4.1 to be the most promising out of available software, and which was used to deploy portal in clouds in the following activities. In PY3, we upgraded the portal deployment recipes (Gromacs and Scipion) to the new major version Cloudify 4. It significantly simplifies usage at the client side (i.e. work of the portal administrator) due to streamlined installation. Even more important is the improved Cloudify Manager service; according to the experience of our usecase, only version 4 became mature enough to work with our applications correctly. We provide the service running at <https://cfm.westlife.dyn.cerit-sc.cz/>. Access to this service is limited to the project partners and

their close collaborators only in order to provide portal deployments; it is not intended to the end users. Both Gromacs and Scipion (see below) are deployed using this service.

3.2.2.2 Gromacs

In PY3 the work on the cloud deployment of Gromacs continued. The Cloudify recipes were update to version 4 (see above) which allowed full-featured deployment with Cloudify Manager service.

In particular, use of Cloudify Manager allows for dynamic scaling of the environment. The deployment consists of a front-end (“all in one”) node, which provides also one-purpose Torque server, and one Torque worker node by default. Once users start submitting jobs and the Torque queue grows, it is detected by a specific sensor and communicated to the Cloudify Manager, triggering the “scale up” action, which creates a new worker node(s) in the cloud, and adds them to the Torque pool. Symmetrically, the nodes are removed when idle.

The user interface of the Gromacs portal was also migrated to the visual style of West-Life. Testing of the whole system, internal and involving a group of pilot users, revealed several bugs, and it brought also a few suggestions for the UI improvement, which were implemented. The service is fully integrated with West-Life AAI.

We provide a testing instance at <https://gromacs.westlife.dyn.cerit-sc.cz/>, and user groups are encouraged to create their own deployments using resources they have access to.

3.2.2.3 Scipion Web Tools

Scipion Web Tools portal has been smoothly running on the EGI Federated Cloud, at the CESNET-MetaCloud site, where it can be reached at <http://scipionwebtools.westlife.fedcloud.eu/m/services/>. The major enhancement since deliverable D4.5 has been to integrate access to the West-Life Virtual Folder as a input source for the web tools, as shown on Figure 8.

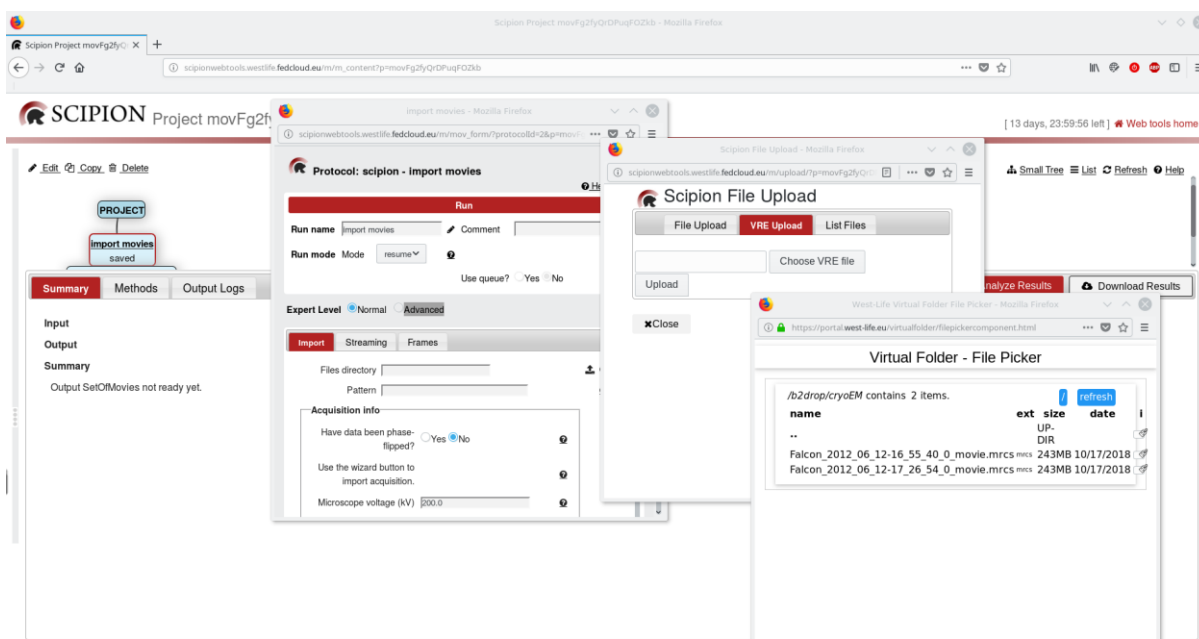


Figure 8 Selecting Virtual Folder files as Scipion Web Tools inputs

3.2.2.4 Scipion Remote Desktop

Scipion is an image processing framework, originally designed as a desktop PC application, for obtaining 3D models of macromolecular complexes using Electron Microscopy. It integrates several software packages allowing scientists to execute workflows combining different software tools. Hardware requirements for real-world EM data analysis scenario may reach beyond usual desktop PC. Many software dependencies also jeopardize its portability, especially in the world of fragmented desktop operating systems. Last but not least, CryoEM datasets are large typically, which complicates their efficient sharing on desktop PCs. We address all these issues by providing an integrated environment which deploys the application in the cloud and exposes a web-only interface to the user.

Scipion deployment, including all the dependencies and VM configuration, is described in a TOSCA blueprint and a set of Puppet recipes. Those are processed with Cloudify to set up the cloud node. In this way we avoid dealing with large pre-canned VM images which are difficult to maintain and to copy over network.

The VM setup with Cloudify and its access with VNC client is still too complicated for the typical user, therefore we wrap it with a thin, application specific software layer, which manages the deployments and their lifecycle, handles errors etc.

The user is exposed to a one-stop web interface (developed using React/Redux framework) where he/she manages the deployed machines and gets instant access to the remote VM desktop in the web browser.

The overall architecture is shown in Figure 9.

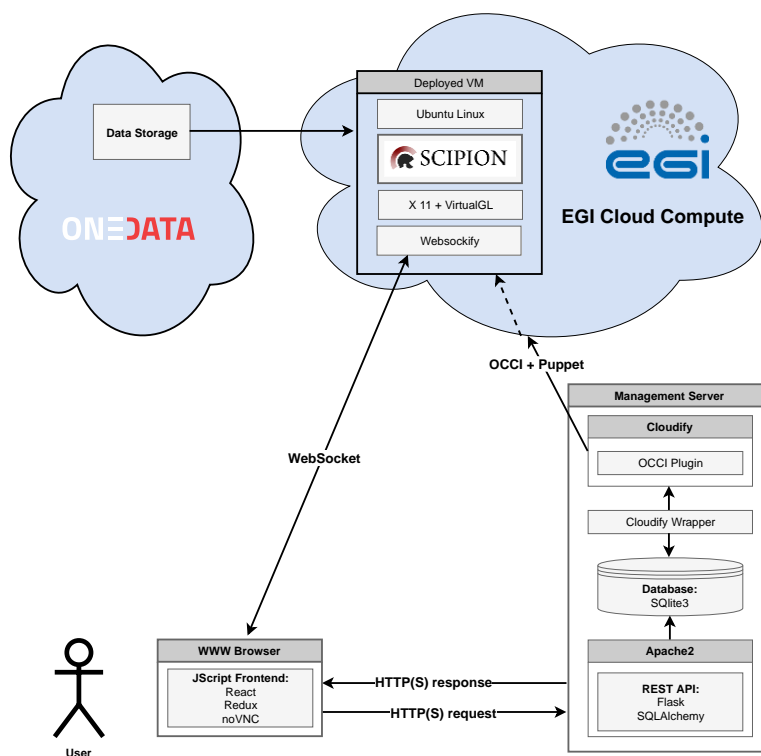


Figure 9 Schema of Scipion deployment in the cloud

Display of reconstructed 3D electron density map requires hardware accelerated rendering. A GPU has to be attached to the virtual machine in the cloud, typically via PCI-passthrough. The VM runs a headless X11 server with hardware accelerated OpenGL 3D rendering using the GPU. Scipion is run in VNC server. OpenGL 3D rendering calls are intercepted by a preloaded VirtualGL library, they are redirected to the accelerated X11 server, rendered to off-screen windows, and the resulting pixmaps are copied back. The same GPU is used for accelerated computation too. Snapshot of the GUI is shown in Fig. 10.

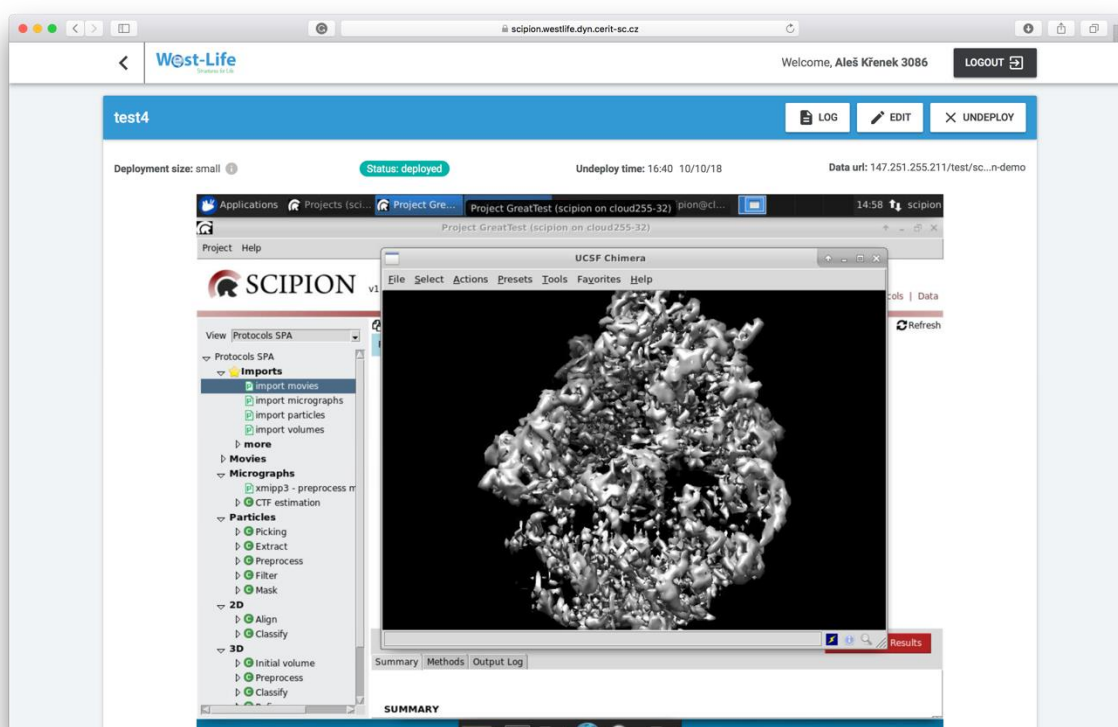


Figure 10 Scipion remote desktop in web browser

Scipion project is a set of files in a single folder with typical size of 10–10,000 GB. We integrate the cloud setup with OneData storage, shielding the user from the need to copy these datasets around manually.

OneData provides a web interface for basic data manipulation. In the VM, a FUSE client for Linux is used to mount the same workspace. Due to performance and stability reasons, the work at the VM is done on a local data copy, which is periodically synchronized back to OneData.

The management server is interfaced with the West-Life authentication infrastructure (AAI), following recommendations of AARC project. The authentication information translates internally to JSON Web token mechanism, completely transparently to the user.

Access to the OneData web interface uses West-Life AAI too. However, in order to keep control on credential delegation, the user is required to generate a specific access token specific to the Scipion project, and to paste it into a dedicated field in the deployment form. The token is passed to the VM and used there to act on behalf of the user.

The work was presented at DI4R 2018 conference as a poster [3], code is available on Github [4].

3.2.2.5 AMPS-NMR

The portal was fully integrated with West-Life SSO as well as Virtual folder in this period.

3.2.2.6 A HADDOCK portal in the cloud (HelixNebula Pilot)

As part of a collaboration with the HelixNebula Pilot, a public-private partnership to provide open cloud services, UU have put some efforts to virtualize the new web portal version of HADDOCK 2.4 (beta version). For this purpose, RHEA / Exoscale, one of the cloud provider of HelixNebula made available a cluster of 26 nodes (1 master node / 400GB / 2*2.2GHz + 25 compute nodes / 50GB / 2.2GHz) as test bed. The objective was to be able to automate the deployment of the new HADDOCK 2.4 virtualized pipeline (web portal / user DB / software) on new resources. The major achievement of this use case was the containerization of the 3 front-end modules within 3 docker images orchestrated by docker-compose and potentially scalable using docker-swarm. Schematic illustration of the modules is presented in Fig. 11.




Figure 11. HADDOCK 2.4 docker containers for the 3 main front-end

The main outcome of the containerisation is the automation of the front-end deployment that now allows deployment of the HADDOCK web portal on any cloud resources within a very short time frame. The main configuration steps are limited to only providing the IP that will be used to access the portal and creating a self-signed certificate (if no signed certificate is associated to the IP that will be used) to encrypt the data communication.

On the back-end side, some efforts still need to be made to fully automate the configuration and installation of the HADDOCK software and its dependencies. This involves compiling and installing the HADDOCK suite and setting up the cron jobs that will monitor HADDOCK workflow + allow for running HADDOCK within docker containers on the compute nodes. This last step is our mid-term plan to obtain a fully-automated pipeline of HADDOCK 2.4 front-end/back-end that could be deployed on any cloud resources available. As a proof of concept,

a working version of HADDOCK 2.4 (from the web portal to the job execution/processing) has been made available and tested successfully. A demo has been given during the last HelixNebula workshop in Amsterdam.

3.2.3 Programmatic access to datasets

The activity of the task 4.3 of the Work Package 4 aimed 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 MS13: *Prototype access to selected dataset* published respectively at month 18 and month 22.

The main conclusion is that the West-Life Virtual Folder web UI integrates datasets from PDB database and PDB-REDO database seamlessly into the appropriate context within a "special folder" allowing browsing of database entries, visualizing and downloading them. Litemol – a web component for visualizing protein structure - is integrated into Virtual Folder's file manager. Additionally an autocomplete component within metadata section of the file manager integrates other components visualizing data from PDB database. Third party applications can navigate to specific part of Virtual Folder as context specific URL are available. VF provides access to datasets via WEBDAV interface or via REST API. Any third party application can be integrated into VF based on its WEBDAV capabilities or via REST API.

The screenshot displays the West-Life File Manager interface. The top navigation bar includes links for Virtual Folder, Home, Services, Support, News, About, Cloud, Developers, Contact, and Introduction. The main content area is titled 'File Manager' and contains two panels. The left panel shows a file list for the directory '/experiment_pcloud/Westlife/notebooks/provenance-test', which contains 5 items. The right panel shows the 'Metadata' tab for the file '434159-1.pdb'. This tab includes sections for generic metadata (name, extension, path, size, date, local url, public url), free text metadata, and related entries from public databases (PDB, Uniprot, etc.). A dropdown menu for 'provenance' is open, showing a list of related PDB entries (1KJ0, 1KJ1) and a 'Generate' button.

Figure 12. Metadata tab available for each file or directory contains generic section, entries section with autocomplete component integrated with PDB and Uniprot Databases and provenance section

At infrastructure level, INFN consolidated its deployment of the Onedata by updating the software to version 18.02.0-rc8 in July 2018. Onedata is a distributed storage system based on a network of storage providers (Oneprovider service) federated together through a central manager (Onezone service). The Onezone service is responsible for the user authentication and authorization; it delivers a token to the authenticated user, which grants him the access to the spaces on any federated Oneprovider. The Oneprovider service hosted at INFN-PADOVA data center offers up to 14 TB of disk storage. West-Life users can ask for a storage space at INFN-PADOVA by connecting to the Onezone server (onezone.cloud.cnaf.infn.it) hosted at INFN-CNAF data center (located in Bologna).

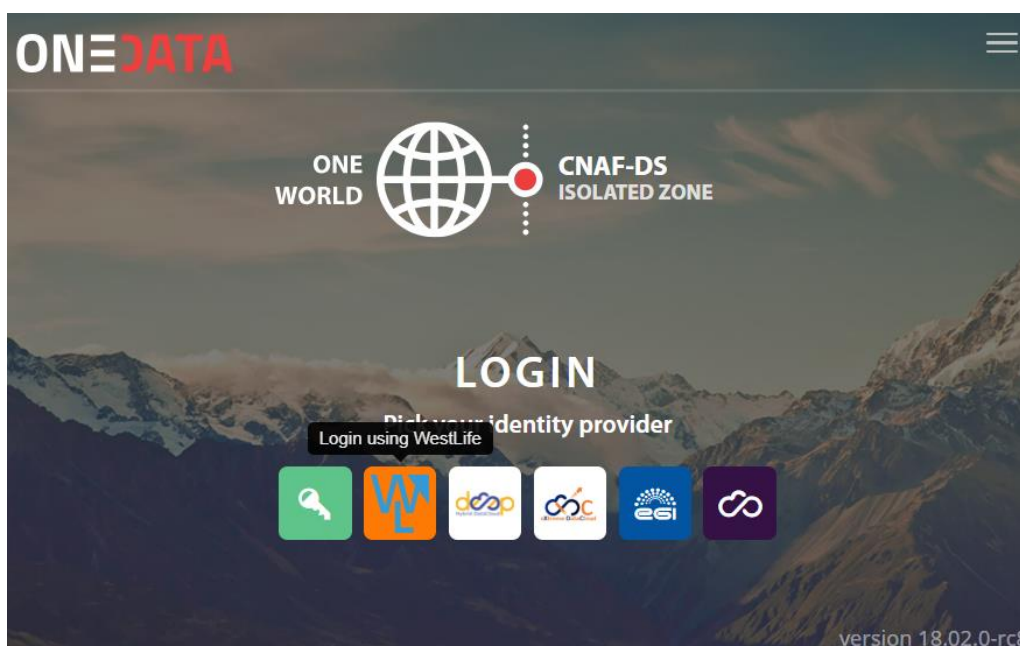


Figure 13: Onedata login dashboard

INFN implemented the support to West-Life SSO as authentication method for Onedata. It required the development of a new plugin specific for the West-Life SSO service that was contributed upstream to Onedata software. Figure 8 show the Onezone LOGIN web page, with a list of supported identity providers. By selecting the orange button, the users can login using the West-Life SSO. The documentation with the instructions for West-Life users on how to create, manage and use Onedata storage is available at the West-Life main web site.

Furthermore, INFN and STFC worked on the integration of West-Life Virtual Folder with Onedata, i.e. to enable Onedata storage system as additional back-end, other than Dropbox, B2Drop and WebDAV supported systems. This activity has been described in detail in the recent MS15 report, so will be only summarized and updated here. The integration between Virtual Folder and Onedata consists in a plugin for the Virtual Folder framework, which is able to get any information about files and directories from a given Oneprovider. The plugin is able to mount on demand any user space available on the provider. The operation requires the execution of the specific command line tool, Oneclient.

During the last two project months, the plugin has been deeply tested and some concerns have emerged:

- several compatibility problems between Oneclient and Oneprovider, forcing to install the same release of the service
- dependency on third party libraries which are not stable enough

- Onedata software maturity still at the level of release candidate

Therefore, for the time being we must consider the integration with Virtual Folder as a prototype. The integration will be "production ready" only when a final release of Onedata is available, and we will support only the stable release.

3.3 Unified security and accounting model

3.3.1 The West-Life security model

The VRE established facilitates access to a number of services. Together with flexibility, security aspects need to provide a sufficient level of security, which will protect users, service providers and operators of the infrastructures. The security model of West-Life handles establishment of a trust relationships between the service and user and provides a unified mechanisms handling authentication and controlling access of West-Life users to services. The West-Life AAI architecture reflects the AARC Blueprint Architecture, which provides a reference architecture to help design proper AAI solutions for researchers. The design is interoperable and based on current approaches, which ensures compatibility and ease sustainability. The architecture addresses identity management, authentication, attribute management, and integration with end services. The solution is based on open standards, which eases the integration of end services.

The West-Life AAI was implemented during the project, which helped unify various approaches utilized by existing West-Life services and provide a single solution. The implementation is based on state-of-the-art advances in terms of technologies employed. As the key corner stone the AAI solution bases on the concept of an SP/IdP Proxy, which provides the interface for both users and services and controls the how flow. The proxy has been integrated with a number of identity providers and a number of West-Life services were connected. The activities performed over the reporting period focused on provisioning of a stable environment and accessing the key features of West-Life AAI to the end users and service providers. Effort was dedicated to extend the utilization of AAI, the number of user registered with the identity management system has been growing since the service was put in production in summer 2017, with more than 60 users being registered with the system at the moment. The AAI system was used to provision access to the West-Life services during training events (e.g. Scipion course at MU), during which the registration procedure was verified to work for a number of attendees.

The West-Life AAI technology was extended with support for the OpenId Connect (OIDC) protocol, which extended current SAML-based interfaces. OIDC is becoming a popular technology to integrate new services and is widely supported. Introducing the new interface extended the flexibility of the service and eased its adoption by end services. At the moment, the system registers 12 end services using SAML and 8 using OIDC, which are exposed to users. In addition to that, a number of testing and development services are registered, too.

Development and preparation of West-Life AAI is pursued with close links to similar developments in the field of life sciences (esp. the Life-Science pilot in AARC, development in ELIXIR and BBMRI) with the vision of providing a compatible solution that will be potentially be possible to merge with a more general solution applicable for the Life-Science domain. First communication about potential utilization of the EGI CheckIn service was performed also. In addition, all Utrecht based portals were migrated to a GDPR-compliant central user management system. As part of the EOSC-Hub project this system was connected to the EGI-CheckIn service with the OpenId Connect (OIDC) protocol. This has been put into production last September in all Utrecht-based portals (Figure 14). Users can now register for access using various credentials provided by the CheckIn and also directly submit jobs to the portals using their SSO credentials. The registration portal can be found here: <https://wenmr.science.uu.nl/>

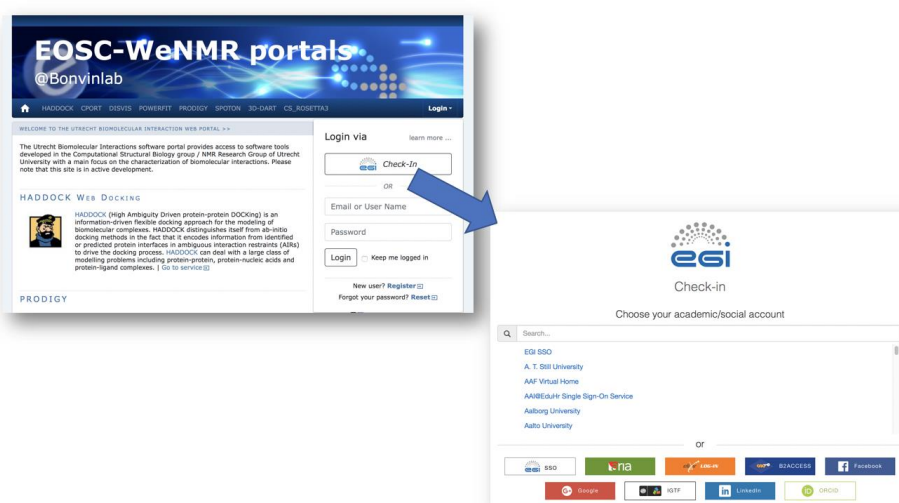


Figure 14: View of the CheckIn implementation in the Utrecht portal front-end

The key concept and current state of the West-Life AAI service was presented at the Digital Infrastructures for Research 2017 (DI4R) conference in Brussels (November 2017).

In close cooperation with MU and CESNET teams, the installation of proxy was moved under a joint operation that is responsible for a number of other proxy instances serving other projects and infrastructures (e.g. ELIXIR and national e-infrastructure).

3.3.2 Data Protection and Privacy

With the adoption of EU regulations on data protection and privacy, the West-Life AAI components were reviewed to ensure compliance with the existing requirements. The “West-Life AAI Acceptable Usage Policy” (AUP) and “West-Life AAI Privacy Policy” documents were updated and published from the West-Life portal (<https://auth.west-life.eu/>). Both the documents follow closely rules stipulated by other life-science infrastructures, trying to retain compatibility in the domain. The registration process was extended to clearly communicate the procedures to new users. Existing users were notified about information that is collected and processed for the need of West-Life AAI operations.

3.3.3 West-Life Accounting

Accounting complements security model by mechanisms to collect records on usage of West-Life services. Information about what resources users consumed is crucial to provide overviews about utility of services and planning further development. West-Life services usually maintain internal overview with usage records that are available to the service owners. There is no mechanism however to collect usage records on project level.

The West-Life accounting aggregation service was developed to provide a simple way of collecting usage records from services in the VRE. The main part of service is composed of an aggregator that is responsible for collecting data from individual services and storing the data in databases. A simple Rest-like protocol was designed to obtain usage records from West-Life services. In order to ease integration by the end services, a reporting module was developed that provides the endpoint offering usage records. The module is supposed to be integrated with the application and serve its usage records using the designed protocol.

The protocol is protected on transport level using the TLS protocol, where both the communications peers are supposed to use X.509 certificates for authentication. In addition, usage records are delivered as JSON messages signed using the service certificates. The trust between the accounting collector and individual services is based on the relationships established by the West-Life AAI environment. The accounting collector takes X.509

certificates used by the West-Life AAI proxy and expects that signed usage record sent from end services can be verified using them.

The concept gathering accounting data was demonstrated on Gromacs. All source codes are available from GitHub (<https://github.com/ICS-MU/westlife-accounting>) under an open-source license.

References cited

- [1] Updated VO SLA MoBrain/enmr.eu and VO OLAs MoBrain/enmr.eu, <https://documents.eji.eu/document/2751>

- [2] INDIGO-DataCloud: a Platform to Facilitate Seamless Access to E-Infrastructures, Salomoni, D., Campos, I., Gaido, L. et al. J Grid Computing (2018) 16: 381. <https://doi.org/10.1007/s10723-018-9453-3>

- [3] A. Křenek et al., Scipion in the Cloud, User Friendly CryoEM Data Analysis From Anywhere. Poster, DI4R 2018. <https://indico.eji.eu/indico/event/3973/contribution/80>

- [4] Scipion Cloudify deployment recipes, <https://github.com/ICS-MU/westlife-cloudify-scipion-web>

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: - 04.1: Setup the project testbed, define interfaces used to provision hardware resources, and negotiate provisioning with the resource providers at the technical level - 04.2: Define, implement, and deploy consolidated architecture for job submission and data access. - 04.3: Review existing security frameworks and define consolidated solution - 04.4: 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