



D4.5 Periodical assessment of Platform services

Lead partner:	EGI Foundation
Version:	1
Status:	Approved by the EC
Dissemination Level:	PUBLIC
Keywords:	EGI-ACE, Compute, Virtual Access, Platform services
Document Link:	https://documents.egi.eu/document/3794

Deliverable Abstract

The report provides assessment and statistics of all the Platform services provided under virtual access in WP4.



EGI-ACE receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101017567.

go.egi.eu/egi-ace

COPYRIGHT NOTICE



This work by parties of the EGI-ACE consortium is licensed under a Creative Commons Attribution 4.0 International License. (<http://creativecommons.org/licenses/by/4.0/>).

EGI-ACE receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101017567.

DELIVERY SLIP

	Name	Partner/Activity
From:	Enol Fernandez	EGI Foundation/WP4
Moderated by:	Sjomara Specht	EGI Foundation/WP1
Reviewed by:	Gergely Sipos	EGI Foundation/WP1
Approved by:	PMB, SDS, SFG	

DOCUMENT LOG

Issue	Date	Comment	Author
v.0.1	27/06/2023	Template	Hien Bui
v0.2	13/07/2023	First version ready for review	Enol Fernandez
v1	23/08/2023	Addressed Gergely Sipos comments	Enol Fernandez

TERMINOLOGY

<https://confluence.egi.eu/display/EGIG>

Terminology/Acronym	Definition
VA	Virtual Access
EOSC	European Open Science Cloud
PaaS	Platform as a Service

Contents

Executive summary	4
1 Introduction	6
1.1 Installations	6
1.2 Communities.....	8
1.3 Metrics definition.....	9
2 Installations	1
2.1 EGI Notebooks	1
2.1.1 Metrics.....	2
2.1.2 Assessment	3
2.2 EGI – DIRAC	4
2.2.1 Metrics.....	5
2.2.2 Assessment	5
2.3 CSIC DEEP training facility	6
2.3.1 Metrics.....	7
2.3.2 Assessment	7
2.4 LIP DEEP training facility	8
2.4.1 Metrics.....	9
2.4.2 Assessment	9
2.5 DODAS	10
2.5.1 Metrics.....	11
2.5.2 Assessment	11
3 Dissemination	12
4 Satisfaction.....	17
4.1 EGI Customer satisfaction reviews	17
4.2 EOOSC Marketplace orders	18

Executive summary

This report provides an assessment at M30, the end of the EGI-ACE project, of the WP4 installations provided by the project under the Virtual Access (VA) mechanism. These installations represent the Platform-as-a-Service (PaaS) layer in the project service catalogue. This assessment is based on the metrics collected by the five WP4 installations during the project duration in six rounds of 6 months each: M01-M05, M06-M10, M11-M15, M15-M20, M20-M25, M25-M30.

WP4 activity extended the EOSC Compute Platform with new installations previously not available in the EGI portfolio: CSIC DEEP training facility, LIP DEEP training facility (2 installations supporting the same service) and DODAS. The EGI Notebooks and Workload Manager services in the EGI portfolio were supported by the EGI Notebooks and EGI DIRAC installations respectively. A new service spawned from the Notebooks installation: EGI Replay, offering an environment for reproducing and sharing research. All the installations have gone through operations and maintenance tasks and installation of release upgrades following requirements from user communities. All installations are on-boarded in the EOSC Marketplace and have been progressively integrated with the EGI central services.

WP4 services have been used by 22 user communities. These 22 user communities represent 1,011 individual users in total. The user communities in EGI-ACE are mapped into Virtual Organisations (VOs) and normally named following a domain-like schema (e.g. vo.reliance-project.eu supports the Reliance Project user community). For DODAS, DIRAC and DEEP the users are 'expert users' from the scientific communities and interact with these technical platform services to develop/deploy and operate community specific platforms, Thematic services on top of them. Users of the Notebooks service are typically 'data scientists' who want to parse, process, visualise data in an interactive way. In particular:

- EGI Notebooks have been used by 273 users (22% growth over the previous 15 months) belonging to the vo.notebook.egi.eu and the vo.access.egi.eu communities, open to individual users, and to nine discipline/project specific communities: auger, biomed, CESSDA, C-SCALE, eiscat_3d, Disaster mitigation and agriculture, Lethe, PaNOSC, and vo.reliance-project.eu. EGI Notebooks outperformed the expected number of new users (150 at the end of the project) as it has become a popular entry point for getting access to the EGI infrastructure.
- EGI DIRAC has been used by 20 different communities (7% growth over baseline in previous 15 months), with 4 new communities incorporated during the project. DIRAC has supported 747 users belonging to these communities (10% over previous period). The target number of communities for DIRAC over the project period went slightly over the expectations thanks to the interest of communities incorporated during the project.
- The CSIC/LIP DEEP training facility has received usage from seven use cases from communities (a 700% growth over the previous period). The usage of the service was low until M15, and thanks to several dissemination activities it reached the target number of use cases for the project).

- DODAS usage has remained stable with one community: `fermi-lat.infn.it` with an increase of the number of users from 3 to 4. Although the number of users increased, DODAS was not able to meet the target of 100 deployments (reached 51 deployments) due to the lack of additional communities interested in the service.

During the M15-M30, the promotion of the installation has continued with dedicated presentations at different workshops and conferences, including the EGI Conference 2023 which featured a full day devoted to EGI-ACE with presentations, demos, tutorials and panel discussions.

Section 4 finally describes the level of satisfaction by checking the orders received via the EOSC portal and the EGI Customer satisfaction reviews, which showed an average level of 5 out of 5 during the reference period.

1 Introduction

Virtual Access (VA) is financial instruments to reimburse the access provisioning costs to access providers. This instrument is provided by the European Commission to increase the sharing of research infrastructures and services that otherwise would not be available to international user groups.

In VA, the services – also called “installations” – must be made available ‘free of charge at the point of use’ for European or International researchers. VA access is open and free access to services through communication networks to resources needed for research, without selecting the researchers to whom access is provided.

Virtual Access to services of the EGI-ACE catalogue applies to the following four categories:

1. Infrastructure Services WP3 - the Cloud Compute (IaaS) and High Throughput Compute services of the EGI portfolio supported by a set of 16 datacentres from the EGI Federation. The enabling components that support the Cloud Compute service: AppDB, for resource discovery and software catalogue; Dynamic DNS, for user-managed DNS provision of domain names for VMs and services running on the e-Infrastructure; and Infrastructure Manager (IM) for the basic orchestration of IaaS resources.
2. Platform Services WP4 - mature software tools offering generic capabilities to facilitate the usage of the underlying infrastructure for EOSC users and Data Spaces.
3. Federated data spaces WP5 - services provided by major European research collaborations, research infrastructures and research institutes, and are composed of mature software tools, datasets and services that offer science discipline specific processing and data analysis capabilities for EOSC users.
4. Federated Access Services WP6 – services providing secure access to other services and enabling large-data analysis workloads in the distributed infrastructure. Included services are delivered by major European research institutions using mature open-source software with already established user communities from multiple scientific disciplines.

This document provides Virtual Access metrics and assessment for WP4.

1.1 Installations

There are 5 service installations under VA in EGI-ACE WP4. The following of these 5 installation have been subject to change since the beginning of the project:

- DODAS target metrics were updated to consider the new operational model of the installation: In order to provide a higher level of flexibility, DODAS enabled the possibility to update the number of components and their configuration supporting a given use case from the upper layers of the service architecture (the one closest to users). This, in turn allows for the dynamic adjustment of the computing and storage resources managed by DODAS without re-deployment, when in the past any modification would imply a recreation of the cluster. With the current model, when a

user community decides to add features to an existing deployment, e.g. deploy a cache close to a on-demand batch, it does not need anymore to re-deploy everything from scratch but it can compose the system on the existing deployment. As such the expected number of deployments has been reduced from 900 to 100. These changes are now formalised in a project amendment.

- WP4 services have been integrated with the EGI-ACE Key Exploitable Result 2 (Services enabling federated computing in EOSC): Table 1 summarises the integration of the WP4 services with KER2. EGI Notebooks and EGI - DIRAC installations were part of the EGI portfolio as the Notebooks and EGI Workload Manager services respectively and as such were already integrated with most of the EGI core services. During this period, Notebooks completed the integration with accounting and piloted the integration with EOSC accounting system. Check-in integration for the Workload Manager depends on the transition from X.509 certificates to token-based authentication in the EGI High Throughput Computing service, which is a still ongoing activity and has prevented the completion of the activity. DODAS have progressively completed the technical integration during the course of the project. With the start of the AI4EOSC project, DEEP are being re-implemented with a more modern technology thus the integration effort stopped in EGI-ACE and has shifted to other development projects that will achieve complete integration of the newer AI4EOSC project toolset and EGI.

Table 1: WP4 integration matrix with EGI core services

Installation	Check-in	Helpdesk	Monitoring	GOCDB	Accounting
EGI Notebooks	pre EGI-ACE	pre EGI-ACE	pre EGI-ACE	pre EGI-ACE	DONE
EGI - DIRAC	Not completed	pre EGI-ACE	pre EGI-ACE	pre EGI-ACE	n/a
CSIC DEEP training facility	DONE	Not completed	Not Completed	DONE	n/a
LIP DEEP training facility	DONE	Not completed	Not Completed	DONE	n/a
DODAS	DONE	DONE	DONE	DONE	n/a

1.2 Communities

Table 2 summarises the usage of the WP4 installation by the EGI communities (both existing and new communities).

Table 2: Communities integration matrix with EGI-ACE WP4 installations.

Community type¹	Community	EGI Notebooks	EGI - DIRAC	CSIC DEEP training facility	LIP DEEP training facility	DODAS
WP2	EISCAT_3D	X	X			
	auger	X	X			
WP5	VIP	X	X			
	WeNMR		X			
	LOFAR		X			
	Disaster mitigation and agriculture					
	UseGalaxy.eu		X			
	OpenCoast		X			
LToS	vo.access.egi.eu	X				
	vo.notebooks.egi.eu	X	X			
	fedcloud.egi.eu		X			
	training.egi.eu		X			
Open Call	vo.reliance-project.eu	X				
	EMPHASIS			X		
	FERMI-LAT					X
	CESSDA	X				
	C-SCALE	X				
	Lethe Project	X				
	PaNOSC	X				
Existing	km3net.org		X			
	lsst		X			
	virgo		X			

¹ WP2 - Early Adopter in WP2; WP5 - Thematic Service in WP5; LTOS - Long tail of science community from EOSC Portal or similar channels; Open Call - Approached EGI-ACE via the Open Call or via related EOSC projects initiatives. Existing - Communities that existed in the service and are generating new usage as part of EGI-ACE.

	vo.grand-est.fr		X			
	vo.hess-experiment.eu		X			
	vo.sbg.in2p3.fr		X			
	vo.france-grilles.fr		X			
	beapps		X			
	vo.complex-systems.eu		X			

1.3 Metrics definition

For each installation several metrics has been defined between the provider and WP4 leader, taking into account following categories:

- **Number of users** – depending on the nature of installation, number could be defined based on accounts (if registration was required) or number of unique IPs (if registration is not needed to benefit of the service)
- **Usage** – the goal of this metric is to report how much the service is used. This metric depended on functionality provided by the service.
- **Number and names of the countries reached** – the goal of this metric was to report how broadly the service is used and how the geographical coverage is changing with time.
- **Marketplace orders** – the goal of this metrics is to provide information about how often the service is being ordered via EOSC Marketplace
 - This metric is not applicable to federation services due to the nature of the service. Federation services are enabling federation and are supporting delivery of customer facing services. Thus, cannot be ordered.

2 Installations

2.1 EGI Notebooks

Description	Notebooks is a browser-based tool for interactive analysis of data using EGI storage and compute services. Notebooks is based on the JupyterHub technology.
Task	4.1
URL	
Service Category	Platform Service
Service Catalogue	https://www.egi.eu/services/notebooks/
Location	Czech Republic
Duration	M1-M30
Modality of access	All the elements of the service are free at the point of use. Valid EGI user registered in Check-in is needed. There are 2 modes: notebooks for researchers available for users upon registration, and notebooks for communities with access controlled via the marketplace.
Support offered	Technical support is provided via the helpdesk central support team, and by the individual service providers. EGI Outreach activities also include webinars, training, and hands-on sessions during conferences and events.
Operational since	Nov 2019
User definition	single researchers and communities

2.1.1 Metrics

Metric name	Baseline	Define how measurement is done	Period 1 M1-M5	Period 2 M6-M10	Period 3 M11-M15	Period 4 M16-M20	Period 5 M21-M25	Period 6 M26-M30
No of users accessing	90	Internal service database	73	86	117	95	95	147
No of notebook sessions/month	100	Internal logs	66	104	106	104	154	236
No of countries reach	19	Check-in	21	31	29	32	25	55
Names of countries reach	UK, FR, ES, BE, NO, HR, NL, HU, SE, DE, PL, ZA, TW, AT, CH, IT, GR, TR, US	Check-in	Algeria, Croatia, Czechia, Denmark, Estonia, Finland, France, Germany, Greece, Indonesia, Italy, Netherlands, North Macedonia, Poland, Singapore, Slovakia, Spain, Sweden, Switzerland, United Kingdom, United States	Algeria, Austria, Belgium, Brazil, Canada, China, Croatia, Cyprus, Czechia, Finland, France, Germany, Greece, Hungary, Indonesia, Ireland, Italy, Japan, Netherlands, Norway, Poland, Portugal, Romania, Spain, Sweden, Switzerland, Turkey, United Arab Emirates, United Kingdom, United States, Vietnam	Austria, Belgium, China, Croatia, Czechia, Denmark, DR Congo, Egypt, France, Germany, Greece, Hungary, India, Indonesia, Italy, Netherlands, Norway, Poland, Portugal, Russia, Singapore, Sri Lanka, Sweden, Tunisia, Turkey, United Kingdom, United States, Vietnam	Algeria, Austria, Bangladesh, Belgium, Canada, Chile, China, Czechia, Finland, France, Germany, Greece, Hungary, India, Indonesia, Ireland, Italy, Latvia, Mexico, Morocco, Netherlands, Norway, Poland, Russia, Singapore, Somalia, Spain, Sweden, Switzerland, Turkey, United Kingdom, United States	Austria, Bangladesh, Belgium, China, Czechia, Denmark, Ecuador, Finland, France, Germany, Greece, Hungary, India, Italy, Lithuania, Netherlands, Norway, Poland, Spain, Sweden, Switzerland, Turkey, United Kingdom, United States	Algeria, Australia, Austria, Bangladesh, Belgium, Brazil, Bulgaria, Canada, Chile, China, Croatia, Cyprus, Czechia, Denmark, DR, Congo, Ecuador, Egypt, Estonia, Finland, France, Germany, Greece, Hungary, India, Indonesia, Ireland, Italy, Japan, Latvia, Lithuania, Mexico, Morocco, Netherlands, North Macedonia, Norway, Poland, Portugal, Romania, Russia, Serbia, Singapore, Slovakia, Somalia, Spain, Sri Lanka, Sweden, Switzerland, Taiwan, Tunisia, Turkey, Ukraine, United, Arab

								Emirates, United, Kingdom, United, States, Vietnam
--	--	--	--	--	--	--	--	--

2.1.2 Assessment

The EGI Notebooks service, hosted at CESNET (CZ), slightly decreased the number of users accessing during M16-M25, keeping the number of session and countries increased in all periods. Number of users increased again during the last period M26-M30. Overall there is a 22% growth in the number of users since the previous reporting period and has seen an increase of usage from 104 monthly sessions to more than 230 monthly sessions. This installation serves mostly individual users that belong to the vo.access.egi.eu and vo.notebooks.egi.eu Virtual Organization (VO)s. Additional VOs have been enabled to access the service: biomed, auger, eiscat_3d and vo.reliance-project.eu, with five additional ones configured during the M16-M30 period: CESSDA, C-SCALE, Disaster mitigation and agriculture, Lethe, and PaNOSC.

Alongside the existing Notebooks, a new service named EGI Replay² was introduced in the EGI portfolio and onboarded in the EOSC portal. Replay is based on the BinderHub technology and allows for replicating previous research with reproducible environments.

The Notebooks service was already onboarded in the EOSC Portal³ at the start of the project. During 2021, the installation was migrated from INFN-CATANIA provider to CESNET with new hardware and updated underlying kubernetes. The service improved its integration with EGI Check-in for better control of the authorised users including profiles enabling access to dedicated hardware and software options. The default RAM and CPU availability was increased from 1GB RAM/1 core to 6 GB RAM/2 cores and a wider range of development environments was added for all users: MATLAB, RStudio, Julia and an extensive collection of libraries for data analytics and Machine Learning. The service was integrated with the EGI DataHub and EUDAT's B2DROP to offer users seamless access to data. The Software Distribution service was also integrated with Notebooks to facilitate access to a wide range of software libraries.

The Notebooks service also started a pilot to push accounting information towards the EOSC accounting system, which successfully sent records periodically using the available APIs.

² https://marketplace.eosc-portal.eu/services/eosc.egi-fed.egi_replay

³ <https://marketplace.eosc-portal.eu/services/egi-notebooks>

2.2 EGI – DIRAC

Description	EGI Workload Manager
Task	4.2
URL	
Service Category	Platform Service
Service Catalogue	https://www.egi.eu/services/workload-manager/
Location	CC-IN2P3/Lyon
Duration	M1-M30
Modality of access	Access for registered EGI users
Support offered	Full user support including documentation and training. Support is offered via the EGI Helpdesk.
Operational since	2014
User definition	All sizes user communities from large HEP communities to multidisciplinary and long tail communities.

2.2.1 Metrics

Metric name	Baseline	Define how measurement is done	Period 1 M1-M5	Period 2 M6-M10	Period 3 M11-M15	Period 4 M16-M20	Period 5 M21-M25	Period 6 M26-M30
No of groups or experiments	17	Internal provider configuration	16	16	18	18	18	20
No of registered users	500	Internal provider database/logs	737	740	735	742	746	747
No of execute jobs	10,000,000	Accounting	4,000,000	5,000,000	4,500,000	5,630,000	6,082,000	6,280,000
No of countries where jobs executed	12	Accounting	12	12	12	12	12	14
Name of countries where jobs executed	NL, IT, PL, DE, US, BE, FR, CN, PT, ES, SK, UK, GR	Accounting	NL,FR,BE,DE,IT,PT,ES,US,PL,CZ,RO,UK	NL,FR,BE,DE,I T,PT,ES,US,P L,CZ,RO,UK	NL,FR,BE,DE,IT,P T,ES,US,PL,CZ,R O,UK	NL,FR,BE,DE,IT,PT, ES,US,PL,CZ,RO,U K	NL,FR,BE,DE,IT,PT, ES,US,PL,CZ,RO,U K	NL,FR,BE,DE,IT,PT ,ES,US,PL,CZ,RO, UK,RU,CN

2.2.2 Assessment

The EGI - DIRAC installation supports the EGI Workload Manager service to manage and distribute computing tasks in an efficient way in the distributed EGI infrastructure. During EGI-ACE, the service was consolidated into a single provider (CC-IN2P3) that took over the responsibility of operating the service for the project, involving DIRAC experts and developers in the support and maintenance of the service. The service is published in the EOSC Marketplace⁴.

During the last period, the installation has increased the number of groups of experiments supported up to 20 (33% increase over the previous project), starting from 16 communities as baseline. The Virtual Organisations supported at M30 are: acc-comp.egi.eu, dteam, auger, biomed, fedcloud.egi.eu, enmr.eu, lofar, km3net.org, lsst, opencoast.eosc-hub.eu, training.egi.eu, virgo, eiscat.se,

⁴ <https://marketplace.eosc-portal.eu/services/egi-workload-manager>

vo.access.egi.eu, vo.grand-est.fr, vo.hess-experiment.eu, beapps, vo.usegalaxy.eu, virgo, vo.grand-est.fr, vo.sbg.in2p3.fr, vo.france-grilles.fr and vo.complex-systems.eu. These VOs support a total number of 747 users (118% over baseline) from 14 different countries (+2 increase since last period) who executed a total of 31.4 Million jobs (133% increase over last period)

The service has started its integration with the EGI Check-in for Authentication and Authorization of users and it is already capable of using Check-in in the web portal interface. The integration will be completed as the underlying compute resources of the EGI High Throughput Compute service transitions from X.509 certificates to token based authentication supported by Check-in. DIRAC is actively participating in the transition campaign for this service.

The service has been presented in several workshops and a webinar will be planned for the upcoming period to further promote the installation and engage with new communities.

2.3 CSIC DEEP training facility

Description	Distributed training facility for Machine Learning, Artificial Intelligence and Deep Learning models hosted at CSIC. This service offers a set of tools to build and train Machine Learning, Artificial Intelligence and Deep Learning models in distributed e- Infrastructures. Ready to use models are available for transfer learning or reuse.
Task	4.3
URL	
Service Category	Platform Service
Service Catalogue	
Location	CSIC, ES
Duration	M1-M30
Modality of access	Free at point-of-use. Additional terms: https://confluence.deep-hybrid-datacloud.eu/display/DS/Terms+of+Use

Support offered	Support is offered via the EGI Helpdesk. Detailed documentation about service, APIs, user guides, tutorials, etc. available.
Operational since	Jan 2018
User definition	Single researchers, small and big communities

2.3.1 Metrics

Metric name	Baseline	Define how measurement is done	Period 1 M1-M5	Period 2 M6-M10	Period 3 M11-M15	Period 4 M16-M20	Period 5 M21-M25	Period 6 M26-M30
No of users communities	17	Internal provider configuration	0	0	1	4		7
ML training cycles in CPU/GPU hours	350,000	Accounting	0	0	4,128	97,420		3,638,292
No of countries reach	6	Check-in	0	0	1	3		3
Names of countries reach	US, SP, PT, DE, FR, UK	Check-in	N/A	0	France	FR, USA, ES, IT, EU-wide		FR, USA, ES, IT, EU-wide

2.3.2 Assessment

The CSIC DEEP training facility is a new service added to the EGI ecosystem as part of EGI-ACE. During the initial months of the project, it has integrated with EGI Check-in and EGI Configuration Database (GOCDDB). Further integration into EGI was not pursued as the

service is transitioning into the EOSC4AI platform⁵ and will be decommissioned in the future. This installation is registered in the EOSC Marketplace within a single entry⁶ that also covers installation described in section 2.4.

The service was promoted in a dedicated webinar (see section 3 - Dissemination), triggering the interest from the EMPHASIS community from the 4th Open Call for use cases of the project. This community has consumed 4,128 CPU/GPU hours over M11 to M15 with users coming from France. Thanks to the dissemination effort from the provider, at the end of the project, the service onboarded a total of 7 user communities, matching the target for the project. These communities consumed more than 3 million CPU/GPU hours on the infrastructure supported by the WP3 installations

2.4 LIP DEEP training facility

Description	Distributed training facility for Machine Learning, Artificial Intelligence and Deep Learning models hosted at LIP. This service offers a set of tools to build and train Machine Learning, Artificial Intelligence and Deep Learning models in distributed e- Infrastructures. Ready to use models are available for transfer learning or reuse.
Task	4.3
URL	
Service Category	Platform Service
Service Catalogue	
Location	LIP, PT
Duration	M1-M30
Modality of access	Free at point-of-use. Additional terms: https://confluence.deep-hybrid-datacloud.eu/display/DS/Terms+of+Use

⁵ <https://eosc-portal.eu/ai4eosc>

⁶ <https://marketplace.eosc-portal.eu/services/deepaas-training-facility>

Support offered	Support is offered via the EGI Helpdesk. Detailed documentation about service, APIs, user guides, tutorials, etc. available.
Operational since	Jan 2018
User definition	Single researchers, small and big communities

2.4.1 Metrics

Metric name	Baseline	Define how measurement is done	Period 1 M1-M5	Period 2 M6-M10	Period 3 M11-M15	Period 4 M16-M20	Period 5 M21-M25	Period 6 M26-M30
No of users communities	17	Internal provider configuration	0	0	0	0	0	0
ML training cycles in CPU/GPU hours	290,000	TBD	0	0	0	0	0	0
No of countries reach	6	Check-in	0	0	0	0	0	0
Names of countries reach	US, SP, PT, DE, FR, UK	Check-in	-	-	-	-	-	-

2.4.2 Assessment

The LIP DEEP training facility complements the CSIC DEEP training facility by supporting the storage associated with the use cases supported by the DEEP service, as such the metrics are reported in the previous installation (section 2.3.1). Technical integration is identical and the installation is part of the same entry in the EOSC Marketplace.

2.5 DODAS

Description	<p>DODAS allows to instantiate on-demand complex infrastructures over any cloud with almost zero effort and with very limited knowledge of the underlying technical details. In particular, DODAS provides the end user with all the support to deploy from scratch a variety of solutions dedicated (but not limited) to scientific data analysis. DODAS provides two principal baselines ready to be used and to be possibly extended and customised</p> <ul style="list-style-type: none"> - HTCondor batch system, possibly federated, and integrated with caching mechanism - Spark+Jupyter cluster for interactive and big-data analysis with persistent storage solutions.
Task	4.4
URL	
Service Category	Platform Service
Service Catalogue	
Location	INFN-CNAF e INFN-Bari
Duration	M1-M30
Modality of access	Free at the point of use
Support offered	Support is offered via the EGI Helpdesk. Reference for tutorial https://dodas-ts.github.io/HandsOn-INFN-2019/ , General Documentation: https://dodas-ts.github.io/dodas-doc ; Technical Documentation https://dodas-ts.github.io/docs-templates/
Operational since	Jan 2018
User definition	Researcher, Small communities, big communities and resources providers

2.5.1 Metrics

Metric name	Baseline	Define how measurement is done	Period 1 M1-M5	Period 2 M6-M10	Period 3 M11-M15	Period 4 M16-M20	Period 5 M21-M25	Period 6 M26-M30
No of users using clusters	30	Internal logs	0	2	3	4	4	4
CPU/hours	1,800,000	Accounting	0	2,294.8	464,379	742,941	999,056	788,478
No of clusters deployments	900	Internal logs	0	5	31	12	2	2
No of visits	90	Internal logs	0	10	35	8	3	10

2.5.2 Assessment

DODAS is one of the new services brought to EGI with the start of the EGI-ACE project. The service was already registered in the EOSC Marketplace⁷ and during this first period of the project was integrated with the EGI core services. Currently DODAS is fully integrated with Check-in, Helpdesk, Configuration Database and ARGO Monitoring. DODAS became integrated with the PaaS Orchestrator (installation in WP6) to build a powerful service which can be then exposed via the PaaS dashboard (<https://indigo-paas.cloud.ba.infn.it>). In this way DODAS service can be exposed to communities through a more user-friendly interface. DODAS also has a new section on the EGI documentation portal that introduces the service for EGI users⁸.

During this period DODAS has been requested by the FERMI-LAT user community that applied to the 3rd Open Call for use cases of EGI-ACE. This community has consumed 2,533,000 CPU hours over the reporting period (a growth of 910% from last period) and have triggered the deployment of 52 clusters for running their analytics. This is below the target (100 deployments) due to the stability of the available clusters and the low number of users engaged with the service.

⁷ <https://marketplace.eosc-portal.eu/services/dynamic-on-demand-analysis-service-dodas-portal>

⁸ <https://docs.egi.eu/users/compute/orchestration/dodas/>

3 Dissemination

In this section we report the list of events in the context of EGI-ACE that WP4 partners contributed during the project, reporting the number of attendees to measure the possible user interests .

Table 2: Dissemination activities related to WP4 installations

Type of Activity	Title	Date	Name of Event	Location	Type of Audience	Reach	Scale
Presentation	Infrastructure services: Highlights from the compute, data, security areas	2021/02/05	EGI-ACE public launch event	Online	Scientific Community, General Public	250	worldwide
Workshop	EGI-ACE Communities Workshop	2021/02/16	EGI-ACE Communities Workshop	Online	Scientific Community, General Public	150 people	worldwide
Presentation	EGI and FG DIRAC services Development in EGI-ACE project	2021/05/11	Virtual DIRAC Users' workshop	Online	Developers, service admins, users	55	Worldwide
Webinar	Access and analyze data from the EGI	2021/05/12	EGI Webinar 2021	Online	Scientific communities , and	Num. of Participants: 27	worldwide

	DataHub with Jupyter notebooks and MATLAB				programmers who support research and education.	Num. of Countries: 10	
Webinar	Analyze your data using DODAS generated cluster	2021/09/22	EGI Webinar 2021	Online	Scientific communities, developers, integrators and end users	Num. of Participants: 11 Num. of Countries: 8	worldwide
Presentation	Delivering Services and Solutions - Workshop (EGI Workload Manager Service)	2021/10/21	EGI Conference 2021	Online	IT providers, Research Community reps.	40 People	Global (mostly European)
Webinar	How to train your AI model in EOSC	2021/12/01	EGI Webinar 2021	Online	User communities want to use GPUs in Clouds.	Num. of Participants: 32 Num. of Countries: 9	worldwide
Presentation	Les services DIRAC au	2021/12/15	JCAD2021	Online	IT providers, Research	~100 people	European (mostly

	CC-IN2P3 (Dirac services at CC-IN2P3)				Community reps.		French)
Presentation	Scalable Environments for reproducible Open Science	2022/04/06	MiniGateways 2022	Online	Scientific user communities	~100 people	worldwide (mainly US)
Presentation	EGI Workload Manager Service	2022/05/10	Virtual DIRAC Users' workshop	Online	Developers, service admins, users	46	worldwide
Presentation	DODAS service	2022/05/27	Healthycloud discussion on orchestrators	online	Researchers	9	EU
Training	Reproducible Open Science With Big Data - The EGI Notebooks and Binder	2022/09/23	EGI Conference 2022	Prague	Researchers	10	Europe

	services						
Training	Platform for distributed, big computing - DIRAC User Group meeting	2022/09/20	EGI Conference 2022	Prague	Researchers	10	Europe
Training	How to deploy ready-to-use BigData Platform on top of the EOSC Compute Platform - the DODAS solution	2022/09/23	EGI Conference 2022	Prague	Researchers	10	Europe
Presentation	EGI Notebook and Replay Services	2023-03-15	APAN55	Nepal (remote)	Researchers involved in the Disaster Mitigation sciences	16	Asia Pacific
Presentation	EGI Notebook and Replay Services	2023-03-24	ISGC2023	Taipei, Taiwan	Researchers, Students	10	Worldwide

Demo	EGI & C-SCALE: Notebooks for Earth Observation	2023/06/21	EGI Conference 2023	Poznan	Researchers	20	Europe
Demo	EGI Workload Manager Service: activities in the EGI-ACE project	2023/06/21	EGI Conference 2023	Poznan	Researchers	20	Europe
Training	Getting started with the EGI Workload Manager Service	2023/06/23	EGI Conference 2023	Poznan	Researchers	20	Europe
Training	Processing Data from EOSC on EGI Compute resources training	2023/06/23	EGI Conference 2023	Poznan	Researchers	20	Europe

4 Satisfaction

In this chapter we report the satisfaction on the WP4 installations are reported by EGI Customer interviews and the number of orders coming from the EOSC Marketplace.

4.1 EGI Customer satisfaction reviews

EGI regularly interviews Communities using the services with an active SLA, in order to measure the satisfaction and discuss possible issues. The level of satisfaction is measured from 1 (min) to 5 (max). For what concerns the first period of the EGI-ACE project the communities using EGI-ACE WP4 services interviewed are reported in table 3.

Table 3: Communities interviewed during the last 15 months of EGI-ACE project

Community	WP4 installations used	Level of satisfactions and comments	Issues/feature requests reported with WP4 installations
WeNMR	EGI - DIRAC	5. Very satisfied	Nothing to report even if DIRAC will remain a critical service for this community.
BioMed (VIP)	EGI - DIRAC EGI Notebooks	5. Very satisfied Services run smoothly, less and less issues, things are going well	The replacement of X.509 with tokens requires help, not sure what should be done, need explanation and help with the transition
EISCAT_3D	EGI Notebooks EGI - DIRAC	5 Very satisfied	EGI Notebook: should be auto updated but at the moment it doesn't. Error happens

4.2 EOSC Marketplace orders

For the services that have been registered on the EOSC Marketplace, we report here the statistics of the orders during the 30 months of the project.

Table 4: Number of Orders from the EOSC Marketplace related to WP4 installations

EOSC Marketplace Service - WP4 installation(s)	Number of orders
EGI Notebooks - EGI Notebooks	8
EGI Workload Manager - EGI DIRAC	1
DEEP training facility - CSIC DEEP training facility & LIP DEEP training facility	12
Dynamic On Demand Analysis Service	1