

Inspired

ISSUE 24
JULY 2016

news from the EGI community



TOP STORIES

EMSODEV pilot on the FedCloud

page 2

Engage Year 1: impacts and outcomes

page 3

Latest numbers: EGI beyond 650k CPUs

page 5

EISCAT-3D goes for DIRAC4EGI

page 7

EGI's CONTRIBUTIONS TO PROJECTS

- 08 AARC and EGI: first year of work
- 09 EGI and INDIGO-DataCloud integration
- 11 EDISON: Building the data science profession

MORE

- 01 Computing centres: Okeanos
- 06 DI4R 2016 - Registration is Open!



Engage - Grow - Innovate

www.egi.eu

Welcome to issue 24!

In this Summer edition of our newsletter, we cover the results of the first years of work of many of the projects with EGI involvement. Read all about it!

Your feedback and suggestions are always welcome!

Send an email to Sara & Iulia at: press@egi.eu



EGI's next event will be the Digital Infrastructures for Research 2016 in Kraków Poland. The event is co-organised by EUDAT, GÉANT, OpenAIRE and RDA Europe

Computing centres: Okeanos

Kostas Koumantaros presents the GRNET cloud infrastructure

GRNET offers cloud computing services to all members of the Greek research & education community via the Okeanos infrastructure.


With Okeanos, users can create a multi-layer virtual infrastructure and virtual computing machines, local networks to interconnect them, and a reliable storage space within seconds.

Thousands of academic users have already used virtual machines in the course of their research, experimental, educational or other activities. The cloud computing infrastructure and services of GRNET S.A. have been made available to the pan-European R&E community via the 'okeanos-global' service. A part of Okeanos is available through open standards such as OCCl and CMDI in the EGI Federated Cloud.

Okeanos makes it possible to:

- > Create virtual machines with custom images

- > Use extra resources for limited time frames e.g. to create virtual laboratories for a project
- > Use advanced and diverse security policies
- > Manage the virtual machines and networks via a user-friendly interface
- > Synchronise user files through the storage capabilities offered (Dropbox-like functionality)
- > Share files among different users
- > Backup as a Service (BaaS)
- > Image as a Service (ICaaS)
- > ~ORKA! a cloud-based platform for big data analytics



The Okeanos computing cluster

- > The Lambda on-demand framework for ingesting real-time streams with efficient stream and batch analytics

More information

GRNET
<http://www.grnet.gr>

Okeanos
<http://okeanos.grnet.gr>

Lambda
<https://lambda.grnet.gr>

EMSODEV pilot running on the EGI Federated Cloud

The EMSODEV team writes about the first step towards a production Data Management Platform

The EMSODEV project started a collaboration with EGI to develop cloud services and resources for the design and implementation of the EMSODEV Data Management Platform (DMP). EMSODEV developed the DMP platform on top of the EGI Federated Cloud with the EGI support on virtualisation, storage, networking and security. The experimental prototype of the DMP is now fully integrated with the EGI Federated Cloud services and is hosted at the INFN Bari site.

The design of the DMP is almost complete. The platform enables the exploitation of capabilities of e-Infrastructures to develop a flexible and scalable Data Management Service for a long-term, high-resolution, (near)-real-time monitoring by providing a coordinated approach for data capture, archiving, management and delivery based on OGC standards.

The DMP is based on Hadoop



and includes a set of common services, compliant to the ENVRI Reference Model.

Sensors data will be coming both in asynchro-nous/batch and real-time mode, from two different data sources: a Sensor Observation Service (SOS) deployed at each EGIM node and the EMSO regional nodes. Both historical time series and real time data will be accessible and searchable via Application Programming Interfaces (APIs)

built on top of NoSQL databases (e.g., HBase, OpenTSDB).

EGI and EMSODEV are moving towards a production environment for the DMP. The first step is a Service Level Agreement to guarantee IT resources and services to run the DMP for a long period. In the meantime, the technical collaboration will continue to investigate in particular, federated deployment of DMP in multiple EGI sites.

EMSODEV & DMP

EMSODEV is a Horizon 2020 project set up to implement the European Research Infrastructure EMSO, a partnership between France, Greece, Ireland, Italy, Portugal, Romania, Spain and the United Kingdom.

EMSODEV focuses on the development, testing and deployment of an EMSO Generic Instrument Module (EGIM) which will ensure accurate long-term measurements of ocean parameters. Collecting and analysing ocean data on a day-to-day basis is challenging, but crucial to monitor environmental processes. The Data Management Platform (DMP) will implement Data Acquisition and Transformation, Data Curation, Data Processing and Data Access phases collect and analyse parameters, to address the needs of a diverse user community.

More information

EMSO: European Multi-disciplinary Seafloor and water-column Observatory

<http://www.emso-eu.org>

This article was contributed by :

Lucio Badiali (INGV),
Pasquale Andriani and
Massimiliano Nigrelli
(Engineering) and Diego
Scardaci (EGI Foundation)

First year of EGI-Engage: impacts and outcomes

Tiziana Ferrari summarises the achievements of the first project year

The EGI-Engage project was set up to expand the capabilities of a European backbone of federated services for advanced scientific computing, storage, data management and knowledge transfer.

The first year concluded in February 2016 and I am pleased to report that we are steadily progressing towards our goals. At the end of the project we aim to have secured and number of impacts and here is a summary of how the work of the first year of EGI-Engage contributed to this.

O1: Ensure the continued coordination of the EGI Community in strategy and policy development, engagement, technical user support and operations of the federated infrastructure in Europe and worldwide.

> Strategy. The EGI Council adopted the new EGI strategy in May 2015 and its new vision, mission and strategic goals with objectives organised in five strategic themes.

> Governance. The EGI Council approved of a new governance model and new statutes for the EGI Foundation. The new model

is more flexible and more welcoming to national and international organisations and user communities.

> The Open Science Commons vision was endorsed in May 2015 by the European Council and the vision was adopted by the main European e-Infrastructures (EUDAT, GEANT, OpenAIRE) and LIBER as the foundation of the new European Open Science Cloud initiative (EOSC) of the EC.

> The EGI Council and all EGI-Engage RIs provided input to the EGI contribution to the European Open Science Cloud consultation.

> Security. The Security Incident Handling Procedure was updated and approved by the EGI OMB. This work exploited new possibilities in incident response required by the newly integrated technologies, primarily cloud IaaS.

O2: Evolve the EGI Solutions, related business models and access policies for different target groups aiming at an increased sustainability of these outside of project funding. The solutions will be offered to large and medium size RIs, small research communities, the long-tail of science, education, industry and SMEs.

> New Business Engagement Programme.

> Thematic solutions. We started to set up procedures for SLA negotiation with providers of thematic services and user communities.

KPI 2 Number of RIs and e-Infras integrated with EGI

15



> New service portfolio approved by the EGI Council in November 2015.

> EGI marketplace: first concept defined. The first prototype of an online marketplace platform is expected in late 2016.

> The EGI Federated Cloud now supports a larger number of user communities and the OCCI standard was extended to provide new capabilities. Cloud federation via native OpenStack is now possible.

> The pilot for a Long Tail of Science platform was created to provide simplified access procedures to individual users or small research groups

> Accelerated Computing is being developed as a new platform to provide distributed access to local GPGPU computing facilities.

> The High-Throughput Data Analysis solution saw eight new releases of the Unified Middleware Distribution (UMD).

> We established 27 collaborations with RIs/FETs, 11 projects/communities and numerous individual researchers.

KPI 5 Number of researchers served by EGI



46,250

Objective 3 (O3): Offer and expand an e-Infrastructure Commons solution

- > New AAI architecture and prototyping of technical solutions. Piloting activities with research communities, primarily ELIXIR, allowed the integrated testing of the new services in the context of the AAI of a large RI.
- > Service request management and resource allocation. The pay-for-use group produced a beta production version of the e-GRANT tool to support allocation to user groups.
- > Technical roadmap for the accounting and monitoring systems (backend and portal). We launched a redesigned accounting portal we defined a multi-tenant provisioning model aiming at offering the monitoring system (ARGO) as a service.
- > The GOCDB was enhanced to support multiple projects and extend its data model. GOCDB was integrated with the EBI IdP to demonstrate the tool capabilities to the ELIXIR community.
- > The Operations coordination defined a new set of core activities and services to support the EGI federation, which were assigned to the service providers through a bidding process.

KPI 14 Compute CPUs available to research



651,748

Objective 4 (O4): Prototype an open data platform and contribute to the implementation of the European Big Data Value.

- > The EGI Foundation became a member of the Big Data Value Association (BDVA), an industry-led consortium for the implementation of the Big Data Value public-private partnership.
- > The EGI Data Hub was defined as a new EGI service to the development of the Data Commons across different scientific disciplines and support the big data value chain by facilitating the commercial exploitation of research data.
- > We defined use cases and the architecture for the Open Data Platform and started the preparation of the prototype.
- > The EGI-EUDAT collaboration continued to provide end-users with a seamless access to paired data and HTC resources. The roadmap is defined by a set of user communities who are already collaborating with both infrastructures in the field of Earth Science (EPOS and ICOS), Bioinformatics (BBMRI and ELIXIR) and Space Physics (EISCAT-3D).
- > A MoU was signed with the PRACE-4IP project.

Objective 5 (O5): Promote the adoption of the current EGI services and extend them with new capabilities through user co-development

- > We established 8 Competence Centres (CCs) linked to 8 RIs and communities: BBMRI, DARIAH, EISCAT-3D, ELIXIR, EPOS, MoBrain/INSTRUCT, LifeWatch, and disaster mitigation.

KPI 17 Number of services, pilots and projects on EGI for SMEs

36

10101
01110
10101



- > We developed Service Level Agreements as a framework to formalise collaborations between user communities (VOs) and EGI providers. So far, four communities have benefitted from this framework:
 - >> MoBrain, BILS, DRIHM and Terredue
- > We expanded the network of cloud user support teams across the NGIs. Now we have 18 national user support teams from 12 countries:
 - >> Czech Republic (CESNET), Croatia (SRCE), France (CNRS), Greece (GRNET, IASA), Hungary (MTA SZTAKI), Italy (INFN Bari, INFN Padova), Macedonia (UKIM), Poland (CYFRONET), Portugal (LIP), Slovakia (IISAS), Spain (BIFI, BSC, CESGA, CIEMAT), Sweden (KTH);
- > 8 Virtual Organisations were set up in EGI for gaining access to cloud resources

More information

EGI-Engage

<http://go.egi.eu/eng>

Tiziana Ferrari is the Technical Director of the EGI Foundation

EGI infrastructure expands beyond 650,000 CPU cores


Alessandro Paolini rounds up the latest EGI figures

The EGI infrastructure builds on 15 years' experience in design, development, and deployment of distributed data analysis services. Today, the production infrastructure federates hundreds of resource centres that in turn serve hundreds of research communities.

The federation members

EGI's main resource providers are the 24 EGI Council Members: 22 national e-Infrastructures, CERN and EMBL. The EGI infrastructure integrates resources provided by an additional 34 worldwide organisations.

In addition, EGI collaborates with the US's Open Science Grid and Compute Canada to support common user communities through inter-operational agreements.



Storage and compute capacity (CPU cores) available in EGI 2011-2016

The federated services

The services offered by the EGI resource centres can be grouped in two macro-categories: High Throughput Computing (HTC) and Cloud Computing.

The overall total capacity available for the research communities supported by EGI has more than doubled in the

past six years, surpassing the threshold of 650,000 cores during 2016.


In terms of storage, EGI federates more than 260 PBytes of disk storage and nearly 240 PBytes of tape storage.

These numbers make EGI the largest international e-infrastructure in Europe.



EGI Council members

Federated Resource Centres



The top-five contributors to the computing capacity of EGI are three National Grid Initiatives (Germany, Italy and the United Kingdom), CERN and the Nordic Data Grid Facility (an international collaboration between Denmark, Finland, Norway and Sweden).

Use of the resources

The resource providers federated in EGI support more than 200 Virtual Organizations (VO). We estimate that EGI is serving a total of 46,000 users: users directly registered in VOs and users who access EGI resources through portals and other high-level services.

Users of the EGI Federated Cloud instantiated 343,200 Virtual Machines from January 2015 to January 2016. This represents 2.31 million hours of

CPU wall time consumed and a 79% yearly increase. The federated cloud, with its 21 providers of which one is commercial, is a new platform that started its production activities in May 2014.

Together with 580 million HTC computing jobs this means that, last year, the EGI users consumed more than 4,500 million CPU core hours.

Top VOs by resource consumption (million CPUs)

atlas / 1,210	}	LHC experiments
cms / 1,033		
alice / 605		
lhcb / 513		
belle / 211	}	biomedical sciences Cherenkov Telescope Array gravitational waves
biomed / 47		
vo.cta.in2p3.fr / 21		
virgo / 11		

More information

Alessandro Paolini is a member of the Operations Team

Digital Infrastructures for Research 2016: 28-30 Sept

Registration for the conference is now OPEN!

The Digital Infrastructures for Research 2016 (28-30 Sept, Kraków) is the first user forum jointly organised by Europe's leading e-infrastructures and projects: EGI, EUDAT, GEANT, OpenAIRE and RDA Europe. Registration for the event is now open and discounted early-bird rates will be available until 5 August 2016.

The **keynote speakers** of the conference will be:

> **Karlheinz Meier** from Heidelberg University, who will speak about "The brain, the universe and the need for integrated infrastructures" (Wednesday 28 September)

> **Laurie Goodman** from the open-data journal GigaScience, who will speak about "Data Citation: Credit Where Credit is Due"

The DI4R 2016 conference is hosting three co-located events (to take place on Tuesday, 27 September):

> **3rd WISE Workshop** - Wise Information Security for collaborating E-infrastructures
> **Design your e-Infrastructure Workshop** - where user communities can get together with experts to discuss requirements and design the e-infrastructure they need for their scientific use case.

> **FitSM Foundation Training** - aims to provide the basic IT service management concepts and terms

The DI4R 2016 is hosted by ACC Cyfronet AGH between 28 and 30 September in Krakow, Poland.

More information

DI4R 2016

www.digitalinfrastructures.eu

Early bird registration:
5 August

EISCAT-3D goes for the DIRAC4EGI Service

Yin Chen gives the latest updates on the EISCAT-3D Competence Centre

Having understood the requirements of the EISCAT-3D Competence Centre, we helped them to select DIRAC4EGI as the technology for they need for the upcoming implementation phase.

The EISCAT-3D Competence Centre (CC) was established within the EGI-Engage project in March 2015 to develop a web-based user interface for searching, retrieving and re-processing (visualisation, analysis) data products generated by EISCAT-3D radar stations.

The EISCAT-3D radar is a next generation incoherent scatter radar system under construction by EISCAT association. Once assembled, it will be a world-leading infrastructure using the incoherent scatter technique to study the atmosphere in the Fenno-Scandinavian Arctic and to investigate how the Earth's atmosphere is coupled to space.

The challenge of the EISCAT-3D data management system is to find an efficient way to handle experimental data that will be generated at great speeds and volumes. During its first operation stage in 2018, EISCAT-3D will produce 5PB data per year, and the total data volume will rise up to 40PB per year in its full operations stage in 2023.

After analysing the functional and performance requirements of the data portal, we proposed to use the DIRAC4EGI framework for the development.

DIRAC is a software framework for building distributed



computing systems. The DIRAC project began in 2002 within the LHCb High-Energy Physics collaboration at CERN to manage Monte Carlo simulations and was later extended to cover all distributed computing operations including raw detector data distribution, data management and end-user analysis. The DIRAC project introduced several innovations focusing on the needs of large and distributed scientific collaborations. Today it is being used or evaluated by a number of global research infrastructures, such as Belle II, BES III and CTA.

The DIRAC4EGI service was developed in 2014 to integrate DIRAC in the EGI FedCloud infrastructure.

With the support of the DIRAC technical team, an EISCAT data portal is customised using DIRAC4EGI and will be up and running within a couple of months time. The live EISCAT portal running on top of the EGI FedCloud was successfully shown to the participants of the EISCAT_3D user meeting, 18th -19th May 2016 in Uppsala, and

the demonstration of metadata-based data searching, file downloading and plotting of data works well and received good feedback from the community.

The data portal's File Catalogue allows researchers to search metadata, download and plot the selected data files. Selected data files can be added to processing algorithms and submitted as a job to generate data plots.

In the next step, new application features will be added to the data portal to support various requirements from EISCAT researchers.

More information

EISCAT-3D Competence Centre

https://wiki.egi.eu/wiki/CC-EISCAT_3D

Yin Chen is a member of the User Community Support Team

AARC and EGI: first year of work

Peter Solagna writes about EGI's contribution to the AARC project



The EC-funded project AARC started in 2015 to promote interoperability of Authorisation and Authentication systems across e-Infrastructures and research collaborations. The results of the first AARC year of work concentrated in the following objectives:

Design the AARC architecture

The AARC team has designed a high-level AAI architecture blueprint based on requirements and expertise of many e-infrastructures, research infrastructures (RIs), research communities and AAI architects.

The EGI Foundation led the process to gather requirements from the communities, including many Research Infrastructures in the ESFRI roadmap.

As a result of these discussions, AARC drafted an initial 'AARC Blueprint Architecture' highlighting a list of building block requirements common among all stakeholders and presents various implementation possibilities.

Harmonise procedures and policies

One of the hurdles on the way of improved collaboration and integration is incident response. Although computer security incident response procedures often exist at a local level or at federation level, there is no best practice guidance for security incidents involving several federations, for example an IdP federation and an SP federation,

spreading across multiple administrative domains.

The AARC project, together with a number of stakeholders including the EGI federation direct involvement, has contributed to The Security Incident Response Trust Framework for Federated Identity (Sirtfi).

Sirtfi offers guidelines for incident behaviour to determine whether an organisation can be effective in incident response.

Develop training modules for different user communities

AARC organised a training at the University of Manchester, from 15 to 16 of March, dedicated to organisations that provide resources and services for the ELIXIR (life sciences) and DARIAH (art and humanities) RIs. The main objective was to help service providers to make their resources and services available to the users of the RIs.

Pilot policy frameworks and critical components

The pilots for the e-infrastructures and research infrastructures focused on enabling federated identity on distributed service providers. This activity is led by the EGI Foundation is driven by the

requirements and use cases of the EGI federation. The ultimate goal of this pilot is to demonstrate how collaborations can use their institutional credentials to both manage their organisation and access distributed services.

AARC and EGI

The most important result of the collaboration is the architecture for the new EGI AAI services, which are fully aligned with the AARC blueprint and in the best position to interoperate with new and existing communities and infrastructures. The smooth integration with the ELIXIR AAI, achieved in May 2016, is a good example.

EGI has now more information about the AAI requirements of the research communities, thanks to our involvement in the gathering process during the early stages of AARC.

More information

AARC Project
<https://aarc-project.eu/>

Peter Solagna leads the EGI Operations team

EGI and INDIGO-DataCloud integration

Davide Salomoni and Giacinto Donvito on the EGI & INDIGO integration roadmap

INDIGO-DataCloud is an EU-funded project set up to develop a new cloud software platform for the scientific community, deployable on multiple hardware and provisioned over hybrid (public/private) e-infrastructure. The INDIGO team has recently completed a deliverable outlining the possible roadmaps for technical integration with different infrastructures. The full "Technological Methods of Integration with other e-Infrastructures" deliverable is available online. One of the chapters is dedicated to potential routes of technical integration between EGI & INDIGO.

Authentication and authorisation

AA in EGI is commonly based on personal X.509 certificates and on the user's membership in a Virtual Organisation: the community attributes are managed by the Virtual Organisation Management System (VOMS) that issues an attribute certificate (AC) describing the user membership. EGI AAI is evolving and currently alternate AAI technologies are available for users: username/password authentication based on SAML and OIDC protocols, support for third parties attribute management, and credential translation services.

The INDIGO Identity and Access Management (IAM) will provide a set of advanced interoperable solutions:

- > a login service (based on OpenID-Connect) able to federate different identity providers
 - > a group membership service that will allow to group users in organisations
 - > an authorisation service that provides the tools to define and enforce authorisation policies over the protected resources
 - > support for controlled delegation of privileges across services
 - > a Token Translation Service, bridging the gap between Authentication mechanisms that do not support OpenID-Connect.
- INDIGO IAM supports the legacy authentication technologies as well as the protocols recently added to the EGI AAI offer. The flexibility of the INDIGO tools will allow an easier integration and support for potentially all the use cases by providing end-users with a unified view of identities and privileges, and the tools needed to enable a secure composition of services from multiple providers in support of scientific applications.

Computing

Grid-job submission via the INDIGO platform

Integrating INDIGO capabilities with the EGI grid infrastructure will allow users to submit jobs and manage virtual machines via already available portals using the advanced features provided by the INDIGO platforms. Users



communities could have an easy and straightforward way of instantiating a community portal in a cloud environment letting the INDIGO PaaS layer to deal with high-availability and scaling e.g. to accommodate unpredictable workloads.

Moreover, already available portals could easily exploit the new INDIGO IaaS and PaaS features by implementing simple APIs provided by the INDIGO SaaS layer.

Application/service deployment on cloud via the INDIGO platform

The scientific communities that need to run a long running service or an application on the EGI Federated Cloud will be able to deploy them in a simple and transparent way, using both APIs and web user interfaces, thanks to the functions provided by the INDIGO platform. The

integration in the EGI Federated Cloud will occur at both IaaS and PaaS levels and will allow end users to dynamically specify the composition of the cluster of services requested for a given use-case or application such, as for example, databases, load balancers, automatic scalability features or proximity to certain datasets.

The INDIGO PaaS will then compose the requests and automatically instantiate the needed services taking care of monitoring the status of each service, automatically scaling the computational resources needed, and restarting the services in case of failure, exploiting for example the IaaS resources available via the EGI Fed Cloud.

Storage

EGL grid storage infrastructure

The storage management of the EGI grid infrastructure is based on a specific class of services called "storage elements". The INDIGO project extends this concept and makes it possible to access files on the EGI grid storage elements in terms of both transfers and read access, exploiting the protocols supported by EGI and providing a unified view of storage resources. End users exploiting the INDIGO platform will be able to import/export data transparently from/to the EGI infrastructure and process it using application and/or services deployed through the INDIGO PaaS layer. The INDIGO PaaS Layer will take care of doing the needed data transfers on behalf of the users, before/after executing the requested application/service.

EGL cloud storage infrastructure

The storage layer of the EGI Federated Cloud provides Block and Object Storage capabilities.

Block Storage: for this type of storage, the INDIGO project will enable seamless use of Block storage resources available at each site. Computing resources instantiated through the INDIGO PaaS will be able to connect to distributed cloud block storage, presenting it to applications as if it was local. Replication and caching facilities are also provided by the INDIGO unified storage layer. Moreover the need of local block storage facilities could be easily exploited by means of the automatism provided by TOSCA Templates and INDIGO PaaS layer. This way, the user can compose a cluster of services requesting the block storage as part of the cluster configuration. All the needed preparatory steps will be automatised by INDIGO services.

Object Storage: in order to access this storage in the EGI Federated Cloud, a CDMI interface is suggested as the standard that each site should adopt, regardless of the technological implementation defined at each site. The INDIGO project is building on top of that a standard CDMI implementation to also support advanced features such as QoS and Data Lifecycle Management so that, for example, different types of



INDIGO - DataCloud

storage technologies (cold or inline storage) or time-dependent access policies to data can be required by users. Moreover, all the available CDMI endpoints could be easily federated into a single geographical storage federation, that is also able to provide seamless access to the storage in order to run legacy applications. The INDIGO components will contribute to the technical realisation of the EGI Open Data Platform, which is a big milestone of the EGI strategy.

Network virtualisation

Regarding networking, integration of INDIGO products may add a new set of capabilities for EGI Federated Cloud users. Contributions by INDIGO will allow them to set up, manage and use private networks on demand in participating cloud sites - a feature that is currently unavailable.

More information

INDIGO-DataCloud
www.indigo-datacloud.eu/

Davide Salomoni is the INDIGO-DataCloud coordinator
Giacinto Donvito leads the project's technical development

Building the data science profession

Themis Athanassiadou writes about the achievements of the EDISON project



A data scientist is an expert in extracting value from existing data and managing the whole data lifecycle, including supporting scientific data e-Infrastructures.

The Horizon 2020 project EDISON will put in place foundation mechanisms that will increase the number of qualified data scientists in Europe, benefiting both academia and industry.

To achieve this, EDISON will:

- > define the skills and competences of the new profession
- > construct Model Curricula to help universities implement courses and programs for data science
- > design a framework for re-skilling and certifying data science skills for graduates and professionals
- > create an online community portal for data science in Europe.

Year 1 achievements

1. Defining skills and competences

EDISON collected and analysed information to establish the competences and skills for data scientists and identified knowledge domains to be included in education and training programmes. This work defines the Competences Framework for Data Science (CF-DS).

The CF-DS describes five groups of competences for data science: data analytics, data science engineering, domain knowledge (cf. NIST definition of data science), and adds two new groups: data management and scientific methods.

The competence framework and its mapping into knowledge domains are critical to the assessment and design of data science curricula.

2. Building model curricula

EDISON produced an inventory of current data science offerings (programs, courses, books) in academia and industry from across the world. Analysis of the inventory revealed both regional variations and competencies and knowledge domains gaps.

The outcomes of this analysis will lead to recommendations for further development of well-balanced curricula, which will be tested next year in three pilot cases.

3. Building a community platform

EDISON is building an online community platform that will include a training marketplace, a community-building area, a job centre and a blog/news portal. A key feature of the portal will be a benchmark to allow dynamic navigation based on user

profiles. The benchmark will identify knowledge gaps, and present training material to help the users achieve their learning goals.

For educators and training organisations, the portal will provide a space to share course information, and allow for authoring on-line material. An important resource will be the virtual labs, and computational resources provided by the EGI FedCloud. They will allow educators to create hands-on tutorials and host datasets on the cloud, which can be accessed by everyone.

EGI in EDISON

EGI is part of the EDISON consortium and will contribute to the data science training activities with the EGI Federated Cloud as environment to host applications and data for training. EGI is also responsible for the certification scheme and for developing the sustainability model of the project outcomes.

More information

<http://edison-project.eu/>

The EDISON project is funded under grant agreement No. 675419.

This article reflects only the author's view and the EC is not responsible for any use that may be made of the information it contains.