Inspired

news from the EGI community

From the first grid jobs to the EOSC

Building the grid: pretending it was easy

EGI in numbers: 2011-2018

And...

EMSO Data Management Platform
EOSCpilot Science Demonstrators
SLAs: EGI's mission into practice
EGI Operations Team: 2019 goals

Advanced Computing for Research

From the first grid jobs to the EOSC

Arjen van Rijn reflects on the challenges ahead

This issue of the EGI newsletter is dedicated to 15 years of operations. The EGI Federation is now an established computing platform for research, with demand growing at an average 20 percent per year.

These 15 years have been memorable, for me personally. In this period the successful BiG Grid project took place from 2007 to 2012: the roll-out of the Dutch distributed computing infrastructure for research.

From 2013 the BiG Grid infrastructure, which included an LHC Tier-1 facility, was adopted by SURF, the Dutch e-infrastructure organization. In 2014 Nikhef acquired co-funding for its share of the Dutch Tier-1 for the period until 2019. Meanwhile, Amsterdam Science Park was selected to host the EGI Foundation, established on 8 February 2010.

I was a member of the first EGI Executive Board until 2014. In 2012, I was appointed one of the Dutch e-IRG delegates and then I became involved in evaluation and monitoring of ESFRI projects and proposals.

This episode has taught me much about the interplay between generic (horizontal) and disciplinary (vertical) infrastructures, both on the national and the European level.

And now we have the European Open Science Cloud.

In its Staff Working Document (SWD), the European Commission has defined the EOSC as a virtual environment to offer free at the point of use, open and seamless services for storage, management, analysis and re-use of research data, across borders and scientific disciplines.

The challenges for the EOSC reside on the interface between vertical and horizontal e-infrastructures. The latter have the potential of being efficient and effective, pooling hardware, software and, even more importantly, people and expertise together instead of building disciplinary pillars.

EGI has been experiencing how integration of computing with data and discipline-specific analytics services, is key for the successful adoption by users.

The current use by research communities in physics, earth sciences, neuro-informatics, chemistry, biology and many more, are proof of this success.

According to the SWD, the EOSC will most likely be implemented as a federation of infrastructures. In my view, the generic part of the EOSC will be the federation of national e-infrastructures.

e-IRG has been advocating governments and funding agencies to improve e-infrastructure coordination at the national level. These strong national e-infrastructure building blocks will enable frictionless cross-border research collaborations and contribute to European e-infrastructure landscape.

The EGI Federation has become a strong and visible part of this landscape. 23 NGIs and CERN are today contributing to build a federation spanning hundreds of datacenters, operating as a remarkably smooth whole, to serve research communities in Europe. A federation par excellence!

I expect that the EOSC will be the incentive for further consolidation in the European e-infrastructure ecosystem, extending well into its national foundations. This consolidation has also been mentioned by Andreas Veispak (DG CONNECT), in his talk at the recent EOSC-hub week in Prague.

Challenging times lie ahead to explore with our friends in this ecosystem, EUDAT, GÉANT, OpenAIRE, PRACE, and our collective constituency, the best possible way forward to make and keep our researchers happy!

Arjen van Rijn is the Institute Manager of Nikhef and chairs the EGI Council since January 2019.

Building the grid: pretending it was easy

Sara Coelho and Jeff Templon look back at 15 years of milestones

The distributed computing infrastructure we now call EGI is a success. The installed High-Throughput Compute capacity has surpassed one million cores earlier this year, and as of March we have about 690 petabytes available for disk and tape storage. Our federated e-infrastructure processes more than half a billion compute jobs per year and serves research communities across all fields of science.

The scale and range of the achievements are impressive but, like Rome, none of it was built in a day. Taking the EGI federation off the ground involved groundbreaking work and dozens of teams across Europe.

So what were the major steps in our story?

The testbed

The European DataGrid project started with the new century on January 1, 2001 with the aim to develop and implement a computational grid for data-intensive scientific computing models, to be demonstrated on a worldwide testbed.

The testbed was first presented at a project review held on March 1, 2002. During the live demonstration, a grand total of 15 jobs were submitted by scientists from the LHCb experiment, Earth Observation and Computational Biology research communities. The 15 jobs were distributed across the five computer centers: CERN, CNAF (Bologna), CC-IN2P3 (Lyon), RAL (Didcot), and NIKHEF (Amsterdam).

The overall experiment was a success. The jobs consumed in total less than one CPU-hour, but the demo showed that the grid was capable to parallelise computational tasks.

The distributed e-infrastructure was born.

Scaling up from prototype

This first release of the testbed was a successful proof of concept. The e-infrastructure needed to be scaled up by several orders of magnitude by the time the LHC started operation. The DataGrid project turned its attention to this challenging task.

As the testbed increased in size (more sites and more machines), problems began to appear with various components on which it was built. An example was the Information System - the list of all services available and information on how to use them. The original Globus component (the Grid Information Index Service or GIIS) became unusable once the grid grew past 7 sites.

During a lunch discussion at Nikhef one day, somebody asked the question "how hard could it be?" (to make a stable information system). David Groep answered that the next day with the first release of the "fake II" (later renamed the BDII). After a few years of use and a factor-of-ten increase in the scale of the grid, the BDII was re-engineered (at CERN) in the EGEE and LCG projects and is still in use today.

The Compute Element was another example: the original could not scale past 100 running jobs per site and had to be redeveloped. It was replaced first with the LCG-CE, later with the CREAM-CE which is itself about to be replaced as of this writing.

The Service Challenges

As the e-infrastructure grew, the LHC Computing Grid (LCG) project was concerned with making sure that the grid would cope with the volumes of data expected with the start of LHC operations.

To do this the LCG designed a series of Service Challenges:

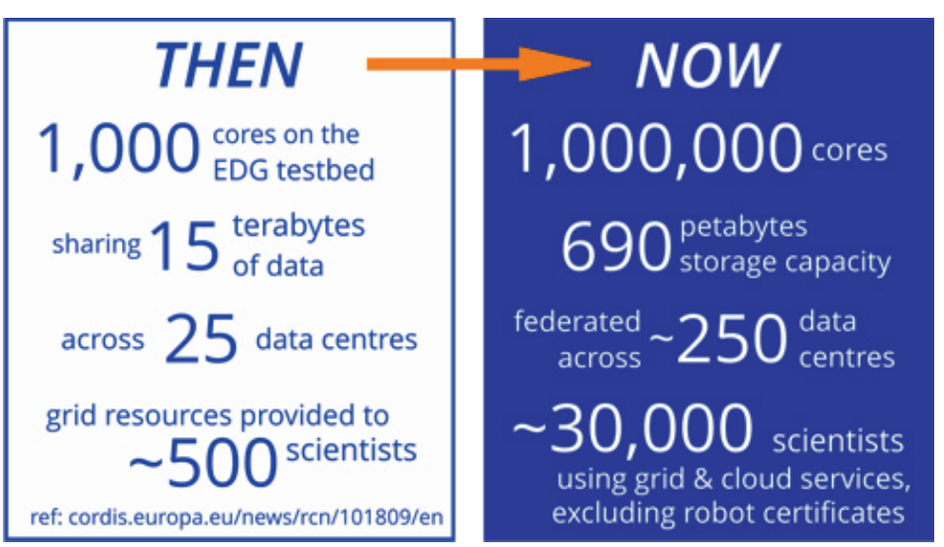

THE MAIN GRID PROJECTS

European DataGrid - EDG
January 2001 – December 2004

EDG built the first test computing infrastructure capable of distributing jobs and sharing data and computing resources across sites.

EGI-Design Study (EGI_DS)
September 2007 – December 2010

EGI_DS established the framework for a permanent organisation to oversee the grid at an European level – this is now the EGI Foundation.

Enabling Grids for E-Science in Europe - EGEE
EGEE I, II and III: April 2004 – April 2010

The EGEE projects worked with LCG to expand the EDG proof of concept. During EGEE, the grid moved from testbed to production environment and saw the development of mature monitoring, accounting, support and other systems required to run the grid.

EGI-InSPIRE
May 2010 – Dec 2014

Supported the operational development of the EGI federation and deployed the EGI Federated Cloud in May 2014.

EGI-Engage
Mar 2015 – Aug 2017

Expanded the backbone of federated services for compute, storage and data in Europe, setting the scene for the EOSC.

scenarios to test real production capabilities in preparation for the start of LCG operations in April 2007. The Service Challenges started in 2004 focusing the development work on specific goals. For example:

> Demonstrating sustained data transfers between grid sites, and

> Providing a reliable base service from all the Tier1s

The Service Challenges progressed with increasing complexity (not unlike leveling up on a computer game) and showed that coordination was key for everything to work smoothly.

A production infrastructure at your service

Through the LCG Project, the EGEE series and now the EGI era, the grid continued to grow in almost every dimension: number of sites, countries, users, cores, petabytes, communities. For the entire e-infrastructure to work reliably and professionally, we needed to develop the support tools necessary for smooth operations. These included, for example, the GGUS helpdesk, configuration management or the monitoring system; here we will focus on accounting.

The original accounting "system" consisted of sysadmins manually sending the numbers recorded on their site by email. Such a procedure is not scalable, desirable or useful in the long term. A central database was developed along with a messaging service; sites could publish their numbers into the messaging service, to be collected into the database. Many sites made use of the APEL tool, developed by STFC, to generate records automatically and publish them; other sites published records directly from their own site-level repository. The central database is then used to generate the summaries available through the Accounting Portal. The jobs recorded in the central accounting database go back to January 2004: about 4,000 jobs submitted by NIKHEF.

The start of the accounting records 15 years ago is a major milestone as in establishing the grid as a mature, professional system, where sites could demonstrate precisely how much resource was used at their site by whom, which is in turn crucial to showing funding agencies the justitfication of their investments.

15 years later...

Now that we are in 2019 with an e-infrastructure in production, supporting wonderful scientific discoveries, it is easy to be amazed at the technical achievement that was to put this together.

But we would argue that it's not just that. Building the grid was also a great sociological experiment: getting so many different teams to sit together to work on one common goal was perhaps our greatest accomplishment.

Sara Coelho is the EGI Foundation Communications Manager and did not know what middleware was until 2010

Jeff Templon leads the Physics Data Processing programme at Nikhef and has worked on grid computing since July 2001

EGI in numbers: from small steps to giant leaps

Alessandro Paolini and Vincenzo Spinoso review the figures from the last years

Over the last 15 years, EGI has been supporting science at all scales, from individual researchers to research infrastructures and large collaborations, by offering computing resources throughout Europe and worldwide.

We have been monitoring the evolution of the EGI e-Infrastructure since the start of the EGI-InSPIRE project, through the period of EGI-Engage and now in EOSC-hub times.

This article highlights a couple of the trends we have recorded.

Of particular note is the 1,000,000 federated cores milestone which was reached earlier this year, marking a decade of cointinued growth.

We are now looking forward to the 2 million core mark!

Installed compute capacity (number of cores), 2011-2019

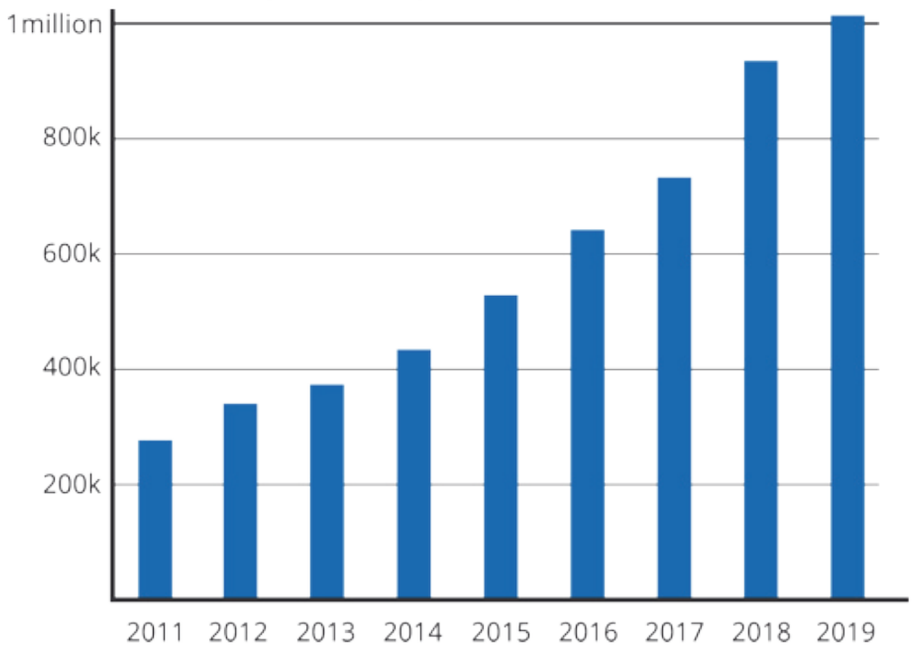

The number of cores installed on the data centres of the EGI Federation has consistently increased since the current measurement system was implemented in 2011, during the EGI-InSPIRE project.

The federation crossed the **1,000,000 cores** milestone earlier in 2019.

Storage capacity (Petabytes), 2012-2019

The EGI Federation offers online storage (in disk) and archive storage (on tape) to meet requirements for data sharing.

The available capacity has been increasing steadily and today is **696 Petabytes**.

The capacity for archive storage decreased slightly in 2019 due to the decommissioning of the service on a specific data centre.

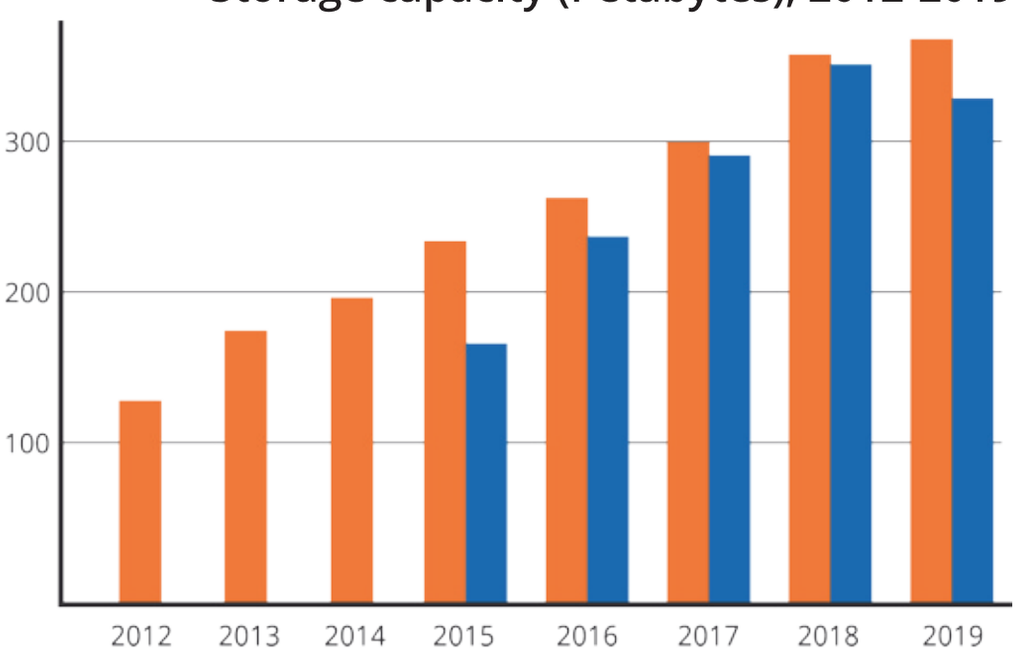

User distribution per field (log scale), 2011-2019

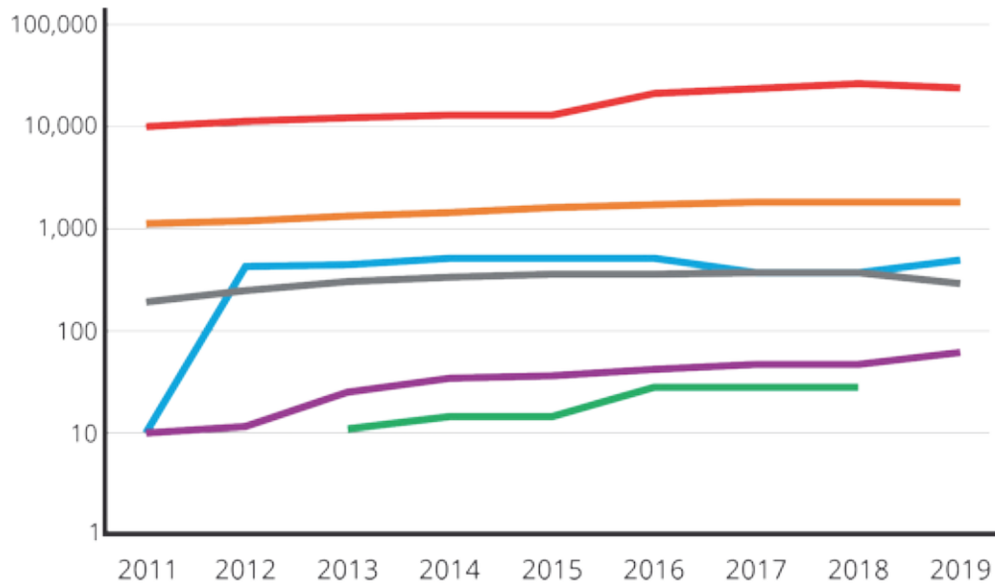

The number of users of the EGI Federation has been increasing steadily over the years.
The highest increase was recorded in the Humanities field.

Consumed HTC resources (elapsed time in years x cores), 2011-2018

This graph shows the years of not normalised wall-clock time consumed by jobs, multiplied by the number of logical cores used by the job.This gives a measure of how much resources are consumed.

In total, since 2011, EGI users have consumed **2,3 million years of CPU time**. With one single core laptop, you would need to go back to the Pliocene, when mammoths and saber-tooth cats ruled the earth.

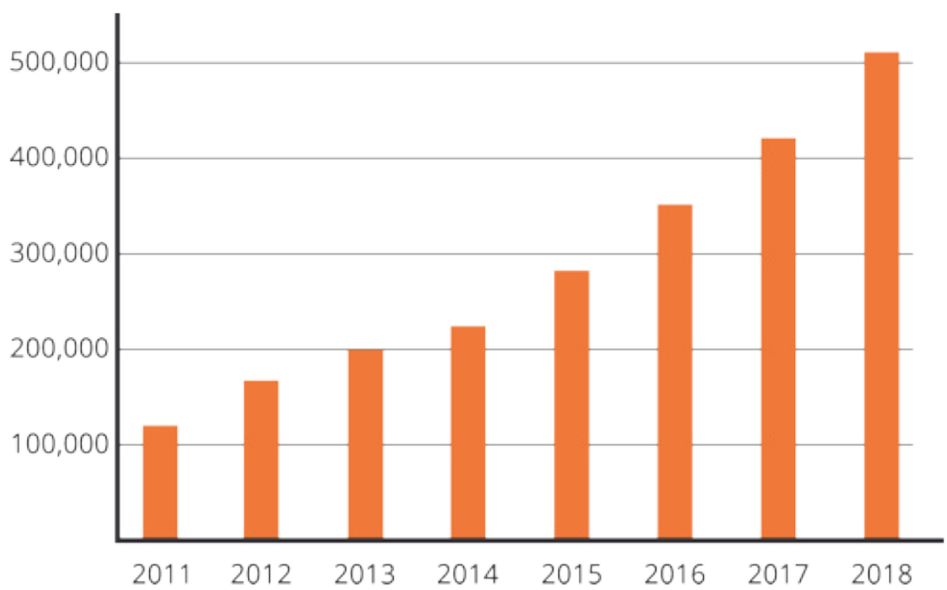

Cloud statistics, 2015-2018

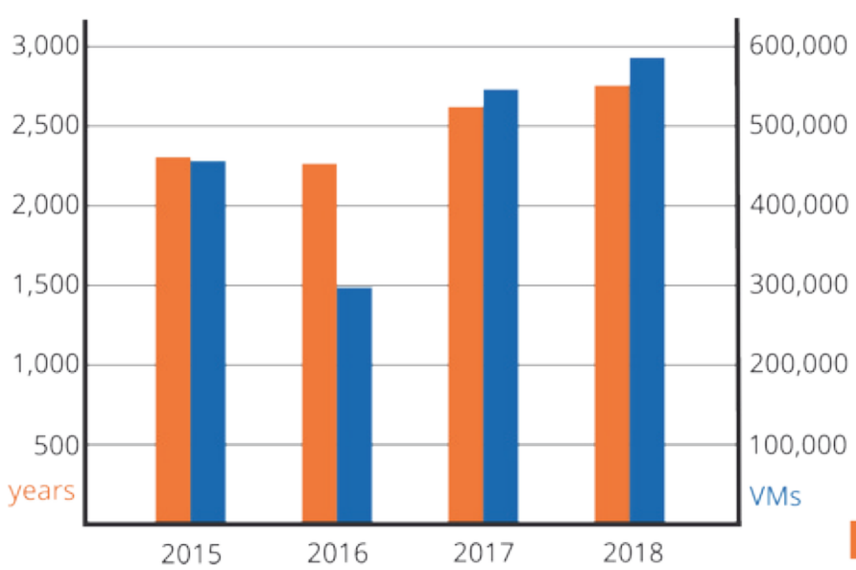

The EGI Federated Cloud was launched in 2014 and went into production in 2015.

Since then, cloud users have ran **1.9 million Virtual Machines** (VMs) and consumed a over **10,000 years** of compute time.

The decrease in 2016 is due to an accounting problem that was corrected in 2017.

The EMSO Data Management Platform: from prototype to full production

Björn Backeberg writes about how EGI is supporting EMSO ERIC

The ocean plays an integral role in regulating Earth's climate and weather patterns, including the heat, freshwater and carbon cycles but despite the ocean's recognised impact, it is still poorly understood.

The European Multidisciplinary Seafloor and water column Observatory (EMSO) aims to explore the oceans and to explain their role in the broader Earth systems, focussing on climate change, risks for biodiversity and natural hazards. EMSO's observatories are platforms equipped with multiple sensors to measure chemical and physical parameters, for example ocean temperature or dissolved oxygen concentration.

EMSO aims to offer distributed data and services to its community, and the EGI Foundation assisted them in this challenge during the EMSODEV project (EU Grant No: 676555).

EMSODEV was set up to develop and deploy the EMSO Generic Instrument Module (EGIM) as a fully operational distributed Research Infrastructure. A key component of providing accurate measurements of ocean parameters, is the Data Management Platform (DMP). The prototype DMP ingests, consolidates, processes and archives data from EGIM, integrates the data management architectures of the regionally distributed EMSO nodes and makes data available to the community.

In addition to a data portal, which provides early access to quality controlled EMSO data, the DMP provides scientific user with the following a set of tools:

> EMSODEV API allows scientific users and other European initiatives in the ocean sciences to interface with DMP data.

> MOODA (Module for Ocean Observatory Data Analysis) is a python framework for direct data access, with data analysis methods.

During the EMSODEV project, EGI assisted EMSO in implementing a DMP prototype using a subset of the EMSO data. To support this implementation, the EGI Foundation brokered access to cloud computing resources made available by RECAS-BARI, NCG-INGRID-PT, INFN-PADOVA-STACK and CESGA. A total of 340 vCPU cores and 9TB storage were made available.

EMSO is now planning to transition the prototype to production by the end of 2019. The operational system needs to be running on a robust infrastructure, with well curated data and no data loss, providing consistent data delivery to the user community. This will require pledged allocation, where resources are reserved and the job will be executed right after submission.

Together with EMSO, the EGI Foundation is scoping the requirements needed for EMSO to run its DMP in full production mode. The fully operational system will provide accurate, long-term measurements of ocean parameters. This, in turn, will lead to increased interoperability of EMSO nodes and the consistent collection of ocean essential variables.

Björn Backeberg is part of the EGI Foundation User Support team.

EMSO-ERIC *http://emso.eu/*

EMSODEV project: *http://www.emsodev.eu/*

EMSODEV team:
P. Andriani (Engineering), L. Badiali (INGV), R. Bardaji (CSIC), R. Huber (U. Bremen), N. Lanteri (Ifremer), J. Piera (CSIC), E. Martinez (UPC), J. del Rio (UPC), H. Ruhl (NOC), J.J. Dañobeitia (EMSO ERIC), P. Favali (INGV)

VIP and the biomed community: achievements and new challenges

Sorina Camarasu Pop writes about a successful science gateway

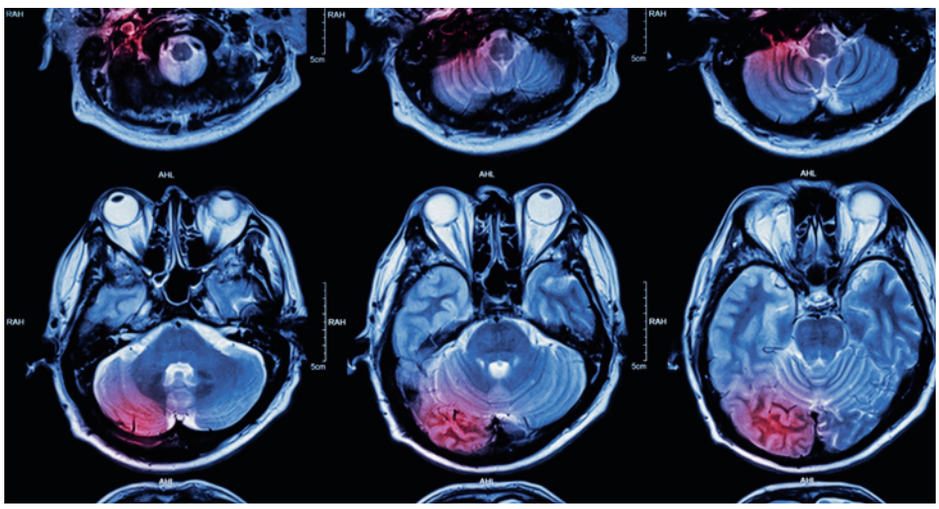

biomed is a large-scale international Virtual Organization (VO) supporting communities from the Life Sciences sector focusing on three main groups:

> medical image analysis,

> bioinformatics and

> drug discovery.

The biomed VO is operated on the EGI infrastructure and supported by more than 50 sites, delivering access to a large number of heterogeneous resources. In 2018, biomed users executed more than 2.5 million jobs corresponding to 439 years of CPU.

One of the flagship science gateways of biomed is the **Virtual Imaging Platform** (VIP). Through VIP, researchers can access multiple applications, as well as important amounts of storage and computing resources, without required technical skills beyond the use of a web browser.

VIP relies on the France Grilles DIRAC instance for job submission and data management on EGI biomed resources. The model has proved to be successful: VIP currently counts more than 1000 registered users, accessing 20 applications. Since 2011 we have logged dozens of peer-reviewed papers published with results computed through VIP.

Challenges for the future

VIP is now focusing on challenges concerning interoperability and reproducibility, in the larger scope of a FAIR approach to scientific data analysis. Through **Boutiques**, VIP is able to easily describe and integrate new applications, as well as publish them on open repositories, such as Zenodo, to make them findable and accessible.

Through **CARMIN** (Common API for Research Medical Imaging Network), VIP fosters data integration and interoperability among platforms. In the next years, we can imagine a platform of platforms exchanging apps and data effortlessly.

Re-usability of applications across platforms can be achieved with containers, which allow users to build applications on their desktops and run hundreds or even thousands of instances on remote CPUs.

The biomed community has come a long way along with what is now the EGI Federation. Despite the multitude and diversity of our research topics and applications, we share similar challenges and we can join our efforts in order to reach goals surpassing our individual means. We need to work together to build a FAIR and promising future.

VIP in parctice: creating a digital heart

Scientists from the University of Lyon used VIP to develop a framework to generate virtual data to study heart diseases. One 3D sequence took around 6 hours to generate on VIP; on a personal laptop, 280 hours would be needed. doi: 10.1109/TMI.2017.2708159

Sorina Camarasu Pop is a CNRS research engineer working at the CREATIS laboratory in France.

VIP Platform
https://vip.creatis.insa-lyon.fr/

What we learned from the EOSCpilot Science Demonstrators

Giuseppe La Rocca summarizes the recommendations for EOSC

The EOSCpilot project was set up to support the first phase in the development of the European Open Science Cloud. As part of this ambition, the project selected 15 Science Demonstrators from different domains to pilot actual implementations of the EOSC Service Portfolio.

The demonstrators were chosen to provide insight on technical and policy needs and prioritise the integration of the EOSC services to meet requirements.

The Science Demonstrator activity involved many experts from 17 institutions working together to analyse requirements, link research communities to compute and storage providers and to manage the technical infrastructures needed for co-design and piloting.

The activity brought many benefits for the communities involved. For example: the EPOS-VERCE and ViSIVO Science Demonstrators shared best practices to connect their frameworks with the EGI Federated Cloud Infrastructure and allow users to run scientific workflows on the cloud computing infrastructure using their federated credentials. Dedicated Virtual Research Environments (VREs) were set up to help members of the Social Science Communities to share and visualize media files on the web and implement the semantic enrichment of text sources.

The EOSC cloud infrastructure was adopted to scale up the execution of scientific workflows and pipelines. In particular, it contributed to facilitate the set up of a cloud-based workflow system to execute genomics analysis that, ultimately, can contribute to improve patients health care, and enable users of the fusion community to reproduce science efficiently.

Many Science Demonstrators focused on the implementation of the FAIR principles.

> CryoEM extended the Scipion framework to support reproducibility by sharing of detailed information on cryo-electronic microscopy image processing workflows.

> EGA Life Science Datasets managed to reproduce and re-master biological pipelines combining Docker container solutions with Nextflow, an emerging language aimed to ease the interpretation of scientific workflows.

Recommendations for EOSC

The impact of the Science Demonstrators goes beyond the positive effect they had on their communities of practice: what we, the data, services and e-Infrastructure providers, learned from the activity is of great value to the implementation efforts of the EOSC.

And from this work we can list the following reccomendations:

> Implement a distributed data and compute infrastructure providing high performance access to data transfer, mirroring and caching, supported by high speed network connectivity.

> Procure and offer EOSC as a federated infrastructure that integrates existing community resources (data, applications, software, storage & computing) and provides additional adequate capacity to scale up existing in-house IT infrastructures. Procure EOSC as a high-capacity system

Science Demonstrators moving to pre-production

Towards the end of the project the following demonstrators were selected to continue their work-plans for three additional months in order to move their pilot services into a pre-production phase:

> **TextCrowd (Social Sciences)**: a text mining solution to semantically enrich text sources and make them available on the EOSC.

> **Photon & Neutron (Physics):** aims to create a virtual platform where data and analysis tools can be made available to scientists all over the world.

> **Prominence (Energy Research)**: provides access to HPC class nodes for the Fusion Research community through a cloud interface.

> **LOFAR (Astronomy):** provides easy access to LOFAR data and knowledge extraction through Open Science Cloud.

> **VisualMedia (Social Sciences & Humanities)**: a service for sharing and visualizing media files on the web.

that meets the demands of data intensive science.

> Provide easy-to-use environments such as scientific gateways, Virtual Research Environments, as managed services to provide integration as turn-key solution; make service descriptions discoverable; offer ready to use integrated bundles of services with low-barrier procurement processes.

> Provide a federated Authentication & Authorisation solution to allow users to access services and resources from different providers with the same credentials.

> Provide and sustain human networks through competency centres of experts working with scientific application developers in close cooperation.

> Provide support for running standard-based workflows.

> Promote tight integration between services and providers.

> Extend the FAIR concepts currently applied to data to IT services. Propose a set of recommendations for making services FAIR, or to further enable services to make data FAIR.

> Include analysis of network requirements, specifically when designing the interoperability of services across sites and organisations.

Although the work on EOSC is still very much a work in progress, we can already see some of these recommendations coming into play. For example: federated log-ins are available on the EOSC Marketplace and many pilot services developed during the EOSCpilot project have enabled federated authentication access for their end-users and have started to offer EOSC services to their scientific communities. As result of the EOSCpilot project, many pilots started the registration in the EOSC Marketplace as service provider.

We hope to see more following soon.

Giuseppe La Rocca is part of the EGI Foundation User Support team and led the EOSCpilot Science Demonstrators activity.

EOSCpilot project:
https://eoscpilot.eu/

The Science Demonstrators
https://eoscpilot.eu/science-demonstrator-topics

Service Level Agreements: EGI's mission into practice

Iulia Popescu on how EGI bridges services from providers to users

EGI has a mission to deliver open solutions for scientists, research infrastructures and industry by federating digital resources across communities and national boundaries.

One of the biggest challenges in achieving this mission is to facilitate the easiest way for researchers to find and use our services. This is why EGI developed **Service Level Agreements (SLAs)** – a trust-based communication channel between researchers and providers to agree on services, their levels and types of support offered. SLAs are not legal contracts but as agreements they outline clear intentions to collaborate and support research.

Over the years, EGI has established numerous SLAs with a variety of scientific communities, from Physics and Astronomy to Environmental Sciences, Life Sciences and Humanities.

Here are a few examples...

Peachnote - Arts and Humanities
Peachnote is a music score search engine and analysis platform.
Services: Cloud Compute, Online Storage
Providers: CESNET MetaCloud

Terradue - Environmental Sciences
Terradue is an SME with a mission to innovate services in Earth science, tailored for data-intensive applications.
Services: Cloud Compute, Online Storage
Providers: 100 Percent IT, CYFRONET-CLOUD, CESGA, GWDG, RECAS-BARI, BEgrid-BELNET

EXTraS - Astronomy
EXTraS has a mission to harvest the unexplored information buried in the data collected by ESA's X-ray space observatory XMM-Newton, during 16 years of observation.
Services: Cloud Compute, Online Storage
Providers: CYFRONET-CLOUD and INFN-CATANIA-STACK.

MoBrain - Life Sciences
MoBrain provides biomedical scientists with web portals to study molecular interactions that can improve treatments and drug design.
Services: Cloud Compute, Online Storage, High-Throughput Compute
Providers: CESNET-MetaCloud, INFN-PADOVA, INFN-PADOVA-STACK, NCG-INGRID-LP, RAL-LCG2, TW-NCHC, NIKHEF and SARA-MATRIX

National Bioinformatics Infrastructure Sweden - Life Sciences
NBIS is a national bioinformatics infrastructure supporting life sciences in Sweden
Providers: IN2P3-IRES and RECAS-BARI
Services: Cloud Compute, Online Storage

D4Science - Multi-disciplinary
D4Science hosts more than 95 Virtual Research Environments (VREs) to serve the biological, environmental and statistical communities worldwide.
Services: Cloud Compute, Online Storage
Providers: CESGA , INFN-CATANIA STACK

Iulia Popescu is the EGI Communications Officer and lead editor of Inspired

EGI Operations Team: 2019 goals

Matthew Viljoen on the team's upcoming objective

The EGI Central Operations Team has defined its high-level objectives for 2019. These objectives are designed to ensure that high-quality services continue to be delivered to the different user communities and to ensure that we continue to improve our service delivery. In this article, we provide an overview of the ways we are doing this by taking a look at some of the high-level objectives for this year.

This year, we have introduced a new way of ensuring that EGI Operational Tools (e.g. monitoring, accounting, VMops Dashboard) meet user requirements. This is being done by the tool developers in cooperation with the EGI Operations Management Board where the National Grid Initiatives (NGIs) are represented. By providing a forum where development roadmaps and priorities are discussed, the EGI Operations Team will help to ensure that these tools remain fit for purpose for the EGI community into the future.

Continuity of service remains a top priority for the EGI Operations Team, especially when software reaches its end of life. In the light of the recent announcement that the CREAM-CE service will cease to be maintained at the end of the EOSC-hub project, the EGI Operations Team in coordination with the CERN WLCG Operations Team are working hard to ensure that user communities still using CREAM-CE can migrate to alternatives as easily as possible. ARC-CE is already well supported within EGI and during 2019 we are working to ensure that HTCondorCE will be equally well supported.

In any complex e-Infrastructure, monitoring when things go wrong is essential.

An effective testing framework typically needs to be constantly tweaked and maintained to ensure that the right people are alerted when failures happen However, it is often challenging to understand abnormal patterns of behaviour when different user communities are using the e-Infrastructure in many different ways.

The easiest way to deal with this is by using common tests which monitor the basic functionality of services, an already existing activity within the EGI Federation (e.g. there are over a hundred different tests running in ARGO). Nevertheless, a smarter way of monitoring is by better understanding the specific work of the different communities and introducing new tests to alert when things go wrong. This is why the EGI Operations Team is now spending time introducing specialised monitoring.

Internal Services delivered by EGI have Operations Level Agreements (OLAs) are in place between EGI Foundation and service providers. This enables EGI to broker customers' requirements with services being delivered as part of the Federation. The EGI Central Operations Team is reviewing OLAs to ensure that they are aligned with user expectations and will be looking for areas that can be improved. The team is also reviewing procedures to ensure that any problems of operational level targets are adequately dealt with.

Matthew Viljoen is the EGI Operations Manager.

EGI Internal Services for federation providers: *egi.eu/internal-services/*

And a final announcement...

...EGI will soon be looking for its new director

After leading the EGI Foundation for over five years, Yannick Legré has announced his intention to leave by the end of this year.

Yannick has made the EGI Federation into the smoothly operating e-infrastructure it is today, serving thousands of researchers in Europe and worldwide. He has turned the Amsterdam-based EGI Foundation into a stable organization, employing over 20 people.

"There is a time for everything. There has been a time for building EGI and a time for structuring it. Now, it is time to step back and allow a new director to move EGI forward into the challenging future of the European Open Science Cloud and the consolidation it will require", says Yannick.

"We are very grateful for the excellent work that Yannick has given to the organisation. He has made EGI ready to enter the new context of the EOSC, which we are convinced will lead to even more collaboration and consolidation between European e-Infrastructure organisations", says Arjen van Rijn, chairman of the Executive Board.

The EGI Foundation will soon advertise the director position allowing the new director to overlap with Yannick's term so that they can profit from his vast experience.

Inspired, the EGI newsletter, was first published in September 2010, right at the start of the EGI-InSPIRE project (that explains the name!). The newsletter is edited by the EGI Communications Team (Sara Coelho and Iulia Popescu), and builds on contributions from our colleagues at the EGI Foundation and across the EGI Community.
Inspired does not get printed often, so this is a good opportunity to thank the 117 writers that have contributed to our first 34 issues.

Thank you!

Sergio Andreozzi - Jelena Angelis - Owen Appleton - Carmela Asero - Themis Athanassiadou - Björn Backeberg - Daniele Bailo - Antun Balaz - Roberto Barbera - Boris Bauermeister - Ugo Becciani - Ian Bird - Ignacio Blanquer - Alexandre Bonvin - Gonçalo Borges - Jan Bot - Steve Brewer - Magdalena Brus - Stephen Burke - Sorina Camarasu Pop - Isabel Campos - Francisco Castejón - Elisa Cauhé - Daniele Cesini - Marios Chatziangelou - Yin Chen - Daniele D'Agostino - Mário David - Giacinto Donvito - Andreas Drakos - Michel Drescher - Afonso Duarte - Laura Durnford - Karolis Eigelis - Enol Fernandez - Tiziana Ferrari - Nuno Ferreira - Licia Florio - Antonella Fresa - Sven Gabriel - Sara Garavelli - Catherine Gater - Audrey Gerber - Tristan Glatard - Jorge Gomes - Antonio Gómez - Pedro Gonçalves - Hilary Goodson - Ricardo Graciani - Baptiste Grenier - David Groep - Ingemar Häggström - Akos Hajnal - Simon Hettrick - Sy Holsinger - Anni Jakobsson - Javier Jimenez - Peter Kacsuk - Namrata Kale - Christos Kanellopoulos - William Karageorgos - Eleni Katragkou - Dimitra Kotsokali - Kostas Koumantaros - Małgorzata Krakowian - Cyril L'Orphelin - Giuseppe La Rocca - Anders Larsson - Yannick Legré - Viviane Li - Simon Lin - Robert Lovas - Nikos Manouselis - Jesús Marco - Damir Marinovic - Ludek Matyska - Richard McLennan - Franck Michel - Chris Morris - Stefano Nativi - Steven Newhouse - Alex Nicolin - Leif Nixon - Neasan O'Neill - Pasquale Pagano - Alessandro Paolini - Claudia Pierotti - Salvatore Pinto - Roberta Piscitelli - Claudio Prandoni - Michal Prochazka - Fotis Psomopoulos - Geneviève Romier - Barbara Ryan - Davide Salomoni - Diego Scardaci - Michael Schiffers - Gergely Sipos - Annabel Slater - Peter Solagna - Carrie Solomon - Ola Spjuth - Erika Swiderski - Agnes Szeberényi - Jelena Tamuliene - Jeff Templon - Paul van Dijk - Arjen van Rijn - Rosette Vandenbroucke - Kitti Varga - Marco Verlato - Matthew Viljoen - John Walsh - Eric Yen - Panagiotis Zervas - Zhiming Zhao