



EOOSC-hub

D4.4 Capacity plans for services in the Hub Portfolio

Lead Partner:	EGI.eu
Version:	1
Status:	FINAL
Dissemination Level:	Public
Document Link:	https://documents.egi.eu/document/3628

Deliverable Abstract

This document describes the criteria that we adopted to produce the capacity plans for the services in the Hub Portfolio, and presents the capacity plans created so far, providing some recommendations about future capacity requirements for the involved services.

COPYRIGHT NOTICE

This work by Parties of the EOSC-hub Consortium is licensed under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>). The EOSC-hub project is co-funded by the European Union Horizon 2020 programme under grant number 777536.

DELIVERY SLIP

<i>Date</i>	<i>Name</i>	<i>Partner/Activity</i>	<i>Date</i>
From:	Alessandro Paolini Catalin Condurache Isabella Bierenbaum	EGI Foundation EGI Foundation KIT	29/07/2020
Moderated by:	Małgorzata Krakowian	EGI Foundation	
Reviewed by:	Pavel Weber Matthew Viljoen	KIT EGI Foundation	
Approved by:	AMB		

DOCUMENT LOG

<i>Issue</i>	<i>Date</i>	<i>Comment</i>	<i>Author</i>
v. 0.1	24/06/2020	version ready for the review	A.Paolini, C. Condurache, I.Bierenbaum
v. 1	29/07/2020	Final after external review	A.Paolini, C. Condurache, I.Bierenbaum

TERMINOLOGY

<https://wiki.eosc-hub.eu/display/EOSC/EOSC-hub+Glossary>

<i>Terminology/Acronym</i>	<i>Definition</i>
AAI	Authorization and Authentication Infrastructure
CAPM	Capacity Management
CMDB	Configuration Management Database
DPMT	Data Project Management Tool
EOSC	European Open Science Cloud
FitSM	Free standards for lightweight IT Service Management
GGUS	Global Grid User Support
GOcdb	Grid Operations Configuration Management Database
ITMS	IT Management System

OLA	Operational Level Agreement
SDTP	Service Design and Transition Packages
SLA	Service Level Agreement
SLI	Service Level Indicator
SLM	Service Level Management
SMS	Service Management System
SP	Service Provider
SPM	Service Portfolio Management
SVB	EOSC-hub Service Validation Board
UA	Underpinning Agreement
VO	Virtual Organisation
VOMS	Virtual Organisation Management System

Contents

1	Introduction	7
2	Capacity Management (CAPM) process.....	8
2.1	Scope of the process	8
3	Capacity plan structure and criteria adopted	9
4	Capacity plans of the services in the Hub Portfolio	12
4.1	AAI - EGI Check-in.....	12
4.1.1	Software/Service Performance SLIs	12
4.1.2	Human (Operation and Tickets handling) SLIs	15
4.2	Configuration Management Database - EGI GOCDB	16
4.2.1	Software/Service Performance SLIs	16
4.2.2	Human (Operation and Tickets handling) SLIs	16
4.3	Accounting Repository	17
4.3.1	Software/Service Performance SLIs	17
4.3.2	Human (Operation and Tickets handling) SLIs	18
4.4	Accounting Portal.....	19
4.4.1	Software/Service Performance SLIs	19
4.4.2	Human (Operation and Tickets handling) SLIs	19
4.5	Collaboration Software	20
4.5.1	Software/Service Performance SLIs	20
4.5.2	Human (Operation and Tickets handling) SLIs	21
4.6	EOSC-hub Helpdesk.....	22
4.6.1	Software/Service Performance SLIs	23
4.6.2	Human (Operation and Tickets handling) SLIs	23
4.7	EOSC Marketplace.....	24
4.7.1	Software/Service Performance SLIs	24
4.7.2	Human (Operation and Tickets handling) SLIs	25
4.8	EOSC-hub Monitoring	26
4.8.1	Software/Service Performance SLIs	26
4.8.2	Human (Operation and Tickets handling) SLIs	27
4.9	EOSC Portal	28
4.9.1	Software/Service Performance SLIs	28

4.9.2	Human (Operation and Tickets handling) SLIs	30
4.10	Operations Portal	31
4.10.1	Software/Service Performance SLIs	31
4.10.2	Human (Operation and Tickets handling) SLIs	33
5	Conclusion	34

Executive summary

Capacity planning is an important aspect that extends across the service lifecycle: a key success factor in managing capacity is ensuring it is considered during the design stage, with periodic reiterations during delivery of the service, in order to guarantee that the capacity of the services meets the agreed capacity- and performance-related requirements in a cost-effective and timely manner. Capacity management is concerned with meeting both the current and future capacity and performance needs of the business.

The EOSC-hub IT Management System¹ is based on the FitSM standard² so in carrying out this activity we applied the FitSM Capacity Management requirements, as also explained in a previous deliverable³. We have worked on the capacity plans of the services included in the Hub Portfolio, and we present in this document the criteria adopted to produce a capacity plan and the plans that have been completed so far. Our approach for the creation of a capacity plan for a given service is in general to define the most appropriate set of quantitative parameters to measure the capacity consumption, load, and performance of the resources in order to understand if the available supply of the service processing power matches the demands made on it by the business. With the analysis of these capacity monitoring data we also produce forecasts for future capacity requirements to meet future demand of a given service.

Our conclusion is that the current capacity of the several services is sufficient to match the current service demand and to continue to provide the services of the quality defined in the specific OLAs; we don't foresee a great increase in the future service demand, but even if it will grow, the Hub Portfolio services can cope with it maintaining the same quality level than before.

¹ An IT service management system is the entirety of activities performed by service providers to plan, deliver, operate and control services offered to customers

² www.fitsm.eu

³ D4.2 Operational Infrastructure Roadmap: <https://documents.egi.eu/document/3422>

1 Introduction

In this deliverable we present our approach to work on a capacity plan and the capacity plans of the Hub Portfolio services that we have produced according to the FitSM requirements on which the EOSC-hub Federated Management System is based.

After a brief description of the Capacity Management process in section 2, we describe in section 3 the criteria that we agreed for creating a capacity plan, explaining its structure and the decisions that we made about the relevant information to include into it. Section 4 is dedicated to the capacity plans created so far, and for each service we provide an assessment of the current capacity and our recommendations for the next 6/12 months. In the last section we provide a summary of our conclusions about the capacity requirements of our services.

2 Capacity Management (CAPM) process

The goal of this process is to ensure that sufficient capacities are provided to meet agreed service levels and performance requirements for services that are part of the catalogue. Capacity Management (CAPM) is usually triggered before the release of a service into the production environment (during the creation of the SDTP), with a periodic reiteration during the entire lifetime of the services in the catalogue: the process considers all resources required to deliver the IT service, and plans for short-, medium-, and long-term business, capacity, and performance requirements.

The result of this analysis is the creation of a plan that documents the current level of resource utilisation and service performance and, after consideration of the service strategy and plans to forecast the future requirements for new IT resources, that supports the IT services that underpin the business activities. The plan clearly specifies any assumptions made as well as any recommendations quantified in terms of resources required, cost, benefits, impact, etc.

2.1 Scope of the process

The scope of our CAPM process is the Hub Portfolio, a collection of services supporting the delivery of the Hub as a whole. The definition of this group of services changed over the years and Table 1 contains the list of services under scope after setting-up the process and valid when we started to implement it. Very recently, the SVB started some discussions to modify the definition of the Hub Portfolio, proposing the inclusion of some new services like the Data Transfer, the Messaging and the Training services, and the removal of some others like DPMT and the Operations Portal⁴. When the new set of the Hub Portfolio services is approved, we will start to work on the capacity plans of the new services according to the requirements of our process.

Table 1 – List of services under scope so far.

Service name	Service component	Capacity Plan status	Included in the new “Hub Portfolio”
EOSC hub AAI	B2ACCESS	not yet started	No
	EGI ⁵ Check-in	completed	Yes
EOSC hub Accounting	EGI Accounting Portal	completed	Yes
	EGI Accounting Repository	not yet completed	Yes
EOSC hub CMDDB	EGI GOCDB	completed	Yes
	DPMT	not yet started	No
EOSC hub collaboration software	Confluence and Jira	not yet completed	Yes
EOSC hub helpdesk	EOSC hub xGUS	completed	Yes

⁴ A new component of the Operations Portal just released at the moment of writing this deliverable, SOMBO (Service Order Management Back Office) has been proposed as a service in the Hub Portfolio to support the service orders handled through the Marketplace.

⁵ An international e-Infrastructure, member of EOSC-hub, set up to provide advanced computing and data analytics services for research and innovation - <https://www.egi.eu>.

EOSC hub marketplace		completed	Yes
EOSC hub monitoring	EGI ARGO Monitoring service	not yet completed	Yes
EOSC Portal		completed	Yes
EOSC hub operations portal	EGI Operations Portal	completed	No

3 Capacity plan structure and criteria adopted

A capacity plan is used to manage the resources (human, technical, and financial) required to deliver IT services: it contains details of current and historic usage of IT services and components, and any issues that need to be addressed (including related improvement activities). The plan also contains scenarios for different predictions of business demand and cost options to deliver the agreed service level targets.

The capacity plan structure we designed is divided into several sections (Table 2), each one corresponding to some of the steps bringing to the creation of a capacity plan.

Table 2 – Schema of a capacity plan structure with the expected content.

Section	Content
Service Description	Reference to SDTP, OLA/UA, list of known issues related to service capacity
Service Level Indicators	List of defined SLIs along with a description
Capacity strategy	Strategies to adjust the service capacity to respond to changes in utilisation
Capacity monitoring	Data about service utilisation and performance covering up to the moment when the plan is created/reviewed
Forecast	Recommendations on capacity requirements for the next reporting period

First of all, the creation of a capacity plan is usually triggered by the addition of a new service in the portfolio (the Hub, in our case) or because there is a new or modified SLA for a service. The service business case design is reviewed, with a particular focus on the service demand, on the expected cost and expected revenue, and lastly on the Service Requirements and Service Acceptance Criteria. All this information is reported in the SDTP which can be updated if necessary. So, the first section of our

capacity plans, named Service Description, contains a reference to the SDTP, to any existing Operational Level Agreement, and also a list of known issues related to service capacity, e.g. reported incidents/problems, SLA violations attributed to capacity, customer complaints, etc.

Once completed, it is time to define quantitative parameters, the so-called Service Level Indicators (SLIs) relevant for understanding the capacity of the service. For all of our services, we tried to define a series of parameters easy to measure which can provide a direct indication on the load of the service under its normal operation, corresponding to the several capacity aspects of the given service.

A group of SLIs usually included in a capacity plan is the one reflecting the hardware aspect, related to technical resources (either virtual or bare-metal) where the service is hosted: CPU, disk, memory, and network. Measuring these will give a direct information about the load of the hosting machines.

Another SLIs category usually part of a capacity plan is related to the human and financial resources. Concerning our services, where the providers receive yearly contributions (PM, persons month effort) for specific activities to be accomplished in a given international project, we decided to consider as SLIs the “monthly effort” and the “number of staff members” that deliver the service; moreover, the providers are also required to provide support to the users (as stated in the services OLAs), in terms of addressing any incidents and service requests that are filed through the EOSC or the federated member Helpdesk services, so that we added to the list indicators such as the “number of tickets” and the “ticket average response time” (without making distinction between “incidents” and “service requests” and regardless the tickets priority). In this way we can have an idea of the human load of the service, for instance if a particularly high number of tickets produces an increase in the average response time, then it could be the case of increasing the number of staff members.

Besides that, we needed some indicators providing a direct information of the usage of the service itself intended as a software, what is consumed by the users, so we defined a series of parameters to measure the “software capacity performance” of the given service; as shown in the next section, this kind of SLIs are service specific.

The next part of the plan is the “Capacity Strategy” section which describes the strategies to adopt to modify the capacity of a service as a consequence of a change in demand. For each SLI it is reported:

- how it is monitored
- any identified capacity issue
- users’ experience in case of saturation (i.e. service completely interrupted, degraded, etc.)
- the designed capacity, if applicable
- capacity adjustment model (the best model to adopt to adjust the capacity).

Specifically, we defined four strategies to modify the capacity to respond to an increase of demand:

- Lead: adding capacity in anticipation of an increase in demand. The total capacity for the given service will be estimated and will be chosen to be always higher than the anticipated demand.
- Lag: waiting until the current capacity is stretched to its limits before adding more capacity.
- Incremental: adding capacity in small increments when you approach full capacity.
- Dynamic (Predictive): adding capacity, large or small, before it's required based on forecasts.

Given the context in which all these services (part of the Hub Portfolio) are currently being funded (through the EOSC-hub project), we considered a “Lead” capacity strategy type for the “Effort” SLI in the “Human” component of the Service Model, as this is something allocated at the start of the project to cover operations of the service for the entire duration and it is less likely to be adjusted later.

On the other hand, the “Number of staff” indicator was considered across all services as a “Lag” capacity strategy type, since additional individuals can be assigned temporary tasks to provide support to a service experiencing degraded performance at a specific moment.

The same considerations have been taken into account for a “Lag” strategy on “Software/Service performance” capacity area: from an operational perspective, effort to fix the service is allocated when the existing capacity is stretched i.e. service performance becomes poor and the service needs some adjustments at software level.

In the “Capacity Monitoring” section of the capacity plan we report the data collected for the several SLIs: looking at the data, we can analyse the utilisation trends of the several parameters to see how far we are from the capacity limit (if known). The trends analysis will then be used as input for the last part of the capacity plan, the one regarding the forecast service utilisation: based on historical data (and in case of new or changed service requirements) we make assumptions on the future demand of the service and on the needs of adjusting the capacity.

Regarding the “Capacity Monitoring” section, we took an approach valid for all our capacity plans: in general, the data related to the hardware SLIs are collectible through fabric monitoring, which is internal information directly available to the providers; we agreed that it is not necessary to provide these data to us, with the condition that we ensure that the providers are proactively monitoring them. So we asked the providers to notify us any time they detect an unusual hardware consumption and supply the related data of the affected SLI(s) in that particular period of time, in order to record this in the capacity plan, and document any decision taken to cope with the situation. In a similar way, any incident due to capacity issues will also be tracked. In some cases, moreover, there is not much sense in reporting data about the hardware capacity aspects, since the High Availability deployment of the given service can automatically detect any capacity overload and adjust the available resources accordingly.

Because of the specifics of each service under the Hub Portfolio, the corresponding Capacity plans have various complexities (i.e. more than one Service Provider in some cases, separate hardware configurations for multiple deployed sub-services), therefore the timeline of creation varied significantly. Hence while most of the plans have been created, approved and not yet due for their first review, there are plans that have already been revised and updated (as part of the CAPM process), but also a small number still in draft mode and not covered by this document.

As the regular review period is 12-month for each approved plan, not all capacity plans will be reviewed by the end of the EOSC-hub project. It is however foreseen that the CAPM process will be part of the Federated Management System within any continuation of the project, so capacity plans will continue to be reviewed and properly maintained.

We conclude this section saying that the general structure of the capacity plans evolved over the project years, especially because of the useful outcomes of the internal audits and process reviews occurred so far, and we will continue to update it to match new requirements or any time there is an opportunity for improvement.

4 Capacity plans of the services in the Hub Portfolio

In this section we will present a simplified version of the capacity plans of the services included in the Hub Portfolio (see Table 1). Based on explanations provided in the previous section and without going into too much detail, we will focus on the Software and Human SLIs identified for each service, reporting the capacity strategy chosen for each SLI, and providing also some usage data. We also report the plans that at the moment of writing this document have not yet been completed: even if we still miss the capacity data for some of the parameters, it was important to us reporting at least the SLIs we defined for the given service.

Each created capacity plan details information on Human SLIs as “Effort” (yearly PM allocated by the project - currently EOSC-hub) and “Number of staff” (staff members operating the service). However we are not reporting the respective figures in these cases due to project confidentiality aspects. Still, we would like to make the reader aware of the fact that the number of staff put in the Capacity Plans is not a full number of FTEs (Full-Time Equivalent), but a number of individuals operating the service (so it could be 0.4 FTE for a person and 0.8 FTE for another one). Also a service needs to be controlled by more than one person (be it 1.0 FTE or 0.2 FTE) otherwise degraded performance could happen if the operator is unavailable and no replacement is provided. Lastly, the personnel mentioned in the capacity plans are deployed not only for processing tickets and service requests, but also for service/software performance i.e. maintaining or developing the service.

Other general remarks:

- the capacity monitoring data where possible start from the beginning of the project or since when they were available to the provider
- we haven't identified any capacity issue concerning the several SLIs
- when the service capacity is saturated, the users' experience will be degraded, but the service will be still available
- for each service the capacity strategy model is a mix of “Lead” and “Lag”, in some cases also “Incremental” (limited to the hardware SLIs)
- by analysing the usage trends and by estimating the future service demand in the next 6/12 months, we do not foresee that an increase of capacity is required to our services: their consumption seems still quite far from the capacity limit.

4.1 AAI - EGI Check-in

EGI Check-in⁶ is a proxy service that operates as a hub to connect federated Identity Providers (IdPs) with EOSC-hub Service Providers. Check-in allows users to select their preferred IdP so that they can access and use EOSC-hub services in a uniform and easy way.

4.1.1 Software/Service Performance SLIs

- **Number of user registrations**
 - Number of new registered users (on a monthly basis)
 - Capacity strategy: lag

⁶ <https://www.egi.eu/internal-services/check-in>

- **Number of connected SPs**
 - Number of Service Providers connected to Check-in (on a monthly basis)
 - Capacity strategy: lag
- **Number of (not-eduGAIN) connected IdPs**
 - Number of Identity Providers (not in eduGAIN) connected to Check-in (on a monthly basis)
 - Capacity strategy: lag
- **Number of Communities (internal VOs)**
 - Number of Communities (internal VOs), i.e. managed in Check-in COmanage (on a monthly basis)
 - Capacity strategy: lag
- **Number of logins**
 - Number of IdP/SP logins (on a monthly basis)
 - Capacity strategy: lag

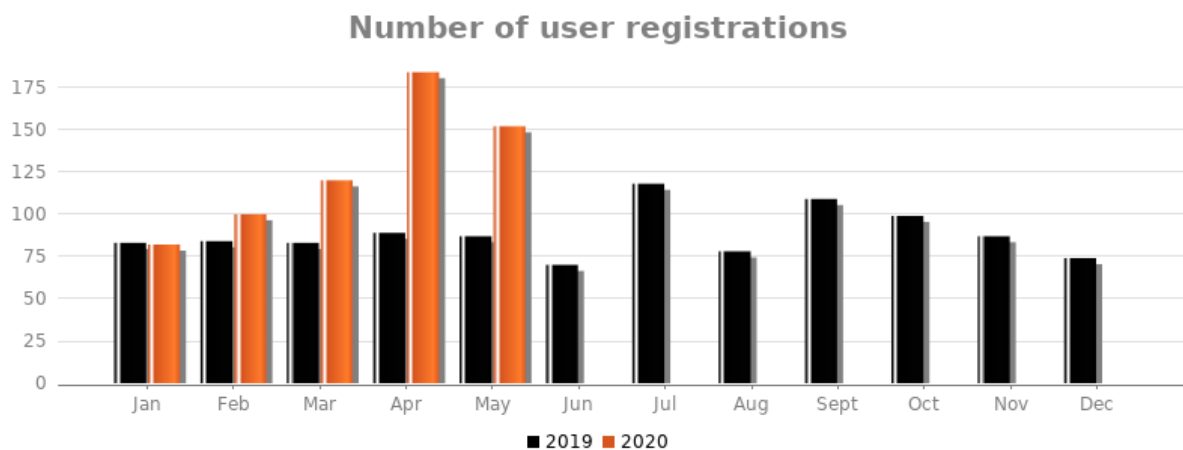


Fig. 1 - Number of new registered users on a monthly basis.

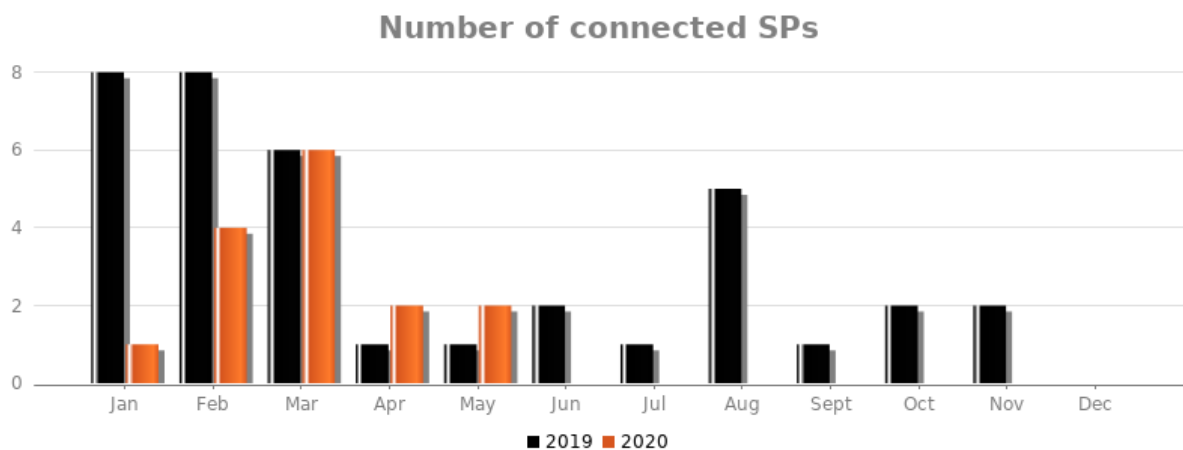


Fig. 2 - Number of new SPs integrated on a monthly basis.

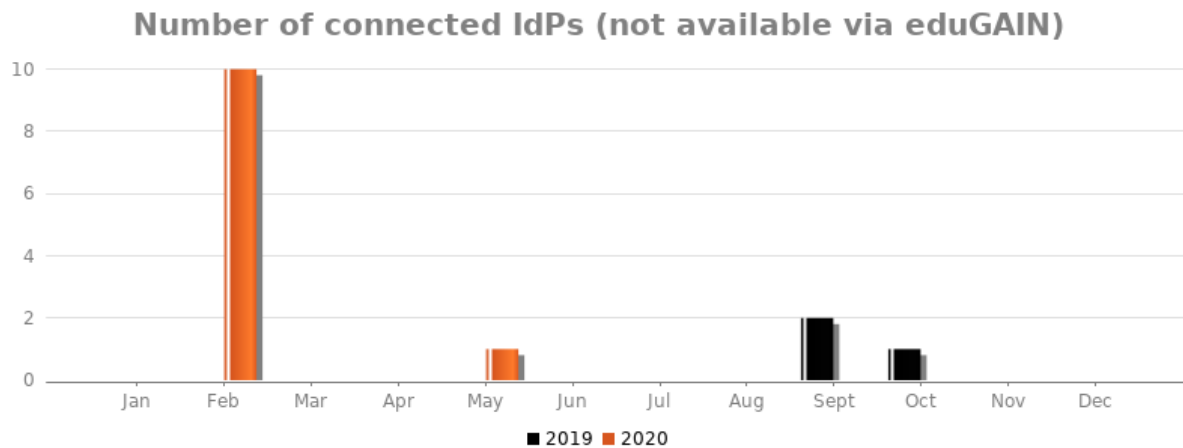


Fig. 3 - Number of new IdPs (not in eduGAIN) integrated on a monthly basis.

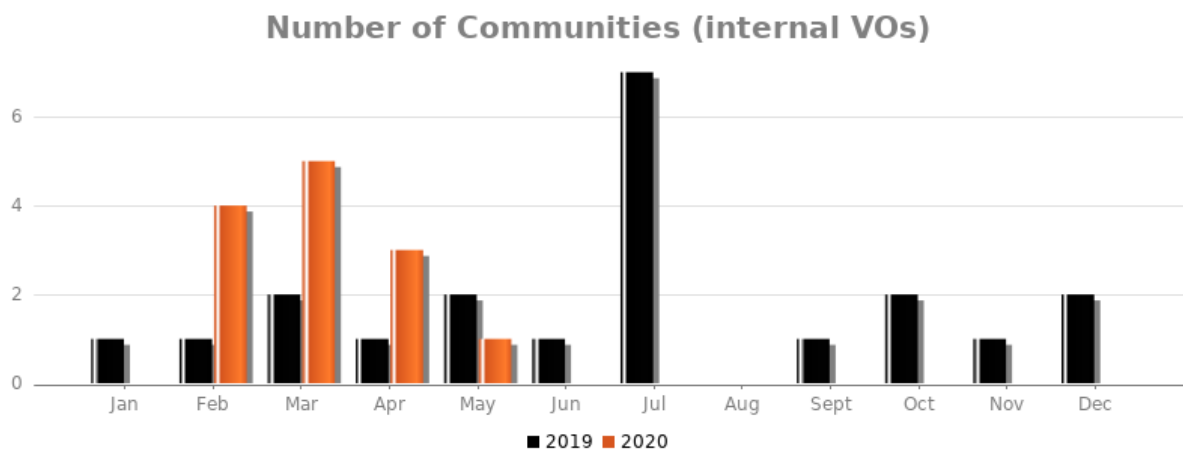


Fig. 4 - Number of new communities registered in and managed through Check-in.

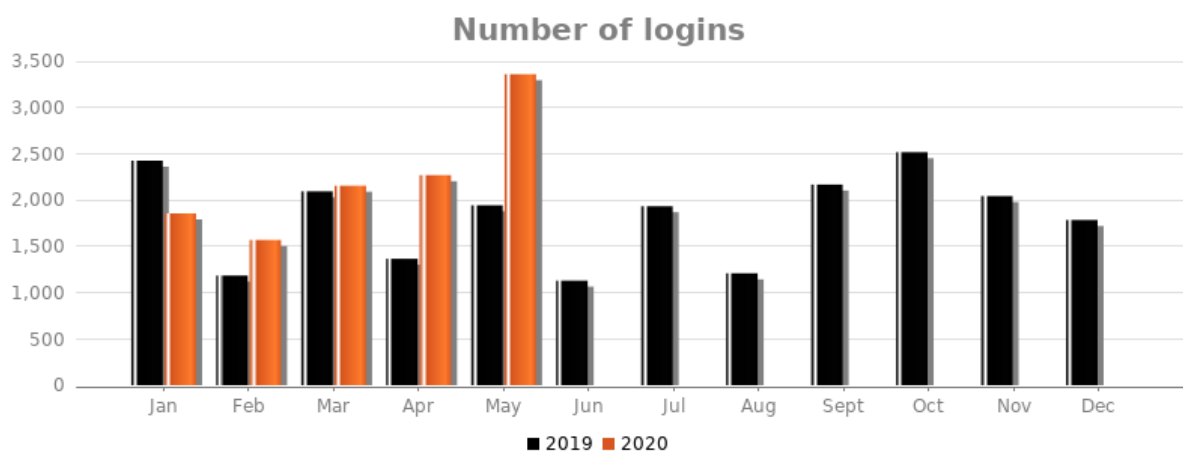


Fig. 5 - Number of IdP/SP logins.

Fig. 1 indicates that every month there is a relevant number of users accessing the Check-in service for the first time, even though this doesn't seem to produce noticeable effects on the number of logins (Fig. 5), either when Check-in is used as Attribute Authority by some federated IdPs or when the users login through Check-in to access the integrated services (the value in May 2020 may be due to the

EOSC-hub Week 2020, when the EOSC-hub services in general received a lot of advertisement and there were also demos and training events). Considering also the other diagrams, we believe that the service demand will not change in the next 6/12 months and the current capacity is sufficient to deal with it.

4.1.2 Human (Operation and Tickets handling) SLIs

- **Number of GGUS tickets**
 - Number of tickets handled by the provider (regardless category and priority). The statistic is collected by the EGI Helpdesk system (GGUS)
 - Capacity strategy: lag
- **Tickets average response time**
 - Average response time (in working days) to the tickets (regardless category and priority). The aim is to check if a relatively high number of tickets can produce an increase of the (average) response time. The statistic is collected by the EGI Helpdesk system (GGUS).
 - Capacity strategy: lag

The monthly amount of tickets (Fig. 6) might seem relatively high, but the support team handled them without particular difficulties, and the service in general didn't suffer capacity issues.




Fig. 6 - Number of tickets handled by the service support unit.

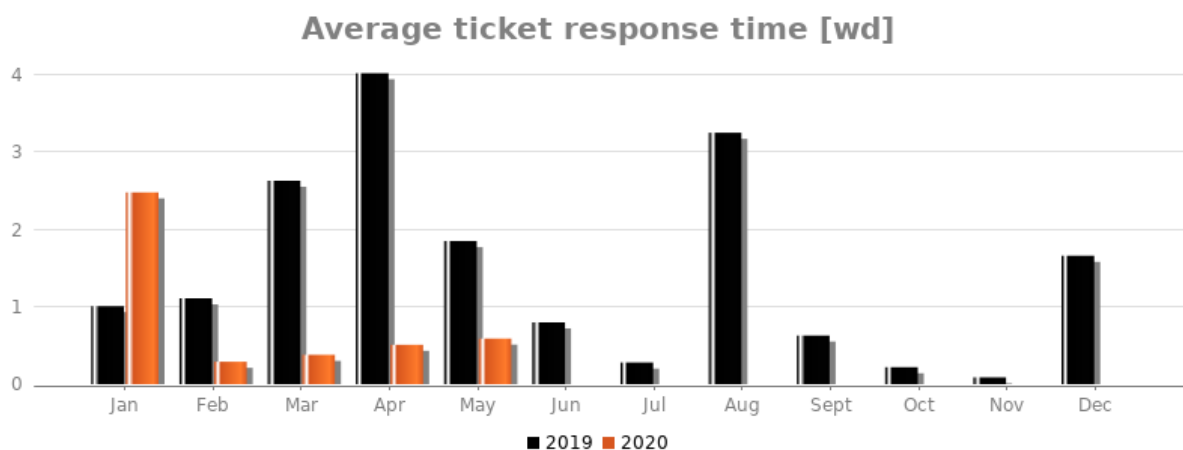


Fig. 7 - Average response time (working days) to the tickets submitted to the service support unit.

4.2 Configuration Management Database - EGI GOCDDB

GOCDDB⁷, a component of the EOSC-hub CMDB, is a central registry to record information about the topology of e-Infrastructures. This includes entities such as Operations Centres, Resource Centres, service endpoints and their downtimes, contact information and roles of users responsible for operations at different levels. The service enforces a number of business rules and defines different grouping mechanisms and object-tagging for the purposes of fine-grained resource filtering.

4.2.1 Software/Service Performance SLIs

- **Average concurrent webserver sessions**
 - The average number of concurrent TCP sessions to production webserver per calendar month - max capacity: 255
 - Capacity strategy: lag

Table 3 – Average number of concurrent webserver sessions per month.

Month	Average concurrent webserver sessions
March 2020	1.14
April 2020	0.60
May 2020	1.01

4.2.2 Human (Operation and Tickets handling) SLIs

- **Incident and service requests**
 - Number of GGUS tickets handled by the provider (regardless category and priority) per calendar month
 - Capacity strategy: lag
- **Tickets average response time**
 - Average response time (in working days) to GGUS tickets (regardless category and priority) per calendar month
 - Capacity strategy: lag

⁷ <https://goc.egi.eu>




Fig. 8 - Number of tickets handled by the service support unit.




Fig. 9 - Average response time (working days) to the tickets submitted to the service support unit.

The number of tickets submitted each month, also the low average response time have not created any capacity issue for the GOCDDB service.

4.3 Accounting Repository

The Accounting Repository stores compute (serial and parallel jobs), storage, and cloud resource accounting data collected from Resource Centres of the EGI Federation. It consists of both hardware and human support for dealing with new requests for integration and fixing problems.

4.3.1 Software/Service Performance SLIs

- **Average concurrent connections to MySQL backend**
 - The average number of concurrent sessions to MySQL backend per calendar month - max capacity: 151
 - Capacity strategy: lag

4.3.2 Human (Operation and Tickets handling) SLIs

- **Number of tickets**
 - Number of GGUS tickets handled by the provider (regardless category and priority) per calendar month
 - Capacity strategy: lag
- **Tickets average response time**
 - Average response time (working days) to GGUS tickets (regardless category and priority) per calendar month
 - Capacity strategy: lag

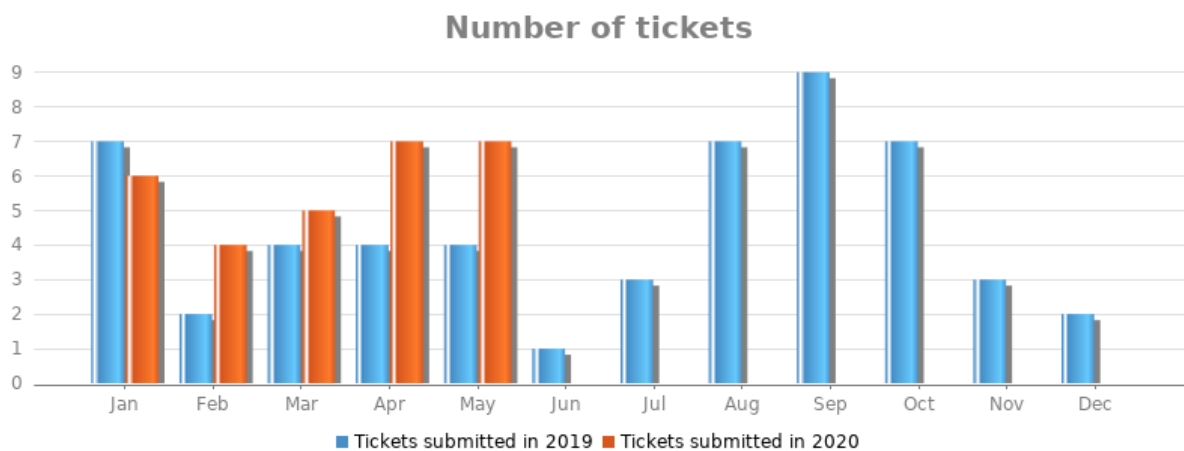


Fig. 10 - Number of tickets handled by the service support unit.

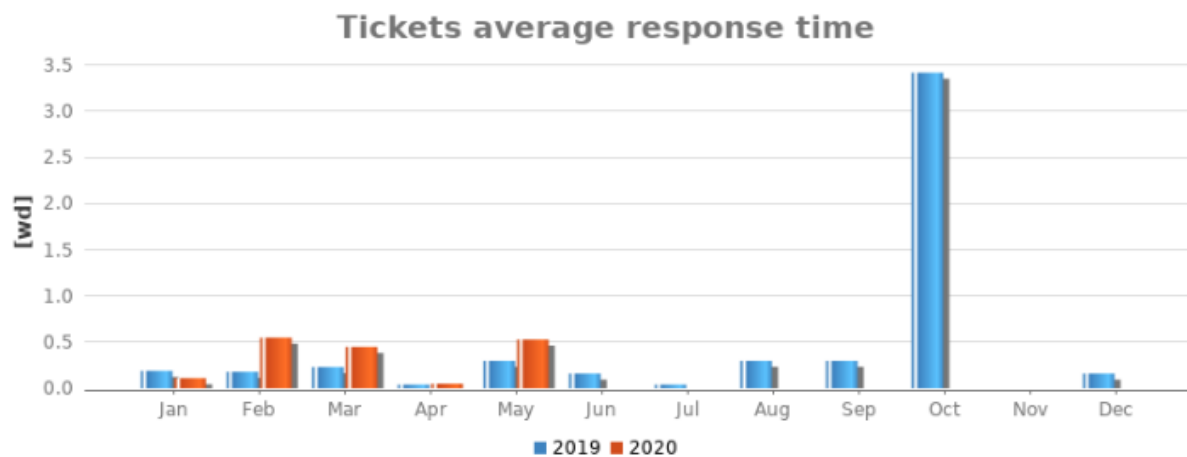


Fig. 11 - Average response time (working days) to the tickets submitted to the service support unit.

The Support Unit for the Accounting Repository had to deal with a relatively small number of tickets each month and the only incident when the response time was higher was due to a ticket not properly progressed in the GGUS (initially wrongly assigned to a different Support Unit).

4.4 Accounting Portal

The EGI Accounting Portal⁸ receives data from APEL and ultimately from sites participating in the EGI and WLCG infrastructures as well as from sites belonging to other Grid organisations that are collaborating with EGI. This is crossed with metadata from other sources to offer an integrated view of accounting data on the EGI Infrastructure.

4.4.1 Software/Service Performance SLIs

- **Total number of hits on the frontend**
 - Number of HTTP requests (Apache hits) in the log for a month
 - Capacity strategy: lag
- **Number of different IP addresses**
 - Number of accesses in the apache logs from different IPs in the month
 - Capacity strategy: lag
- **Total number of visitors**
 - Number of accesses in the apache logs considering also the same IP addresses
 - Capacity strategy: lag

Table 4 – Service Performance SLIs for May 2020 (values available only for the last month - it was agreed with the provider to collect the values on a monthly basis).

SLI	value in May 2020
Total number of hits on the frontend	~ 50000
Number of different IP address	450
Total Number of visitors	11932

The number of hits and visitors recorded every month may seem high but the service doesn't present overload situations, so the resources are enough to guarantee the quality of service agreed in the OLA.

4.4.2 Human (Operation and Tickets handling) SLIs

- **Number of tickets**
 - Tickets directed to the "Accounting Portal" Support Unit in GGUS per month
 - Capacity strategy: lag
- **Tickets average response time**
 - Average response time (in working days) to GGUS tickets (regardless category and priority) per calendar month
 - Capacity strategy: lag

⁸ <https://accounting.egi.eu>

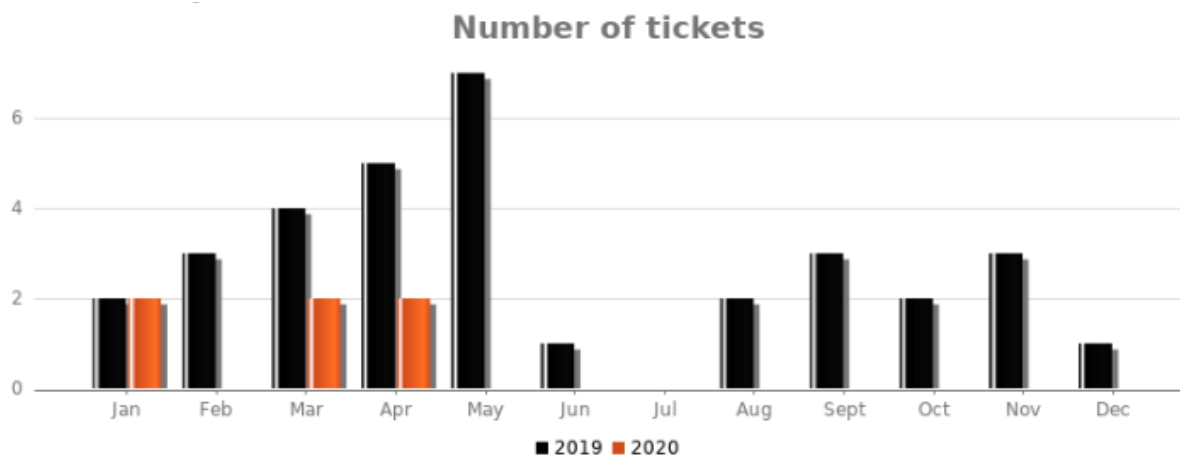


Fig. 12 - Number of tickets handled by the service support unit.

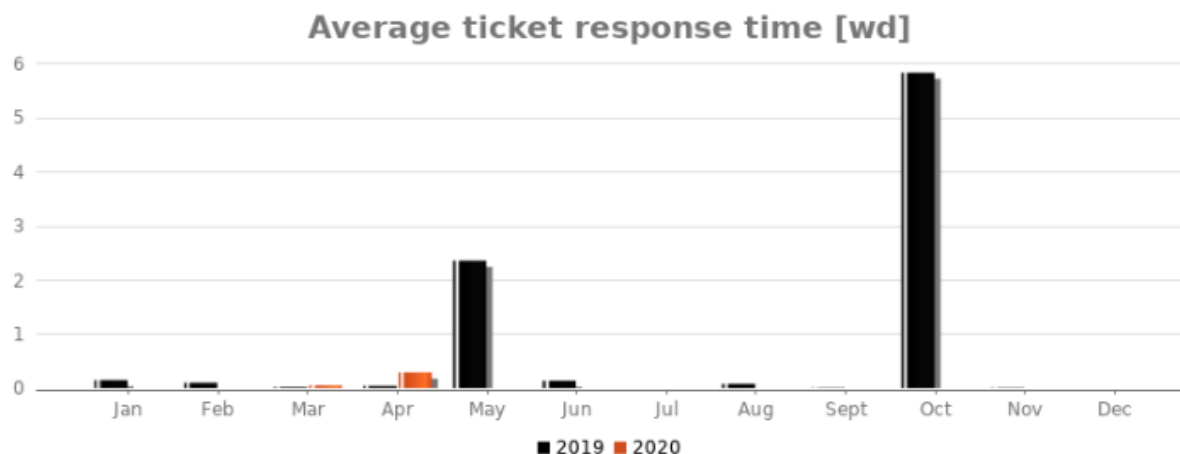


Fig. 13 - Average response time (working days) to the tickets submitted to the service support unit.

The number of tickets submitted on a monthly basis didn't produce capacity issues to the service staff so far; the high value for the average response time recorded in Oct 2019 was due to a ticket properly handled but whose status wasn't changed to "in progress".

4.5 Collaboration Software

It is the software allowing EOSC-hub partners to collaborate. It's composed of the Atlassian Jira issue tracker and Atlassian Confluence wiki.

4.5.1 Software/Service Performance SLIs⁹

- **Number of Jira tickets recorded per month**
 - Total number of tickets opened by month in Jira
 - Capacity strategy: lag

⁹ As reported in Table 1, this capacity plan is not completed yet, so that at the moment of writing this document only a subset of capacity monitoring data was available.

- **Number of concurrent Jira user sessions**
 - Total number of concurrent Jira user sessions
 - Capacity strategy: lag
- **Number of Jira logins per day**
 - Number of Jira logins per day (monthly report of daily average)
 - Capacity strategy: lag
- **Number of Confluence pages**
 - Total number of Wiki pages recorded in Confluence
 - Capacity strategy: lag
- **Number of concurrent Confluence user sessions**
 - Total number of concurrent Confluence user sessions
 - Capacity strategy: lag
- **Number of Confluence logins per day**
 - Number of Confluence logins per day (monthly report of daily average)
 - Capacity strategy: lag

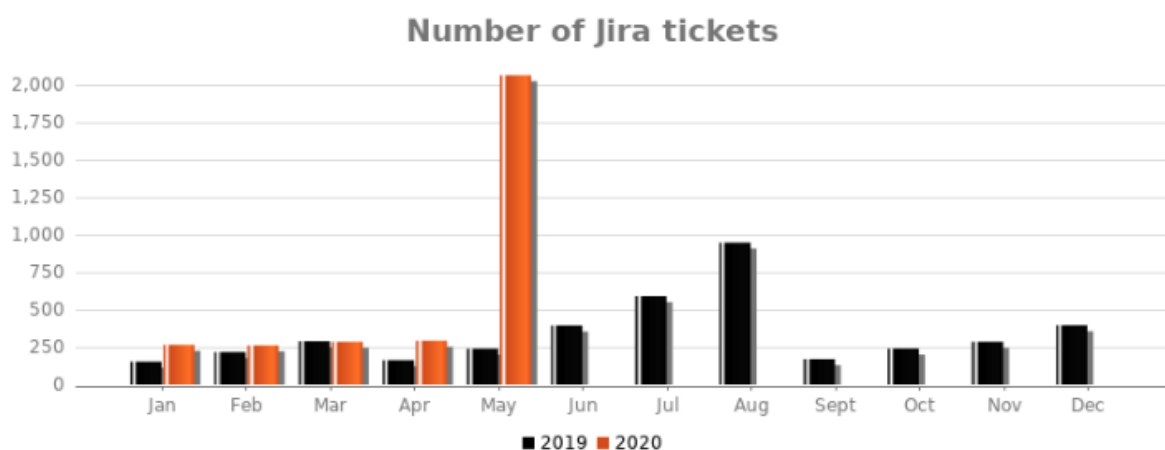


Fig. 14 - Number of new tickets created in Jira by month.

The spike in newly created Jira tickets in May 2020 was due to fake service orders made in the Marketplace by the participants at the EOSC-hub week. Despite this, no hardware or software related capacity issues have been recorded.

4.5.2 Human (Operation and Tickets handling) SLIs

- **Number of tickets**
 - Number of tickets handled by the provider (regardless category and priority). The statistic is collected by the EGI Helpdesk system (GGUS)
 - Capacity strategy: lag
- **Tickets average response time**
 - Average response time (in working days) to GGUS tickets (regardless category and priority) per calendar month
 - Capacity strategy: lag

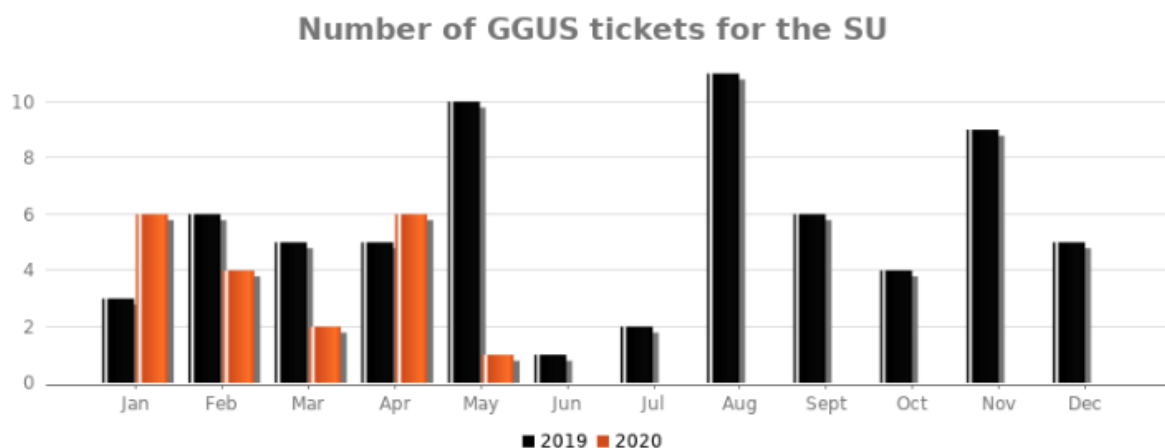


Fig. 15 - Number of tickets handled by the service support unit.

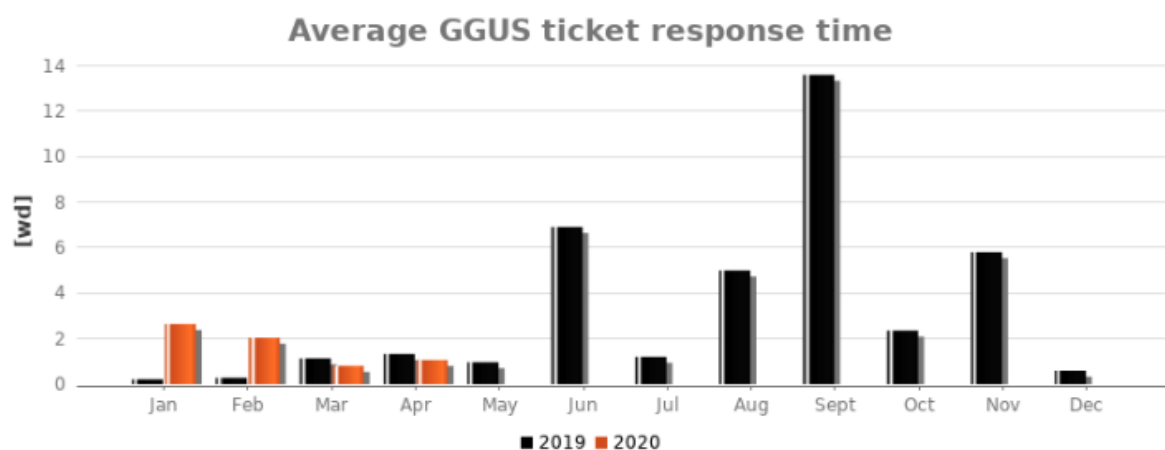


Fig. 16 - Average response time (working days) to the tickets submitted to the service support unit.

The number of tickets handled monthly by the Collaboration Tools SU has been less than 10 in 2020 and most of 2019, and it has not produced any specific capacity issue. Regarding a relatively higher value for the average response time (Sep 2019), this was due mainly to a single ticket that, while the reported issue had been properly progressed and solved, had been poorly dealt with in GGUS portal (late switch to “in progress” status). Overall, we consider that both hardware and human capacities of this service are sufficient to deal with possible surge in demand and to guarantee an acceptable performance level.

4.6 EOSC-hub Helpdesk

The EOSC-hub Helpdesk¹⁰ service is the central point where users can request support for the services included in the EOSC Service Catalogue. The support request is made by creating a ticket (either an Incident or a Service Request) which will be addressed to the affected service support unit.

¹⁰ <https://helpdesk.eosc-hub.eu/>

4.6.1 Software/Service Performance SLIs

- **Number of submitted tickets**
 - Total number of tickets submitted through the EOSC hub Helpdesk
 - Capacity strategy: lag
- **Number of ticket notifications**
 - Number of email notifications sent after a ticket submitted, updated, and closed.
 - Capacity strategy: lag

As it emerges from the diagrams in Figures 17 and 18, the usage of the EOSC helpdesk increased in the first half of 2020 compared to 2019, but the system can handle the demand even if it becomes greater in the next 6/12 months.




Fig. 17 - Total number of tickets submitted through the EOSC Helpdesk.

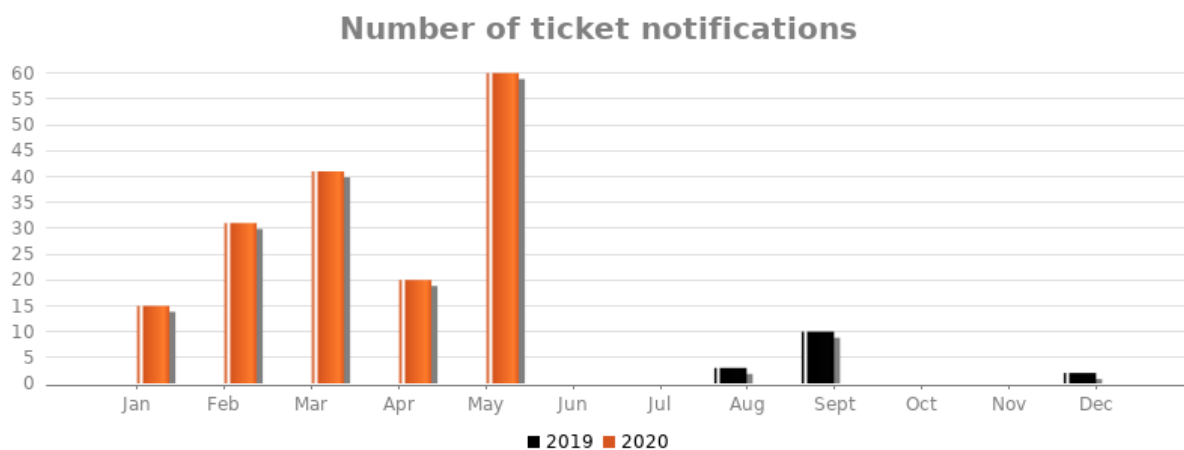


Fig. 18 - Number of ticket notifications sent as a ticket is submitted, updated, and closed.

4.6.2 Human (Operation and Tickets handling) SLIs

- **Number of tickets**

- Total number of tickets assigned to EOSC-hub Helpdesk support unit (the source is the EOSC-hub Helpdesk itself): 1
- Capacity Strategy: lag
- **Average ticket response time**
 - Average response time across all tickets regardless category and priority
 - Capacity strategy: lag

The Helpdesk support unit received only 1 ticket so far and we don't foresee the risk of the human capacity saturation for the next 6/12 months.

4.7 EOSC Marketplace

EOSC-hub Marketplace¹¹ (MP) is a user-facing platform where productional EOSC-hub services can be promoted, discovered, ordered and accessed. Interaction between the user and e-infrastructure involves numerous processes from a service management point of view which requires the Marketplace as a business tool which supports partial service catalogue management, order management, the management of service level agreements (SLAs, OLAs) and the service reporting.

4.7.1 Software/Service Performance SLIs

- **Number of users/visitors**
 - People visiting the service, in the month
 - Capacity strategy: lag
- **Services and resources (numbers of entries)**
 - Number of services and resources added to Marketplace per calendar month
 - Capacity strategy: lag
- **Number of orders handled (per day, the daily average per month)**
 - Number of service orders handled in Jira per calendar month
 - Capacity strategy: lag

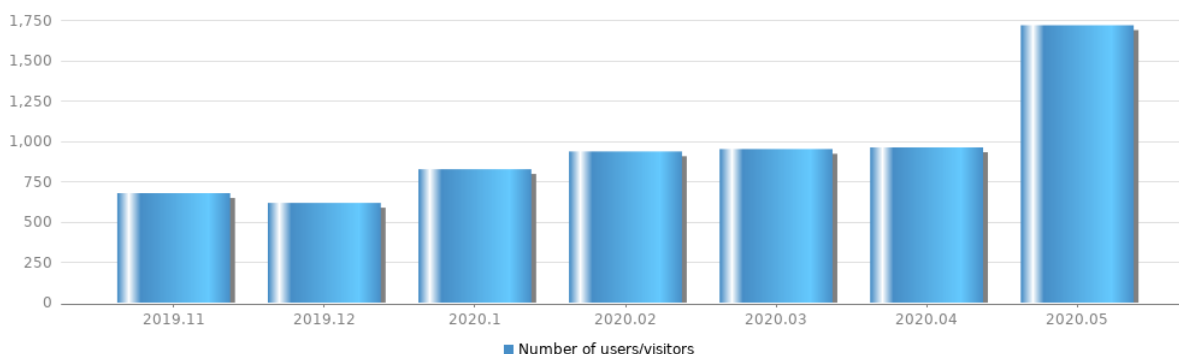


Fig. 19 - Number of users visiting the Marketplace.

¹¹ <https://marketplace.eosc-portal.eu>

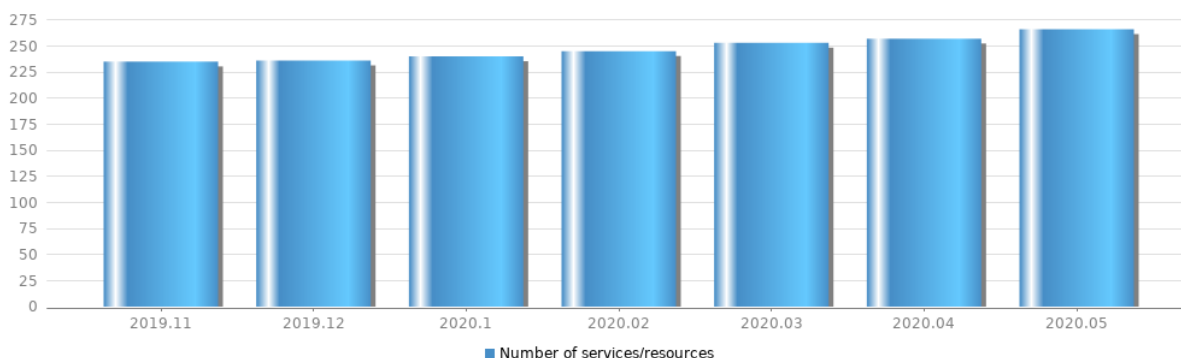


Fig. 20 - Number of services registered in the Marketplace.

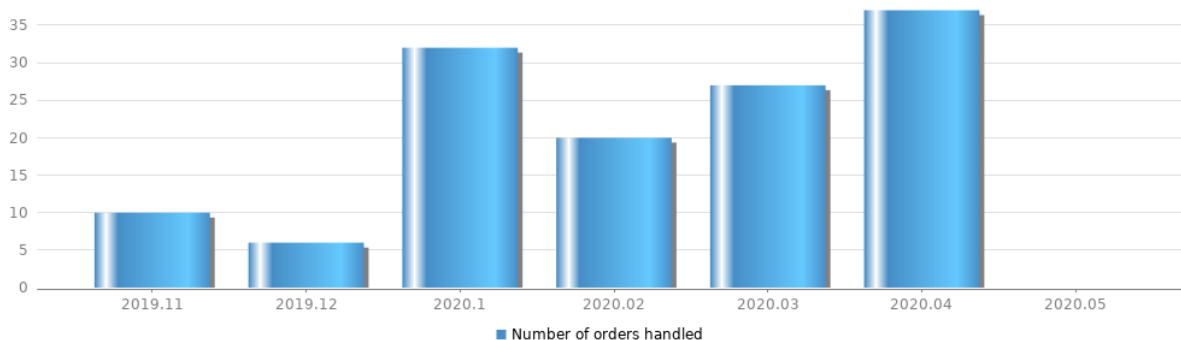


Fig. 21 - Number of service orders handled per month.

As shown in Fig. 19, in May 2020 there was a higher number of users visiting the Marketplace compared to the previous months due to the EOSC-hub Week 2020. For the same reason, the Number of Orders handled in May 2020 is not reported in Fig. 20 because the participants to the EOSC-hub Week made a number of fake service orders which produced noise to the statistics. In general the service didn't show capacity issues and the current capacity should be enough to respond to the future demand even if it will increase.

4.7.2 Human (Operation and Tickets handling) SLIs

- **Number of tickets**
 - Number of GGUS tickets handled by the provider (regardless category and priority) per calendar month
 - Capacity strategy: lag
- **Average ticket response time**
 - Average response time across all tickets irrespective of urgency
 - Capacity strategy: lag

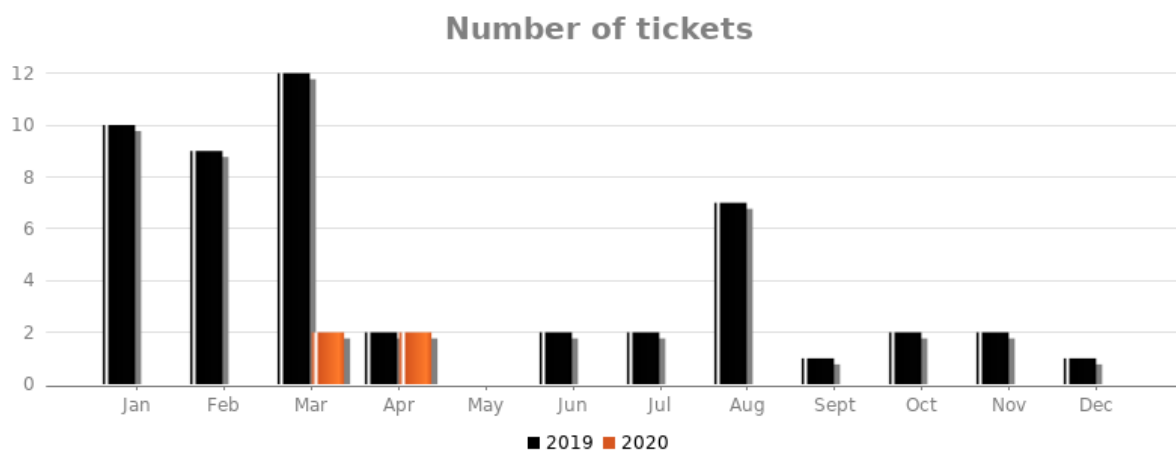


Fig. 22 - Number of tickets handled by the service support unit.




Fig. 23 - Average response time (working days) to the tickets submitted to the service support unit.

Similarly to other services, sometimes the support team replied to the tickets within the service level targets but it didn't properly change the (single) ticket status to "in progress", causing in this way a relatively high value in the average response time (Fig. 23).

4.8 EOSC-hub Monitoring

The EOSC-hub Monitoring service, based on ARGO¹², is a central tool to monitor the several EOSC-hub services and to produce performance reports.

4.8.1 Software/Service Performance SLIs¹³

- **Number of connections / contacts via web**
 - Total number of accesses to the web interface
 - Capacity strategy: lag
- **Number of monitored service endpoints**

¹² <http://argoeu.github.io/index.html>

¹³ As shown in Table 1, this capacity plan is not yet completed: we managed to agree with the provider the SLIs list and we are waiting for receiving the required capacity monitoring data (only the Human SLIs data are already available to us).

- Total number of monitored endpoints
- Capacity strategy: lag
- **Number of probes**
 - Total number of probes in the system
 - Capacity strategy: lag
- **Number of probes executions** (per day, daily average per month)
 - Total number of times that the probes are executed each day
 - Capacity strategy: lag
- **Number of notifications** (per day, daily average per month)
 - Total number of email notifications that are sent to the providers who enabled this feature
 - Capacity strategy: lag

4.8.2 Human (Operation and Tickets handling) SLIs

- **Number of tickets**
 - Number of tickets handled by the provider (regardless category and priority). The statistic is collected by the EGI Helpdesk system (GGUS)
 - Capacity strategy: lag
- **Tickets average response time**
 - Average response time (in working days) to the tickets (regardless category and priority). The aim is to check if a relatively high number of tickets can produce an increase of the (average) response time. The statistic is collected by the EGI Helpdesk system (GGUS).
 - Capacity strategy: lag

The support team handled the tickets (Fig. 24 and 25) without particular problems even when a higher number of tickets than usual was assigned to it, and we can assume that this will continue in the same way also in case of an increase in service demand.




Fig. 24 - Number of tickets handled by the service support unit.

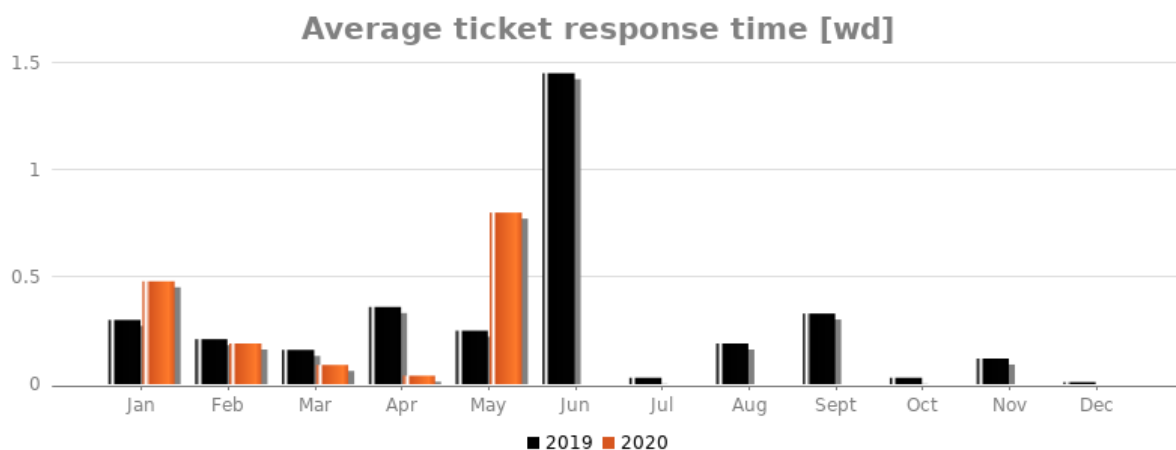


Fig. 25 - Average response time (working days) to the tickets submitted to the service support unit.

4.9 EOSC Portal

The EOSC Portal¹⁴ website is a service complementing the EOSC Marketplace, which provides information content about the EOSC initiative.

4.9.1 Software/Service Performance SLIs

- **Number of Users/Visitors**
 - Users who have initiated at least one session in a given month
 - Capacity strategy: lag
- **Number of Sessions**
 - Total number of Sessions within the month. A session is the period time a user is actively engaged with your website, app, etc. All usage data (Screen Views, Events, Ecommerce, etc.) is associated with a session.
 - Capacity strategy: lag
- **Number of Page Views:**
 - Pageviews is the total number of pages viewed. Repeated views of a single page are counted
 - Capacity strategy: lag

¹⁴ <https://eosc-portal.eu/>

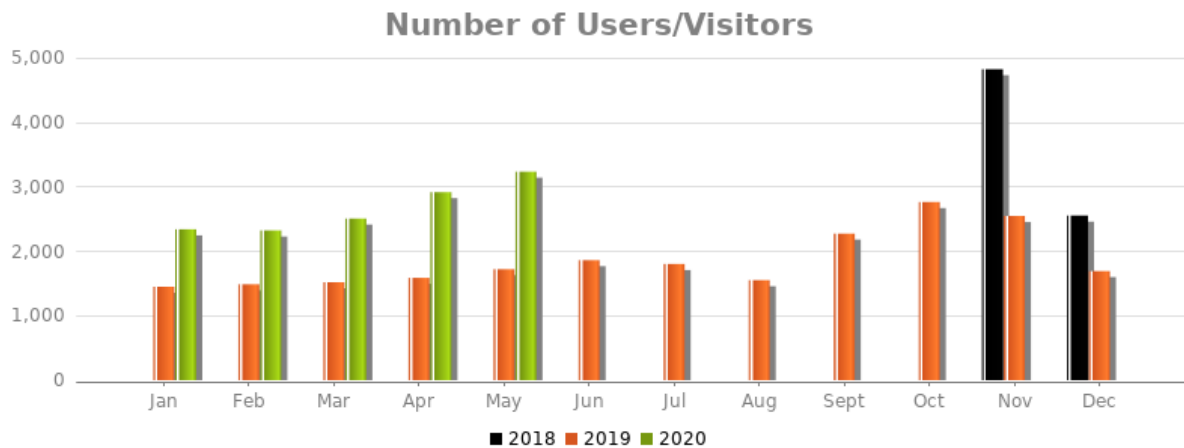


Fig. 26 - Number of visitors to the EOSC Portal over the past months.

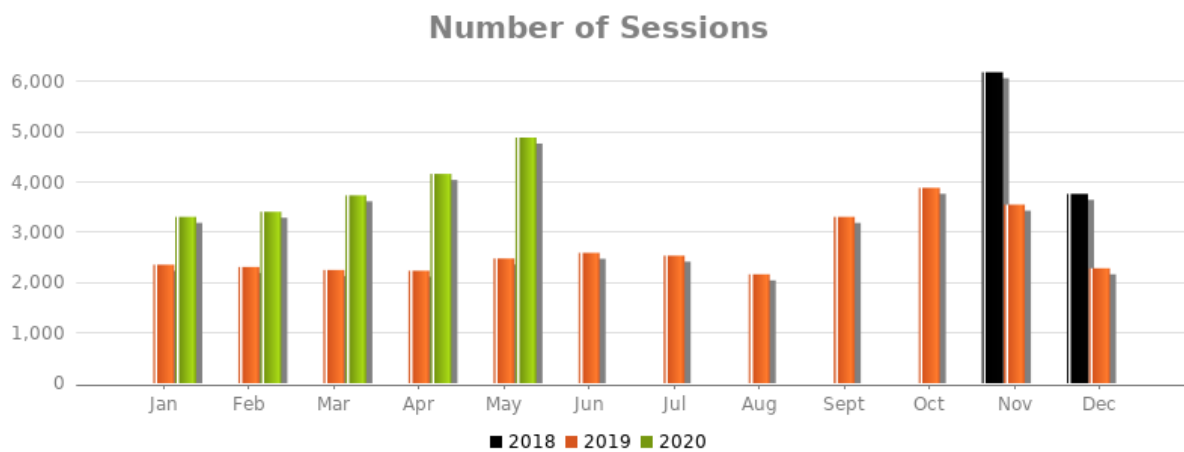


Fig. 27 - Number of sessions over the past months.




Fig. 28 - Number of page views over the past months.

After the spike in the number of accesses to the website recorded in Nov 2018 when the service was launched for the first time, the “number of visitors” and the other two SLIs didn’t register particularly high values. It appears that from the beginning of this year there is an increase of the usage of the

website: even if the increase will continue in the next 6/12 months, we don't think that a capacity adjustment would be necessary.

4.9.2 Human (Operation and Tickets handling) SLIs

- **Number of tickets**
 - Number of tickets handled by the provider (regardless category and priority). The statistic is collected by the EGI Helpdesk system (GGUS)
 - Capacity strategy: lag
- **Tickets average response time**
 - Average response time (in working days) to the tickets (regardless category and priority). The aim is to check if a relatively high number of tickets can produce an increase of the (average) response time. The statistic is collected by the EGI Helpdesk system (GGUS)
 - Capacity strategy: lag

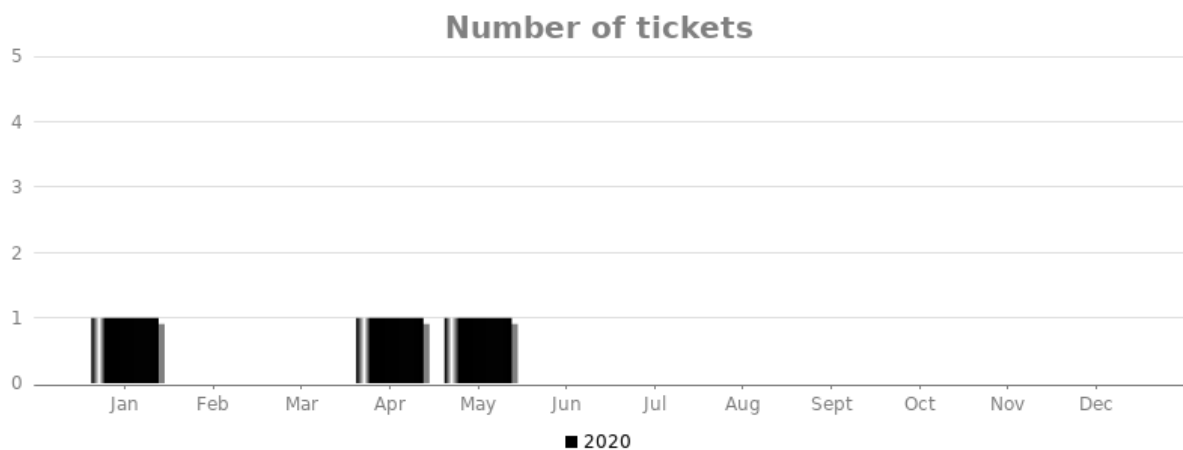


Fig. 29 - Number of tickets handled by the service support unit.

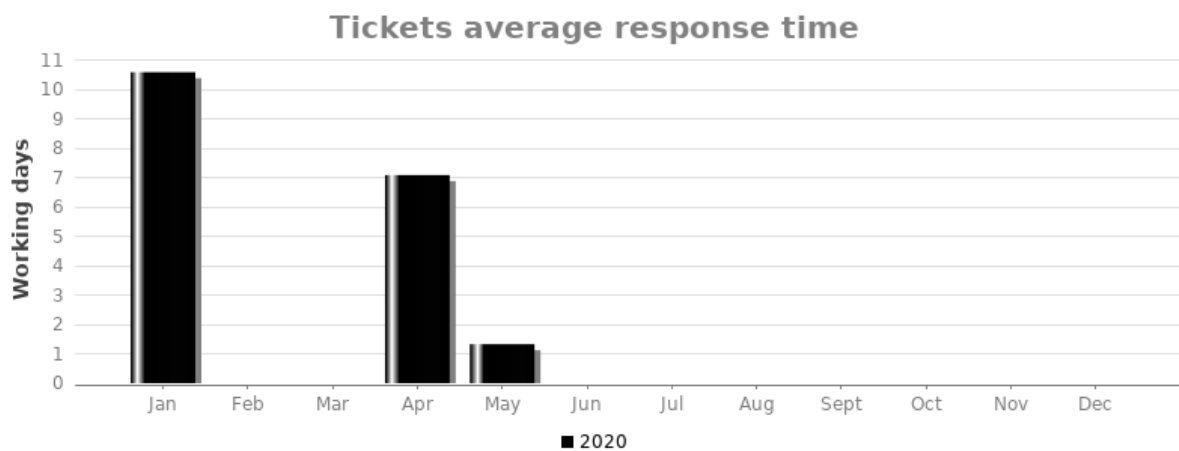


Fig. 30 - Average response time (working days) to the tickets submitted to the service support unit.

As shown by the diagrams in Figures 29 and 30, the EOSC Portal support unit received only 3 tickets since it was created in Dec 2019; the related high values in the response time were produced by not having changed the ticket status to “in progress” even though the tickets were accordingly handled.

4.10 Operations Portal

The Operations Portal¹⁵ provides several capabilities which support the EGI and EOSC-hub daily operations, such as the broadcast tool, VO registration and management facilities, a security dashboard and an operations dashboard that is used to display information about failing monitoring probes and to open tickets to the affected Resource Centres, supporting the central infrastructure oversight activities.

4.10.1 Software/Service Performance SLIs

- **Number of Broadcasts**
 - Number of broadcasts sent through the tool on a monthly and yearly basis
 - Capacity strategy: lag
- **Number of Broadcast's recipients**
 - Total number of potential recipients of the broadcasts. This is a cumulative number.
 - Capacity strategy: lag
- **Number of subscriptions to Downtime tool**
 - Number of recipients getting the downtime notifications. This is a cumulative number.
 - Capacity strategy: lag
- **Number of VOs in production and decommissioned**
 - Number of active and decommissioned VOs. This is a cumulative number.
 - Capacity strategy: lag
- **Number of users captured from VOMS**
 - Number of VOs users registered in the Operations Portal database. This is a cumulative number.
 - Capacity strategy: lag
- **Number of issues in the operations dashboard**
 - Number of issues recorded from the monitoring service. This is a cumulative number.
 - Capacity strategy: lag

¹⁵ <http://operations-portal.egi.eu>

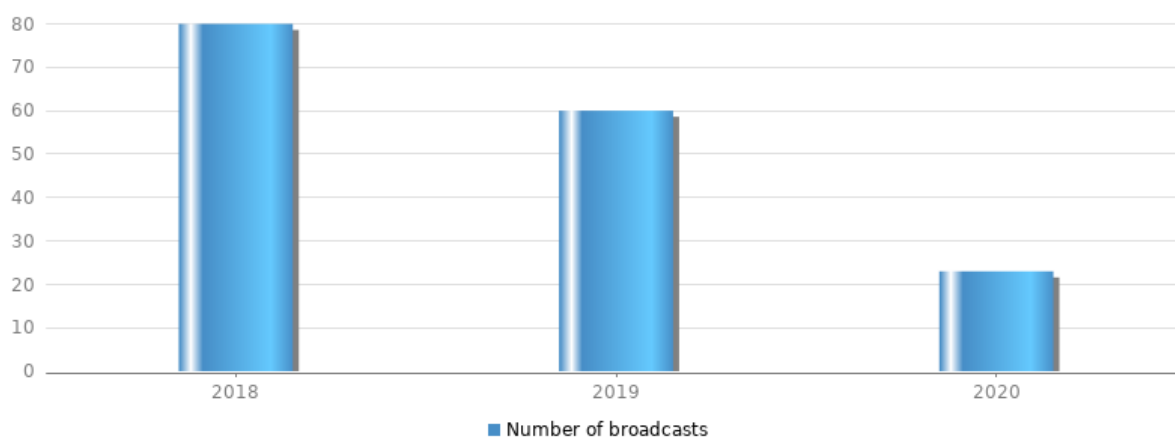


Fig. 31 - Number of broadcasts sent over the past years (until May 2020).

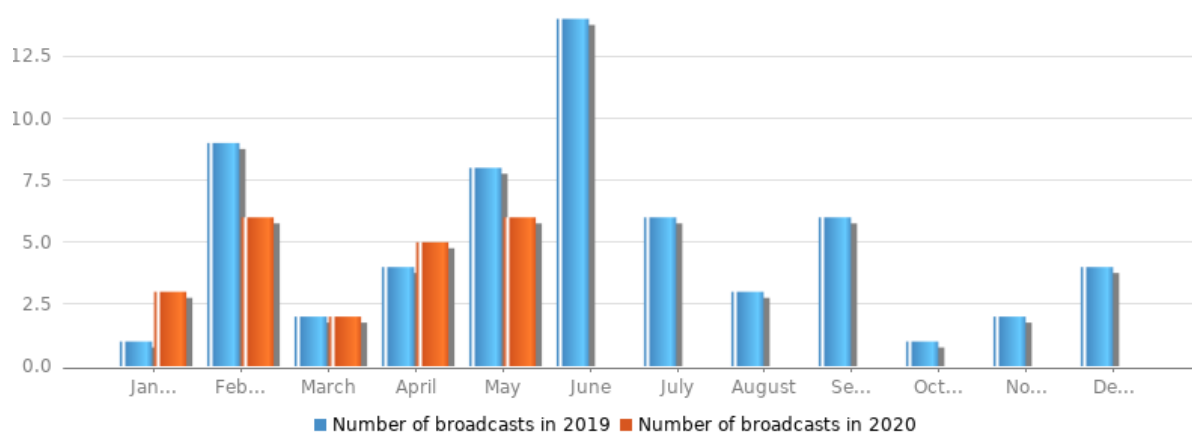


Fig. 32 - Number of broadcasts sent over the past months.

The usage of the broadcast tools, as shown in Figures 31 and 32, doesn't present periods of high demand and we can assume that in the next period it will not deviate from the current trend. Besides, the cumulative numbers provided in Table 5 are relatively stable and their increase (or decrease) rate is very short, so we can assume that the capacity of the service doesn't need to be modified even if there is an increase of demand.

Table 5 – Aggregate values for some Service Performance SLIs.

Broadcasts recipients	1645
Downtime subscriptions	270
VOs in production	258
VOs decommissioned	168

VOs users in DB since 2010	95344
Operations dashboard issues (since 2017)	75k

4.10.2 Human (Operation and Tickets handling) SLIs

- **Number of GGUS tickets**
 - Number of tickets handled by the provider (regardless category and priority). The statistic is collected by the EGI Helpdesk system (GGUS).
 - Capacity strategy: lag
- **Tickets average response time**
 - Average response time (in working days) to the tickets (regardless category and priority). The aim is to check if a relatively high number of tickets can produce an increase of the (average) response time. The statistic is collected by the EGI Helpdesk system (GGUS).
 - Capacity strategy: lag

Looking at the amount of tickets handled by the Operations Portal team (Figure 31), it appears that the human capacity area of the service didn't suffer particular issues; the high value in the average response time (Figure 32) recorded in November 2019, compared to the other months, was due to a single ticket whose status wasn't properly set to "in progress", even though it was processed according to the service OLA. We can assume that the number of staff and the effort allocated to deliver the service are enough to cope with an eventual increase of demand and to guarantee an adequate performance level.




Fig. 33 - Number of tickets handled by the service support unit.




Fig. 34 - Average response time (working days) to the tickets submitted to the service support unit.

5 Conclusion

We have defined the criteria for creating a capacity plan according to the FitSM requirements, and based on this we have produced capacity plans for the services in the Hub Portfolio to assess if the resources provided for the service delivery are sufficient to meet the current demand at the quality level agreed in the several OLAs. At the same time we have provided our recommendations on future capacity requirements that may be necessary to meet an eventual increase in the service demand. We have ensured that the capacity of our services is regularly monitored and we have also planned periodic reviews of the capacity plans: if any capacity issue will be identified, we can proactively intervene, discussing and planning along with the providers either any capacity adjustment or any performance improvement that might be necessary to meet the agreed service level.

From the data collected it appears that current capacity of the services is enough to cope with the current and the future service demand, even in the case of an increase of service utilisation over the next 6/12 months.