Ref. Ares(2022)9020252 - 30/12/2022

Digital Cognitive Biomarkers

Horizon 2020 Project LETHE
"A personalized prediction and intervention model for early detection and reduction of risk factors causing dementia, based on AI and distributed Machine Learning."

Research and Innovation Action
H2020-SC1-DTH-2020-1
GA 101017405

Duration: 48 months from 01/01/2021
Coordinator: Sten Hanke, FH JOANNEUM GESELLSCHAFT MBH

Deliverable ID.:	**D5.2**	
Deliverable title:	**Deployment of the LETHE infrastructure**	
Planned delivery date:	30/12/2022 (M24)	
Actual delivery date:	30/12/2022 (M24)	
Editors:	Ignacio Lamata Martínez (EGI), Catalin Condurache (EGI), Ville Tenhunen (EGI)	
Contributing partners:	FHJ, FORTH, ER, Info	
Internal reviewer:	Samuel Lahtinen (COMB)	
Checked and released by:	Sten Hanke (FHJ)	
Dissemination Level:	**X**	PU = Public;
		CO = Confidential

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101017405.

This deliverable reflects only the authors' view and the Commission is not responsible for any use that may be made of the information it contains.

Type:	Other

Document information and history

Deliverable description (from DoA)
Deployment of the LETHE infrastructure (EGI), This will be the service prototype, with a report to outline the final set-up of the cloud-based infrastructure for the LETHE project including the EOSC services tailored and links to open-source repositories of software used to address the project needs.

Please refer to the Project Quality Handbook for guidance on the review process and the release numbering scheme to be used in the project.

Release number	Release date	Author [Person and Organisation]	Milestone*	Release description /changes made
V. 0.9	22.12.2022	Ville Tenhunen (EGI)	Proposed	Finalized version

** The project uses a multi-stage internal review and release process, with defined milestones. Milestone names include abbreviations/terms as follows:*

- *TOC = ”Table of Contents” (describes planned contents of different sections);*
- *Intermediate: Document is approximately 50% complete – review checkpoint;*
- *ER = “External Release” (i.e. to commission and reviewers);*
- *Proposed: document authors submit for internal review;*
- *Revised: document authors produce new version in response to internal reviewer comments*
- *Approved: Internal project reviewers accept the document.*

Table of Contents

Table of Figures

List of Tables

1 Executive Summary

The infrastructure of the LETHE project has been built and deployed with a set of services and tools based on the European Open Science Cloud (EOSC). This infrastructure is deployed using cloud computing technologies, which allow for rapid adaptation to provide the necessary computing resources as the project evolves.

For solution developers, LETHE's infrastructure basically follows an *Infrastructure-as-a-Service* (IaaS) model, where they have on-demand access to cloud-hosted servers and storages. With this infrastructure, LETHE components such as the Data Lake or the Data Warehouse create a *Platform-as-a-Service* (PaaS) layer where Artificial Intelligence (AI) developers have on-demand access to a complete, ready-to-use, cloud-hosted platform for developing, running, maintaining and managing their applications.

This deliverable focuses on the infrastructure layer of the LETHE platform, describing the actual hardware devices that support the operation of its different components.

2 About this Document

2.1 Role of deliverable

In this deliverable, the LETHE project infrastructure and its components has been documented. It outlines the final set-up of the cloud-based infrastructure for the LETHE project, including the EOSC services tailored and links to open-source repositories of software used to address the project needs.

Note that this document only addresses the description of the infrastructure, but does not describe the LETHE applications, their components or technologies. Those are documented in other reports and deliverables, such as Deliverables 2.2, 2.3, 5.1 and 5.3. Similarly, this document does not describe the actual deployment process of the infrastructure, which is explained in detail in Deliverable 2.4.

2.2 Relationship to other LETHE deliverables

This document is closely related to the following deliverables:

- D2.2 "*LETHE architecture and scenarios*", which describes the provisional architecture of the platform and the potential scenarios for the LETHE applications, and is relevant to understand the different components involved in the platform.
- D2.3 "*LETHE architecture update*", which contains the final architecture specification of the LETHE application and its infrastructure. This deliverable has not been submitted at this stage.
- D2.4 "*LETHE deployment setup and handbook*", which describes the deployment instructions for the different components of the LETHE platform and is relevant for the understanding of deployment concepts and the explanation of the elements involved.
- D5.1 "*Data Lake and Data Warehouse infrastructure*" describes the setup and the methods of the LETHE repositories for unstructured and structured data. This deliverable has not been submitted at this stage.
- D5.3 "*The LETHE AAI architecture*" describes the security measures implemented on the LETHE platform, especially in terms of authentication and authorisation, and is relevant to the understanding of the platform's security. This deliverable has not been submitted at this stage.

2.3 List of acronyms

The acronyms used in this document are listed in Table 1.

Table 1. Acronyms used in this document

Acronym	Explanation
AES	Advanced Encryption Standard
AI	Artificial intelligence
API	Application Programming Interface

DB	Database
DL	Data Lake
DNS	Domain Name System
DWH	Data Warehouse
EOSC	European Open Science Cloud
GPU	Graphics Processing Unit
IaaS	Infrastructure as a Service
IaC	Infrastructure as Code
IM	Infrastructure Manager
OS	Operating System
PaaS	Platform as a Service
TLS	Transport Layer Security
VCS	Version Control System
VM	Virtual Machine
XTS	XEX-based Tweaked-codebook mode with ciphertext Stealing

2.4 Structure of the document

The rest of this document is organised as follows:

- Section 3 briefly explains some important concepts that are useful to understand the rest of the document.
- Section 4 describes the actual infrastructure deployed for the LETHE project.
- Section 5 discusses the EOSC services that have been used for the development of the infrastructure.
- Section 6 lists the open-source repositories that have been part of the development of the infrastructure.
- Section 7 gives some conclusions.

3 Concepts on cloud infrastructure

This section briefly explains some concepts that are useful for understanding the technical parts of the rest of the document. *Cloud Computing* is explained in Section 3.1, and its relevance lies in the fact that the LETHE platform is implemented in a cloud environment. The concept of *Infrastructure-as-Code* is described in Section 3.2, as this is the approach followed to implement the infrastructure of the LETHE platform. Finally, Section 3.3 briefly introduces some of the systems that manage infrastructure, which are widely used in the LETHE platform to create or modify infrastructure.

3.1 Cloud Computing

Traditionally, the computing hardware infrastructure that supports software applications has run on independent physical servers that were purchased, installed and managed by the organisations behind such applications. This process highlighted several problems, including:

- It requires a costly and time-consuming procurement process, especially in highly bureaucratic organisations, and delivery and installation often take considerable time.
- Servers require expertise to be managed and maintained.
- Servers become outdated over time and need to be upgraded.
- Hardware failure affects all computing components sooner or later, and replacing or fixing server components means unacceptable application downtime. The solution is to duplicate servers for high availability, which is more expensive and involves additional servers to maintain.

Cloud technologies are used to address these problems. They have been made possible by recent advances in computing and have become a common trend in recent decades, as organisations have transitioned from physical servers to cloud environments. Cloud technologies are based on the following principles:

- Servers and other hardware components are no longer in the organisation's premises, but are hosted by a *cloud provider*, who rents them for the organisation. This avoids the need to purchase hardware and allows organisations to dynamically use the exact amount of hardware they need at any given time.
- Servers are maintained by the cloud provider, so the client organisation does not have to devote effort and resources to this task.
- From the customer's perspective, servers and hardware components are nor perceived and managed as physical elements, but as *virtual* ones. This means that servers can be requested and served almost instantaneously and on demand, which helps to easily overcome hardware failures.

This virtual allocation of servers is done through **Virtual Machines** (VM), which are software components that run over physical hardware and emulate and provide the functionality of a physical computer system. One of the advantages of cloud technology is that the virtual infrastructure can be created directly through software instructions, without the need to physically access the servers. Moreover, this virtualisation greatly facilitates *elasticity*, a term used to describe the ability of a system to increase or decrease its resources to adapt to the current workload.

The LETHE platform **uses EGI's Cloud Computing to supply its infrastructure**. As described above, this infrastructure is virtual, not physical, being flexible to change according to the project needs.

3.2 Infrastructure as Code

One of the immediate consequences of using a virtual infrastructure is that it can be defined in software instructions. This is known as **Infrastructure-as-Code** (IaC), an approach that allows the management of the infrastructure hardware resources with source code. For this purpose, the necessary infrastructure is defined in text files, following rules similar to those of programming languages, which can then be used by different tools to implement the virtual infrastructure in a cloud provider. Amongst its many advantages are the following:

- The infrastructure is defined in human-readable files.
- It can benefit from all advances in the software development industry, such as version control systems[1] and Integrated Development Environments[2].
- The process of hardware provisioning becomes consistent, repeatable and less prone to human error.
- It greatly improves the efficiency of the infrastructure creation process, enabling its automation.

In particular, the LETHE project uses **code templates to define its virtual infrastructure** needs, allowing the platform to be deployed in a few hours and with minimal human intervention. Complete information about these templates can be found in [1].

3.3 Infrastructure management systems

There are many systems available to manage infrastructure. This section briefly discusses two of the most prominent systems used in LETHE:

- The **Infrastructure Manager** (IM) [2] is a free and open-source application to assist users with the creation of infrastructure in a cloud environment and the deployment of software. LETHE defines its infrastructure in TOSCA templates [3], which are provided to the IM to create the required infrastructure in the cloud provider. Although the LETHE's current infrastructure is deployed in the EGI Federation, the IM can deploy in different cloud providers, including public environments such as those from Amazon, Google and Microsoft. For a more complete description of how the IM is used in LETHE, see [1].
- **Kubernetes** [4] is an open-source container management system[3] that assists developers in deploying software applications. To do this, Kubernetes uses a set of Virtual Machines joined together in a configuration known as a *cluster*. Kubernetes can be instructed to provide high availability of applications automatically by deploying those applications on the different VMs available in the cluster. If the existing VMs are insufficient or, otherwise, idle, it can increase or decrease the size of the cluster (the number of Virtual Machines in the

infrastructure) to match the workload. Kubernetes is used by some components of the LETHE platform.

4 Description of the infrastructure

This section describes the LETHE infrastructure. It begins by describing the components of the platform in Section 4.1. These components are supported by the infrastructure, which is discussed in Section 4.2. Section 4.2.1 presents some preliminary considerations for the infrastructure, such as the need for encrypted storage for user data. Section 4.2.2 lists the hardware requirements for each component of the platform. Finally, Section 4.2.3 presents the actual set-up and allocation of hardware that makes up the infrastructure.

4.1 Brief introduction to the LETHE platform

The LETHE platform is a relatively complex system involving multiple components that intercommunicate to send or retrieve data. A simplified view of the general architecture is presented in Figure 1, where the different data flows are depicted. On the left side of the figure is the ***Data Generation*** layer, which consists of the systems used by patients to provide medical data and accessed by devices physically owned by patients, such as wearable activity trackers and smartphones.

Once these devices collect patient's data, they are sent to a repository of unstructured and unprocessed data called the ***Data Lake***. The mission of the Data Lake is to ingest digital information as it was provided by the source.

These data are then processed and stored in a structured manner in the ***Data Warehouse***, which is a core storage component of the platform and is used by all the other components to store and retrieve important operational data.

As an auxiliary aid, the ***Notifications*** component is responsible for transmitting reminders to users, essentially transmitting and acknowledging receipt of notifications to/from the users/devices.

The data stored in the Data Warehouse is accessed by the ***Artificial Intelligence*** layer of the platform, which has multiple components that use trained models, reasoning and other tasks to analyse data and extract conclusions.

Both physical ***Social Robots*** and ***Smart Glasses*** are hardware devices that interact with patients, and also collect information. Finally, different ***Dashboards*** use all data available to display information to be assessed by clinicians and experts in the medical field.

All these components are supported by the underlying computing hardware, referred to as ***the infrastructure***.

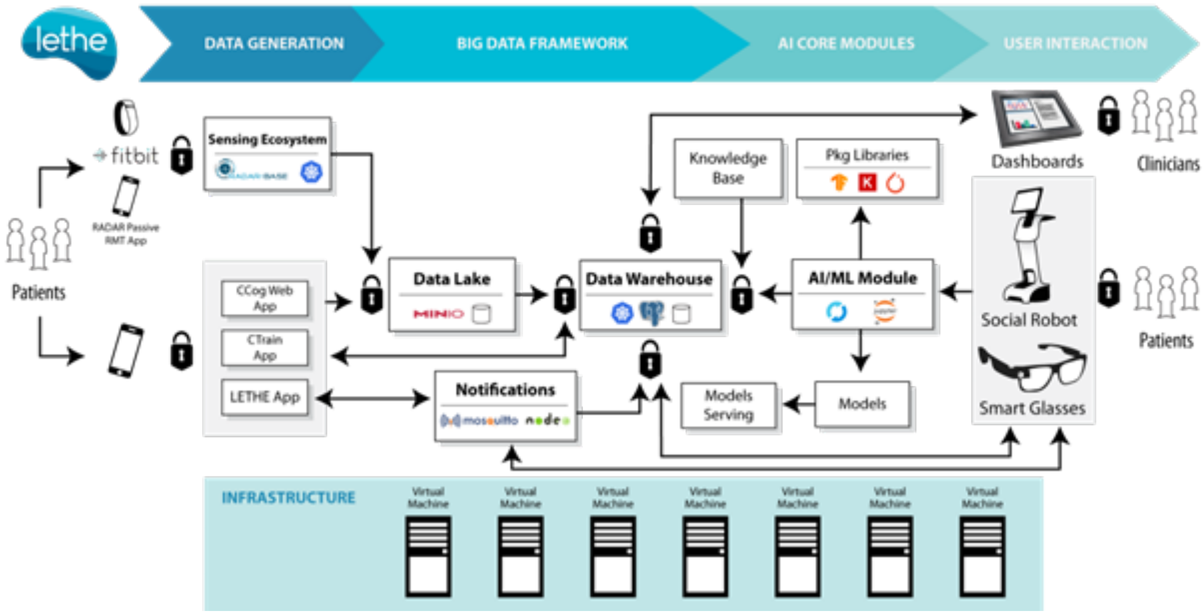

Figure 1. LETHE platform simplified architecture

4.2 Infrastructure overview

The infrastructure is the underlying hardware layer that supports the different components of the LETHE platform. It is created using cloud technologies, as explained in Section 3.1, and can therefore be scaled and restructured by software instructions, on demand. This has multiple advantages, being the most relevant:

- It allows the infrastructure to be **as agile as the software components** of the platform, being able to adapt to rapidly evolving environments.
- The infrastructure can be **designed, operated and managed** in the same way as the platform **software code**. This is the main pillar of the *Infrastructure-as-Code* described in Section 3.2. In fact, the required artefacts are also stored in the same *Version Control System* (VCS) used for the source code of the LETHE components[4].

The infrastructure of LETHE is deployed in one of the providers of the EGI Federation, which uses OpenStack[5] as its cloud technology. Amongst other things, there are three main resources that are managed through the OpenStack platform:

- **Computing power**. The different servers needed by LETHE are implemented as Virtual Machines, each of which is managed independently like a physical server.
- **Storage**. Discs are managed as virtual devices to provide the required storage capacity, and can be attached to the different Virtual Machines as if they were physical drives connected to physical servers.
- **Networking**. Network connectivity is managed and configured to allow Virtual Machines to communicate. This includes traffic routing (to move data packets from a source to a destination) and firewall rules (to block undesired network traffic).

Keeping the OpenStack platform operational is a complex process involving many challenges, and most of this burden is handled automatically by the EGI Federation provider. Thanks to this, operators of LETHE can simply focus on requesting hardware resources from the platform.

Figure 2 depicts an oversimplified version of this hardware allocation: There is a pool of physical hardware resources available in the cloud provider. When a user requests a server, a Virtual Machine is created by the cloud software (OpenStack in this case) and some hardware resources are assigned to it, making a *virtual* allocation of hardware resources. Resources can be added and removed on demand as needed, which enables a more efficient use of hardware resources.

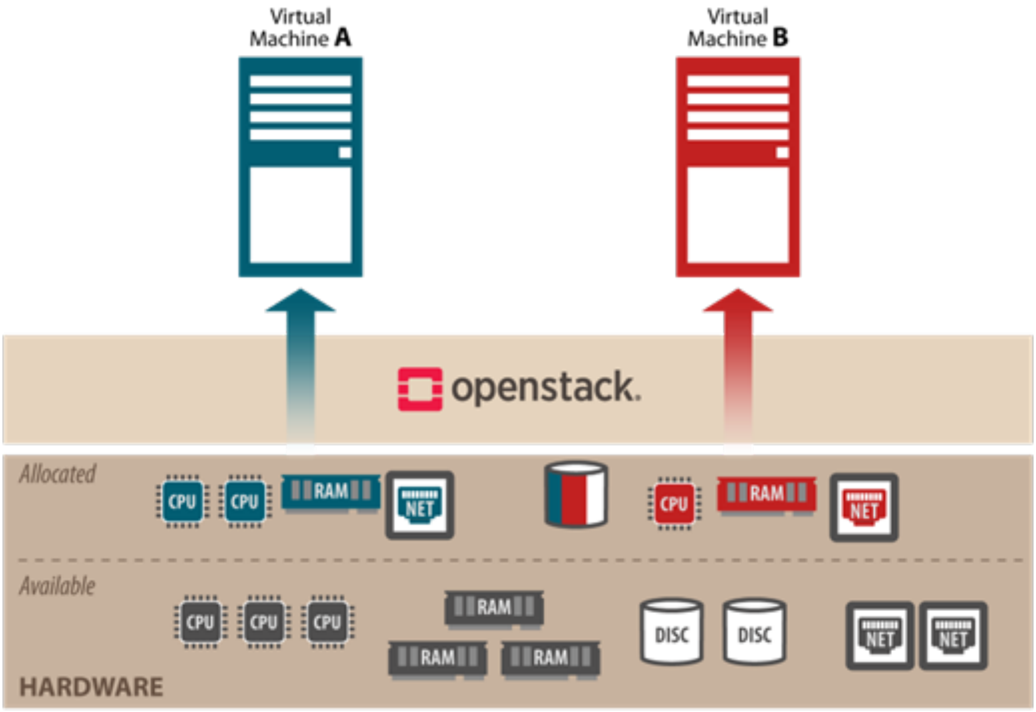

Figure 2. Simplified view of hardware allocation

4.2.1 Preliminary considerations

The LETHE infrastructure is a *disposable* system, meaning that any of the servers may fail at any time and the infrastructure must be ready to be easily recreated. The deployment process of the LETHE platform has been designed for automation, and this includes the capacity to deploy the infrastructure in an automated manner [1]. The only part that is not considered ephemeral and should persist are the users' data, which must be backed-up according to the LETHE back-up policy.

User data managed in LETHE are mostly medical data, which imposes a special treatment on the data, as they are particularly sensitive. Data security is a crucial aspect of the platform, and its hardware infrastructure also plays a relevant role. In terms of infrastructure, this translates into the need to maintain the physical security of data at multiple layers, from limiting physical access to the storage devices to the encryption of the information.

The infrastructure therefore provides special volumes created for the storage of user data. This is achieved through LUKS[6], which ensures the confidentiality of data by encrypting the data volumes with *AES* [5] in *XTS* operation mode (a special mode for data protection on block-oriented storage devices). Note that authenticated encryption is not provided in LUKS by default, and is currently still an experimental feature, meaning that only confidentiality is ensured on the current LETHE platform.

4.2.2 Hardware requirements

The LETHE project is organised in two phases:

- **Phase I** is the model generation, where preparatory work is carried out to harmonise and optimise existing data and design a prediction mode.
- **Phase II** is the intervention implementation and includes the validation and feasibility trials.

Both phases have different needs in terms of infrastructure. The hardware requirements of Phase I are not too demanding, as this is a preliminary phase, and can be found in Table 2. There is a row for each Virtual Machine needed, and the machine specifications are simplified in three categories: (i) CPU cores needed, (ii) Amount of GB of memory RAM needed, and (iii) Additional storage needed (note that this does not include the storage supplied with the VM by default). Additional information is given in the table, such as the Operating System installed or some comments. For Phase I, only two machines with an encrypted volume are needed to conduct its tasks.

Table 2. Hardware requirements for LETHE Phase I

Component	Machine Name	CPU	RAM	STORAGE	OS	COMMENT
Harmonisation	**lethe-01**	8	16	200GB	Ubuntu	
	oneprovider	8	16	500GB	Ubuntu	

The hardware requirements for Phase II are given in Table 3. This phase performs a much heavier workload than Phase I and therefore requires substantially more computing support. Some of these requirements **are adapted in the final provisioning**. For example, VMs requesting three CPU cores are upgraded to four in the actual hardware allocation, which is a more standard measure. Similarly, some storage requests might be excessive and may be reduced in the final allocation.

Table 3. Hardware requirements for the LETHE Phase II

Component	Machine Name	CPU	RAM	STORAGE	OS	COMMENT
Sensing Ecosense	**RADAR-base1**	3	16	100GB + 200GB	Ubuntu	Kubernetes cluster (Control plane)

				(installation)		
	RADAR-base2	3	16	100GB	Ubuntu	Kubernetes cluster (node)
	RADAR-base3	3	16	100GB	Ubuntu	Kubernetes cluster (node)
Data Lake	**Data Lake**	16	32	1TB (data) + 512GB (installation)	Ubuntu	
Data Warehouse	**LETHE-DW-DB-1**	8	32	1TB	Ubuntu	Postgre database Primary
	LETHE-DW-DB-2	4	16	1TB	Ubuntu	Postgre database primary Secondary (Replica)
	LETHE-DW-API-FN	4	8	120GB	Ubuntu	Kubernetes cluster front node
	LETHE-DW-API-WN1	4	8	120GB	Ubuntu	Kubernetes cluster working node
	LETHE-DW-API-WN2	4	8	120GB	Ubuntu	Kubernetes cluster working node
	LETHE-DW-API-WN3	4	8	120GB	Ubuntu	Kubernetes cluster working node
AI Modules	**Knowledge base**	16	16	512GB	Ubuntu	
	Activity Profile Algorithm	16	16	512GB	Ubuntu	Only for prediction (inference) task; no training
	Prediction Server	16	16	512GB	Ubuntu	Only for prediction

						(inference) task; no training
Notifications	**Notification Server**	8	16	256GB	Ubuntu	
Dashboard	**Dashboard**	8	16	256GB	Ubuntu	

These requirements are expected to change slightly during the project. For example, it is foreseen that some **GPU resources** will be needed, which are especially useful for the development of AI models. This is not a necessity at this stage of the project, but the resources will be allocated as soon as the requirement is precisely defined.

4.2.3 Hardware provisioning

The provisioning of the infrastructure consists of the following three aspects:

- **Hardware**: The CPU, RAM, disc, networking and other hardware components, allocated through Virtual Machines.
- **DNS**: The naming system for addressing machines in a human-friendly way. The EGI Dynamic DNS service[7] provides this functionality for LETHE.
- **TLS certificates**: The software artefacts containing cryptographic properties to enable secure network communication. Certificates must be issued by a third-party that is globally trusted, and the LETHE platform uses Let's Encrypt[8] for this.

Note that, strictly speaking, both DNS and TLS certificates are not part of the infrastructure, as they are not hardware components. Likewise, the installation of software, such as Operating Systems and cluster configurations, are not hardware components either. However, for convenience, they are managed in LETHE as part of the infrastructure provisioning.

The summary of the LETHE infrastructure is presented in Figure 3, which depicts the hardware resources allocated, according to the requirements (listed in Section 4.2.2) and classified by the different components of the platform (discussed in Section 4.1). Each VM is described in terms of CPU, RAM and storage.

Special encrypted volumes are attached to some VMs to protect user's data, as described in Section 4.2.1. Some machines have a DDNS flag, to represent machines that have an entry associated in the DNS and can be addressed by a fixed, human friendly name, such as *datalake.lethe.fedcloud.eu* instead of its more complicated and ephemeral network address.

Equally, machines with the TLS flag represent the machines that have obtained a TLS certificate, which normally means that they serve external users, so a secure network communication is required. Note that TLS certificates are typically (and cryptographically) linked to DNS names, so every machine containing the TLS flag should normally have the associated DDNS flag. The contrary is not true, as there can be machines that need to be addressed with a human-friendly name but for some reason do not need a specific TLS certificate to secure the network communication.

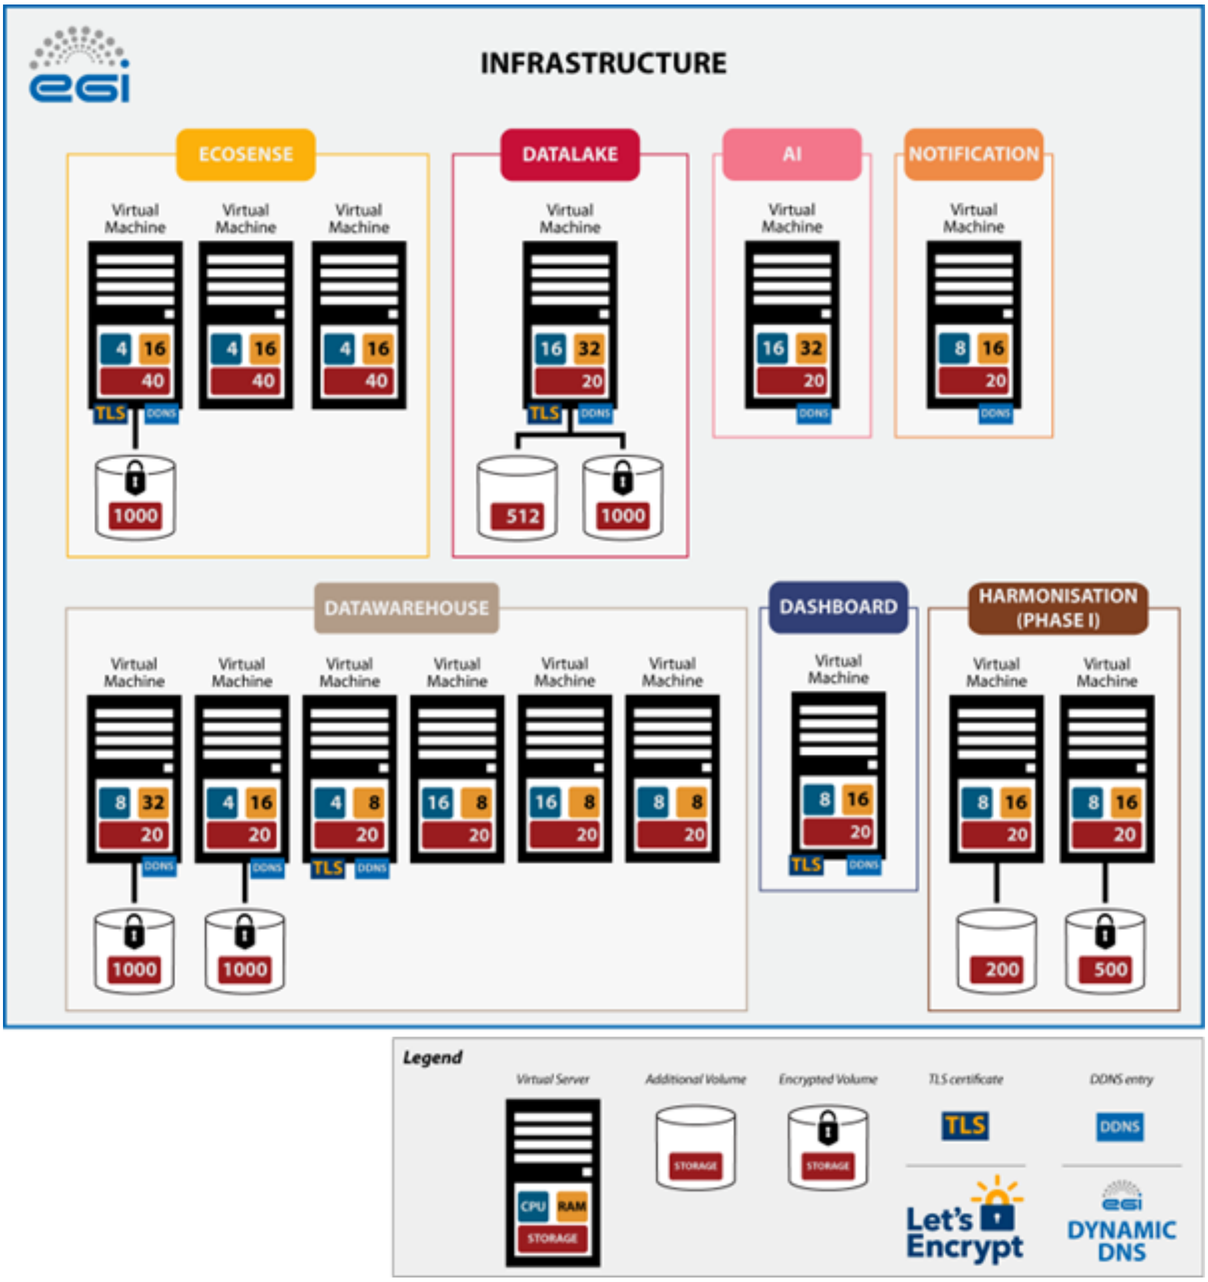

Figure 3. General infrastructure diagram (month 24 of the project)

Note that the LETHE infrastructure is not immutable but flexible and might scale up (increase resource allocation) and down (reduce resource allocation) in future, depending on the needs of the platform.

5 EOSC Services used in the LETHE infrastructure

Infrastructure components described in the previous section of this document based on EOSC services[1].

Following EOSC services are used in the LETHE infrastructure:

5.1 EGI Cloud Compute

With the EGI Cloud Compute[2] on-demand virtual machines have been deployed and scaled. The service offers guaranteed computational resources in a secure and isolated environment. In the LETHE the EGI Cloud Compute is used as a Phase II computing solution.

5.2 EGI DataHub

In the LETHE project The EGI DataHub[3] has been used to bring data to the secure storage and manage data sets on it. The EGI Datahub is used in the Phase I of the LETHE project.

5.3 EGI Online Storage

In the project we store data to the EGI Online Storage[4]. This service is integrated with the EGI Cloud Compute service i.e. Online Storage offers storage for virtual machines in Phase II solution.

5.4 EGI Accounting

EGI Accounting[5] tracks and reports the use of EGI services, offering insights and control over resource consumption. EGI Cloud computing entails the use of different supportive services, such as EGI Accounting.

5.5 EGI Check in

EGI Check-in[6] is a proxy service that operates as a central hub to connect federated Identity Providers (IdPs) with EGI service providers. Check-in allows users to select their preferred IdP so that they can access and use services in a uniform way. In the LETHE project Phase II authentication and authorisation solutions based on the EGI Check-In.

[1] https://marketplace.eosc-portal.eu/

[2] https://marketplace.eosc-portal.eu/services/egi-cloud-compute

[3] https://marketplace.eosc-portal.eu/services/egi-datahub

[4] https://marketplace.eosc-portal.eu/services/egi-online-storage

[5] https://www.egi.eu/service/accounting/

[6] https://marketplace.eosc-portal.eu/services/egi-check-in

5.6 EGI Notebooks

The EGi Notebooks[7] is a browser-based tool for interactive analysis of data using EGI Cloud Computing and EGI Online Storage. Notebooks are based on JupyterHub technology.

5.7 Dynamic DNS service

The Dynamic DNS[8] service provides a unified, federation-wide Dynamic DNS support for VMs in EGI infrastructure.

[7] https://marketplace.eosc-portal.eu/services/egi-notebooks

[8] https://marketplace.eosc-portal.eu/services/dynamic-dns-service

6 Open-source repositories of software

In this section links to open-source repositories of software used to address the project needs are listed. This means those open source repositories which contain software used in the LETHE infrastructure.

Applications such as data lake, data warehouse, dashboard and their integrations might contain other open source components which will be documented in other LETHE deliverables.

This section lists the base elements required to make the LETHE infrastructure ready to accommodate a series of dedicated applications developed within the project.

6.1 OpenStack

The LETHE infra components requiring CPU and storage resources are underpinned by OpenStack[9] [10], a free, open standard cloud computing platform, mostly deployed as infrastructure-as-a-service (IaaS) in both public and private clouds where virtual servers and other resources are made available to users. OpenStack is licensed under the *Apache License 2.0*, a permissive licence whose main conditions require preservation of copyright and licence notices.

6.2 Ubuntu

All virtual machines deployed with OpenStack are based on Ubuntu[11] [12], a popular Linux distro open-source operating system that has been adopted by numerous communities for cloud operations, bare metal provisioning, edge clusters or IoT. More specifically LETHE VMs are running Ubuntu LTS (Long Term Support) releases that are supported and maintained for longer periods of time than the standard editions.

6.3 Onedata

The dedicated storage used during both places of the project to keep data centrally available has been provided by Onedata[13] [14] technology, a global data management system, providing easy access to distributed storage resources. Onedata is licensed under the *Apache License 2.0*.

---

[9] https://www.openstack.org/

[10] https://releases.openstack.org/index.html

[11] https://ubuntu.com/

[12] https://ubuntu.com/download/server

[13] https://onedata.org

[14] https://github.com/onedata/onedata

6.4 Kubernetes

Some of the LETHE components are managed using Kubernetes[15] [16], an open-source container management system for automating software deployment, scaling and management. Kubernetes is licensed under *Apache License 2.0*.

Other applications deployed within LETHE infrastructure and supported by the above listed elements might also contain open source components but they are documented in other LETHE deliverables.

15 https://kubernetes.io

16 https://kubernetes.io/releases/download/

7 Conclusion

The LETHE infrastructure has been described in this document. In line with recent trends in computing, LETHE uses a virtual infrastructure supported by cloud technologies. Herein, technical concepts relevant to understanding the technologies used by the LETHE infrastructure have been briefly explained here.

The LETHE application as a whole is implemented on a platform consisting of different software components, which have been presented in this document. These components must run on the hardware devices that constitute the infrastructure, therefore the hardware requirements for each component have been listed, and the final resource allocation has been described. Due to the sensitive nature of user data, encrypted storage has been provided.

The ESOC services used for the LETHE infrastructure have been listed, together with the open-source software repositories used to meet the needs of the project.

As demonstrated during the project, the LETHE infrastructure is a robust and flexible allocation of hardware resources that can easily evolve with the needs of the project and can be managed as software code. It can be installed from scratch with minimal human intervention in a matter of minutes (or a couple of hours if the encrypted volumes need to be created for the first time) and can be dynamically adapted to match the current workload of the platform.

8 REFERENCES

[1] LETHE project (2022) **Project Dhttps://kubernetes.io/releases/download/eliverable 2.4**.

[2] Caballer M., Blanquer I., Moltó G., de Alfonso C. (2015) **Dynamic Management of Virtual Infrastructures**. J. Grid Comput., vol. 13, no. 1, pp. 53–70, Mar. 2015, doi: 10.1007/s10723-014-9296-5.

[3] OASIS (2018) **TOSCA Simple Profile in YAML Version 1.0** https://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/TOSCA-Simple-Profile-YAML-v1.0.html (accessed Jul. 28, 2022).

[4] Hohn A. (2022) **The Book of Kubernetes. A Complete Guide to Container Orchestration**. No Starch Press. ISBN-13: 9781718502642.

[5] Aumasson J-P. (2017) **Serious Cryptography. A Practical Introduction to Modern Encryption.** No Starch Press. ISBN:978-1-59327-826-7