Large-scale EXecution for Industry & Society

Deliverable D4.5

Definition of Mechanisms for Securing Federated Infrastructures

Co-funded by the Horizon 2020 Framework Programme of the European Union
Grant Agreement Number 825532
ICT-11-2018-2019 (IA - Innovation Action)

DELIVERABLE ID | TITLE	D4.5 | Definition of Mechanisms for Securing Federated Infrastructures
RESPONSIBLE AUTHOR	Frédéric Donnat (O24)
WORKPACKAGE ID | TITLE	WP4 | Orchestration and Secure Cloud/HPC Services Provisioning
WORKPACKAGE LEADER	LINKS
DATE OF DELIVERY (CONTRACTUAL)	31/03/2020 (M15)
DATE OF DELIVERY (SUBMITTED)	29/05/2020 (M17)
VERSION | STATUS	V1.0 | Final
TYPE OF DELIVERABLE	R (Report)
DISSEMINATION LEVEL	PU (Public)
AUTHORS (PARTNER)	O24; Atos; LINKS; IT4I; LRZ
INTERNAL REVIEW	Lukáš Vojáček (IT4I); Marc Levrier (Atos)

Project Coordinator: Dr. Jan Martinovič – IT4Innovations, VSB – Technical University of Ostrava
E-mail: jan.martinovic@vsb.cz, **Phone:** +420 597 329 598, **Web:** https://lexis-project.eu

DOCUMENT VERSION

VERSION	MODIFICATION(S)	DATE	AUTHOR(S)
0.1	Definition of Table of Contents and Partners´ involvement	14/12/2019	Frédéric Donnat (O24)
	Adding Introduction and Partners´ involvement	10/02/2020	Frédéric Donnat (O24)
	Added Section 4.1	17/02/2020	Laurent Ganne (Atos)
0.2	Proposed typo fixing in the introductions (Section 1), written section 2 introduction paragraph and (Section 2.1 about security-by-design)	21/02/2020	Marc Levrier (Atos)
0.3	Modified/added content to all sections where LRZ contribution was asked for.	02/03/2020	LRZ Team
0.4	Added/modified content to several sections. Added missing LRZ sections	31/03/2020	Frédéric Donnat (O24), Stephan Hachinger (LRZ)
0.5	Reviewed several sections and completed some parts. Added RBAC Matrix as ANNEX, updated both architecture and network diagrams. Adjusted Glossary, fixed minor typo.	06/04/2020	Frédéric Donnat (O24), Stephan Hachinger (LRZ)
0.6	Included internal reviewer feedback and comments. Fixed minor typo.	14/04/2020	Frédéric Donnat (O24)
0.7	Prepared public availability by removing sensitive information on HPC centre infrastructure.	14/04/2020	Frédéric Donnat (O24)
0.8	Extended document to explain new HPC centre integration (with small note on decommissioning).	27/04/2020	Frédéric Donnat (O24)
0.9	Internal review, Small additions/corrections proposed, Diagram review.	06/05/2020	Marc Levrier (Atos), Lukáš Vojáček (IT4I), Stephan Hachinger (LRZ)
0.91	Final updates based on review comments.	13/05/2020	Frédéric Donnat (O24)
1.0	Final check and small update of Section 4.2.	28/05/2020	Kateřina Slaninová (IT4I), Jan Martinovič (IT4I)

GLOSSARY

ACRONYM	DESCRIPTION
AAI	Authentication and Authorization Infrastructure
ABAC	Attribute-based Access Control
ACL	Access Control List
ALIEN4CLOUD OR A4C	Application LIfecycle ENablement for Cloud
API	Application Programming Interface
BSD	Berkeley Software Distribution
DDI	Distributed Data Infrastructure
DDOS	Distributed Denial Of Service
DMZ	DeMilitarized Zone
DNS	Domain Name System
DOS	Denial Of Service
FE	Front End
FQDN	Fully Qualified Domain Name
FTPS	File Transfer Protocol Secure
HPC	Hyper-Performance Computing
IT	Information Technology
ISO	International Organization for Standardization
HTTP	HyperText Transfer Protocol
HTTPS	HyperText Transfer Protocol Secure
MWN	Münchner Wissenschaftsnetz (Munich Scientific Network of LRZ and the Munich Universities)
NAT	Network Address Translation
NIC	Network Interface Controller
NTP	Network Time Protocol
OIDC	OpenID Connect
OWASP	Open Web Application Security Project
PRACE	Partnership for Advanced Computing in Europe
RBAC	Role-based Access Control
REST	Representational State Transfer
SAML	Security Assertion Markup Language

SIEM	Security Information and Event Management
SLURM	Simple Linux Utility for Resource Management
SMTPS	Simple Mail Transfer Protocol Secure
SSH	Secure SHell
SSL	Secure Socket Layer
SFTP	Secure File Transfer Protocol
TLS	Transport Layer Security
VLAN	Virtual Local Area Network
VM	Virtual Machine
VPN	Virtual Private Network
YORC	Ystia ORChestrator

TABLE OF PARTNERS

ACRONYM	PARTNER
Avio Aero	GE AVIO SRL
Atos	BULL SAS
AWI	ALFRED WEGENER INSTITUT HELMHOLTZ ZENTRUM FUR POLAR UND MEERESFORSCHUNG
BLABS	BAYNCORE LABS LIMITED
CEA	COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES
CIMA	Centro Internazionale in Monitoraggio Ambientale - Fondazione CIMA
CYC	CYCLOPS LABS GMBH
ECMWF	EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS
GFZ	HELMHOLTZ ZENTRUM POTSDAM DEUTSCHESGEOFORSCHUNGSZENTRUM GFZ
IT4I	VYSOKA SKOLA BANSKA - TECHNICKA UNIVERZITA OSTRAVA / IT4Innovations National Supercomputing Centre
ITHACA	ASSOCIAZIONE ITHACA
LINKS	FONDAZIONE LINKS / ISTITUTO SUPERIORE MARIO BOELLA ISMB
LRZ	BAYERISCHE AKADEMIE DER WISSENSCHAFTEN / Leibniz Rechenzentrum der BAdW
NUM	NUMTECH
O24	OUTPOST 24 FRANCE
TESEO	TESEO SPA TECNOLOGIE E SISTEMI ELETTRONICI ED OTTICI

TABLE OF CONTENTS

LIST OF FIGURES

EXECUTIVE SUMMARY

This document comprehensively describes the measures we are currently putting in place to make the federated infrastructure of LEXIS security. It covers aspects from the concrete design over the implementation to the deployment. With this, it follows the deliverable D4.1 [1], which focused on the first design decisions, including a selection of an actual Authentication and Authorisation system for the federated LEXIS infrastructure.

In order to achieve a high level of security, we followed the principles of "Security by Design" as described by well-known security organization in the security guide (NCSC [2, 3], Infosec Institute [4] and OWASP [5])":

- Minimising attack surface area: restricting access to selected features/applications,
- Following the principle of "Least Privileges": reducing privileges (user rights, access to compute/data resources, etc.) to the minimum needed to accomplish a task,
- Following the principle of "Defence in Depth": validating access in several layers (Portal, Orchestration, Computation, Data, ...) and adding auditing security control,
- Establishing secure defaults: putting in place default configurations guaranteeing a high security level,
- Not trusting services: all services will refer to the LEXIS AAI system, and not to another service, to validate authentication/authorisation,
- Keeping security simple: creating a clear and simple (not too complex) architecture to decrease risks due to misunderstandings or configuration errors,
- Putting in place proper separation of duties: respecting the security concepts of our computing and data centres and relying on clear responsibilities, we aim at preventing unintended cross-system access as well as abuse or attacks from insiders.

Position of the deliverable in the whole project context

The security of the LEXIS platform, which is a federated infrastructure spanning two HPC service providers and one Weather & Climate scientific data provider, is one of the main concerns of the LEXIS project. As a matter of fact, an abuse of the LEXIS platform from the outside, or even from within the project, can result in serious security issues such as denial of service (DoS or DDoS), data breach and cost issues (e.g. "stolen" computing time on HPC systems).

After the first focus on fundamental design decisions, and in particular on the federated authentication and authorisation system - the LEXIS AAI - in D4.1 [1], this deliverable will describe concrete security rules and mechanisms being implemented in the LEXIS infrastructure. It therefore complements, on the security side, the deliverables describing the mid-term infrastructure (D3.3 [6]) of LEXIS and its co-design (D2.3 [7]).

Description of the deliverable

D4.5 starts with an introduction (Section 1). Minimisation of attack surface, avoidance of trust in (non-AAI) services, least privilege principle and defence in depth are, altogether, topics of architectural security, as is the general Role-Based access to LEXIS services. These are discussed in Section 2, which closes with measures increasing availability and resilience. Separation of Duties is not only a concern of the LEXIS platform, but a concept which largely already guides the internal architecture of the LEXIS computing and data centres. The integration of the LEXIS security concept and those of the centres, with appropriate Separation of Duties is therefore extensively discussed in Section 3. Section 4, finally, discusses security of the different systems within the LEXIS platform (YORC and HEAppE, DDI, Portal). It clearly focuses on software issues such as secure defaults and encryption by default. Section 5 summarises our results.

Contributors to the deliverable content are:

- O24 as leader of activities regarding security aspects, responsible for this deliverable,
- IT4I both as HPC centre in charge of deploying and administrating LEXIS infrastructure (connecting to HPC systems), and as provider of HEAppE software in charge of mapping LEXIS infrastructure security policies to the existing HPC (and Cloud-Computing) infrastructure,
- LRZ both as HPC centre (analogous to IT4I), and as WP3 leader in charge of the LEXIS distributed data infrastructure (DDI),
- Atos both as WP2 leader in charge of co-design activities, and technology provider (for YORC and Alien4Cloud software that are used for LEXIS orchestration),
- LINKS as WP4 leader in charge of the LEXIS Orchestration,
- CYC as WP8 leader in charge of the LEXIS Portal where it is crucial to implement security recommendations and to connect to the LEXIS AAI.

1 INTRODUCTION

From the very beginning, the entire LEXIS collaboration has aimed at creating and providing a secure platform, and at leveraging the security of the High-Performance-Computing (HPC) centres instead of compromising it. The co-design and policy-centric work packages (WP2, WP4) have raised awareness for the LEXIS security principles and tackled architectural challenges. The present deliverable D4.5 describes the results of this process.

During the earlier co-design phase, one very important focus has been on the LEXIS AAI as a central component for a secure infrastructure. However, also, general LEXIS security and service-quality principles have been devised (cf. executive summary):

- Minimizing attack surface area,
- Following the principle of “Least Privileges”,
- Following the principle of “Defence in Depth”,
- Ensuring “Secure Defaults”,
- Following the model of “No Trust”,
- Keeping security simple,
- Putting in place “Separation of Duties”,
- Ensuring high availability of services.

Already, in the deliverable D4.1 “Analysis of mechanisms for securing federated infrastructures” [1], these principles are reflected e.g. in a deployment diagram for LEXIS AAI software that shows:

- network isolation: deployment of web proxies (with load balancing and web filtering capabilities) in front of the LEXIS AAI and dedicated servers for the LEXIS AAI (Keycloak, Database, etc.); and
- high availability: overall high available deployment of LEXIS AAI with high available deployment of each of its component (Keycloak, Database, etc.).

During the further co-design up to M15, the architecture of the LEXIS platform has been designed, refined, and implemented taking into consideration the security principles as discussed above. This has resulted in the following measures:

- The LEXIS architecture has been divided into clear “LEXIS Layers” (LEXIS Portal; LEXIS Service layer including LEXIS AAI, LEXIS Portal BackEnd; LEXIS DDI; LEXIS Orchestration; and LEXIS Infrastructure layer). Connectivity between different layers, as also to the internet, is restricted, so as to prevent direct exposure to internet of all the “LEXIS Services”, thus allowing to keep security simple and reducing the attack surface area. Some LEXIS Services will also be deployed behind proxies in order to add an additional security layer for protecting them by adding filtering capabilities (besides load balancing).
- The LEXIS RBAC Matrix that is currently implemented in LEXIS AAI is implementing the principle of “Least Privileges“ by devising three different types of access (list, read and write), according to the need to accomplish a task. Those access types have already been detailed in [1].
- The LEXIS architecture is based on the model of “No Trust”: LEXIS Services do not trust one another. Instead, each LEXIS Service refers to the authentication and authorisation via the LEXIS AAI and its JWT tokens. As an example, a request of the LEXIS Portal to a LEXIS Service (with a JWT token issued from LEXIS AAI and some resource request) triggers a validation of authentication and authorisation by the LEXIS Service before an access to the specific resource is granted.
- The LEXIS infrastructure layer is isolated from the computing centres’ core infrastructure in order to provide proper separation of duties. The HEAppE middleware provided by IT4I is the key component that allows for an interaction between those components and for preventing LEXIS users from accessing lower-level infrastructure in an abusive manner. On the other hand, any interaction of LEXIS Services among each other is controlled by the LEXIS AAI. Therefore, any access to LEXIS resources and services (be it interactive or automatic) involves a validation of LEXIS AAI information (authentication and authorisation), preventing abuse.

All the security mechanisms and controls put in place in LEXIS infrastructure will be detailed in the following sections. Brief security-oriented guidelines on integrating/decommissioning further HPC and Cloud-Computing centres with the LEXIS platform are part of this document as well. These also include measures for the case of centres leaving the platform.

2 ARCHITECTURE SECURITY

An architectural security diagram (Figure 1) describes the three different security zones of the LEXIS platform. LEXIS services are deployed in these security zones according to their requirements in terms of security. All LEXIS services interact with the LEXIS AAI:

- The "LEXIS DeMilitarized Zone (DMZ)" layer contains all components that need direct access to the internet (basically the LEXIS Portal front end), and security equipment/services that are required to:
  - provide a safe entry point to all LEXIS Services (LEXIS AAI, LEXIS Portal back end, LEXIS DDI, LEXIS Orchestrator) - e.g. a Reverse Proxy for protecting Web Applications and APIs,
  - inter-connect HPC supercomputing centres such as a VPN Gateway.
- The "LEXIS Trusted Zone" is the second security layer for LEXIS Services implementing the actual LEXIS functional services which are mostly composed of main LEXIS core components:
  - LEXIS AAI for identity and access management (IAM) providing authentication and authorisation to LEXIS Services,
  - LEXIS Portal back end for providing all business logic to the LEXIS Portal through LEXIS Portal Front End (deployed in "LEXIS DMZ"),
  - LEXIS DDI for (distributed) data management,
  - LEXIS Orchestrator for orchestration of workflows composed of combination of computing, visualisation, and data transfer tasks.
- The "HPC/Cloud Infrastructure" is the last security layer corresponding to both cloud and HPC infrastructures and the related back-end services:
  - site-specific HPC security services such as authentication and authorisation system,
  - HEAppE security middleware bridging cloud and HPC identities and access, and abstracting HPC resource management.
  - approval system for allowing access to HPC/Cloud Infrastructure resources (mainly compute and storage ones),
  - accounting and billing system.

In addition to the zones defined in the LEXIS architecture and as a direct consequence of the "No Trust" principle, LEXIS Services are all untrusted by default and, as such, must either be directly registered in the LEXIS AAI (as Keycloak "Client") or interact with the HEAppE middleware. In the latter case, HEAppE implements a secure bridge between the service or user accessing it and the HPC/Cloud resources which are not directly accessible by the end users.

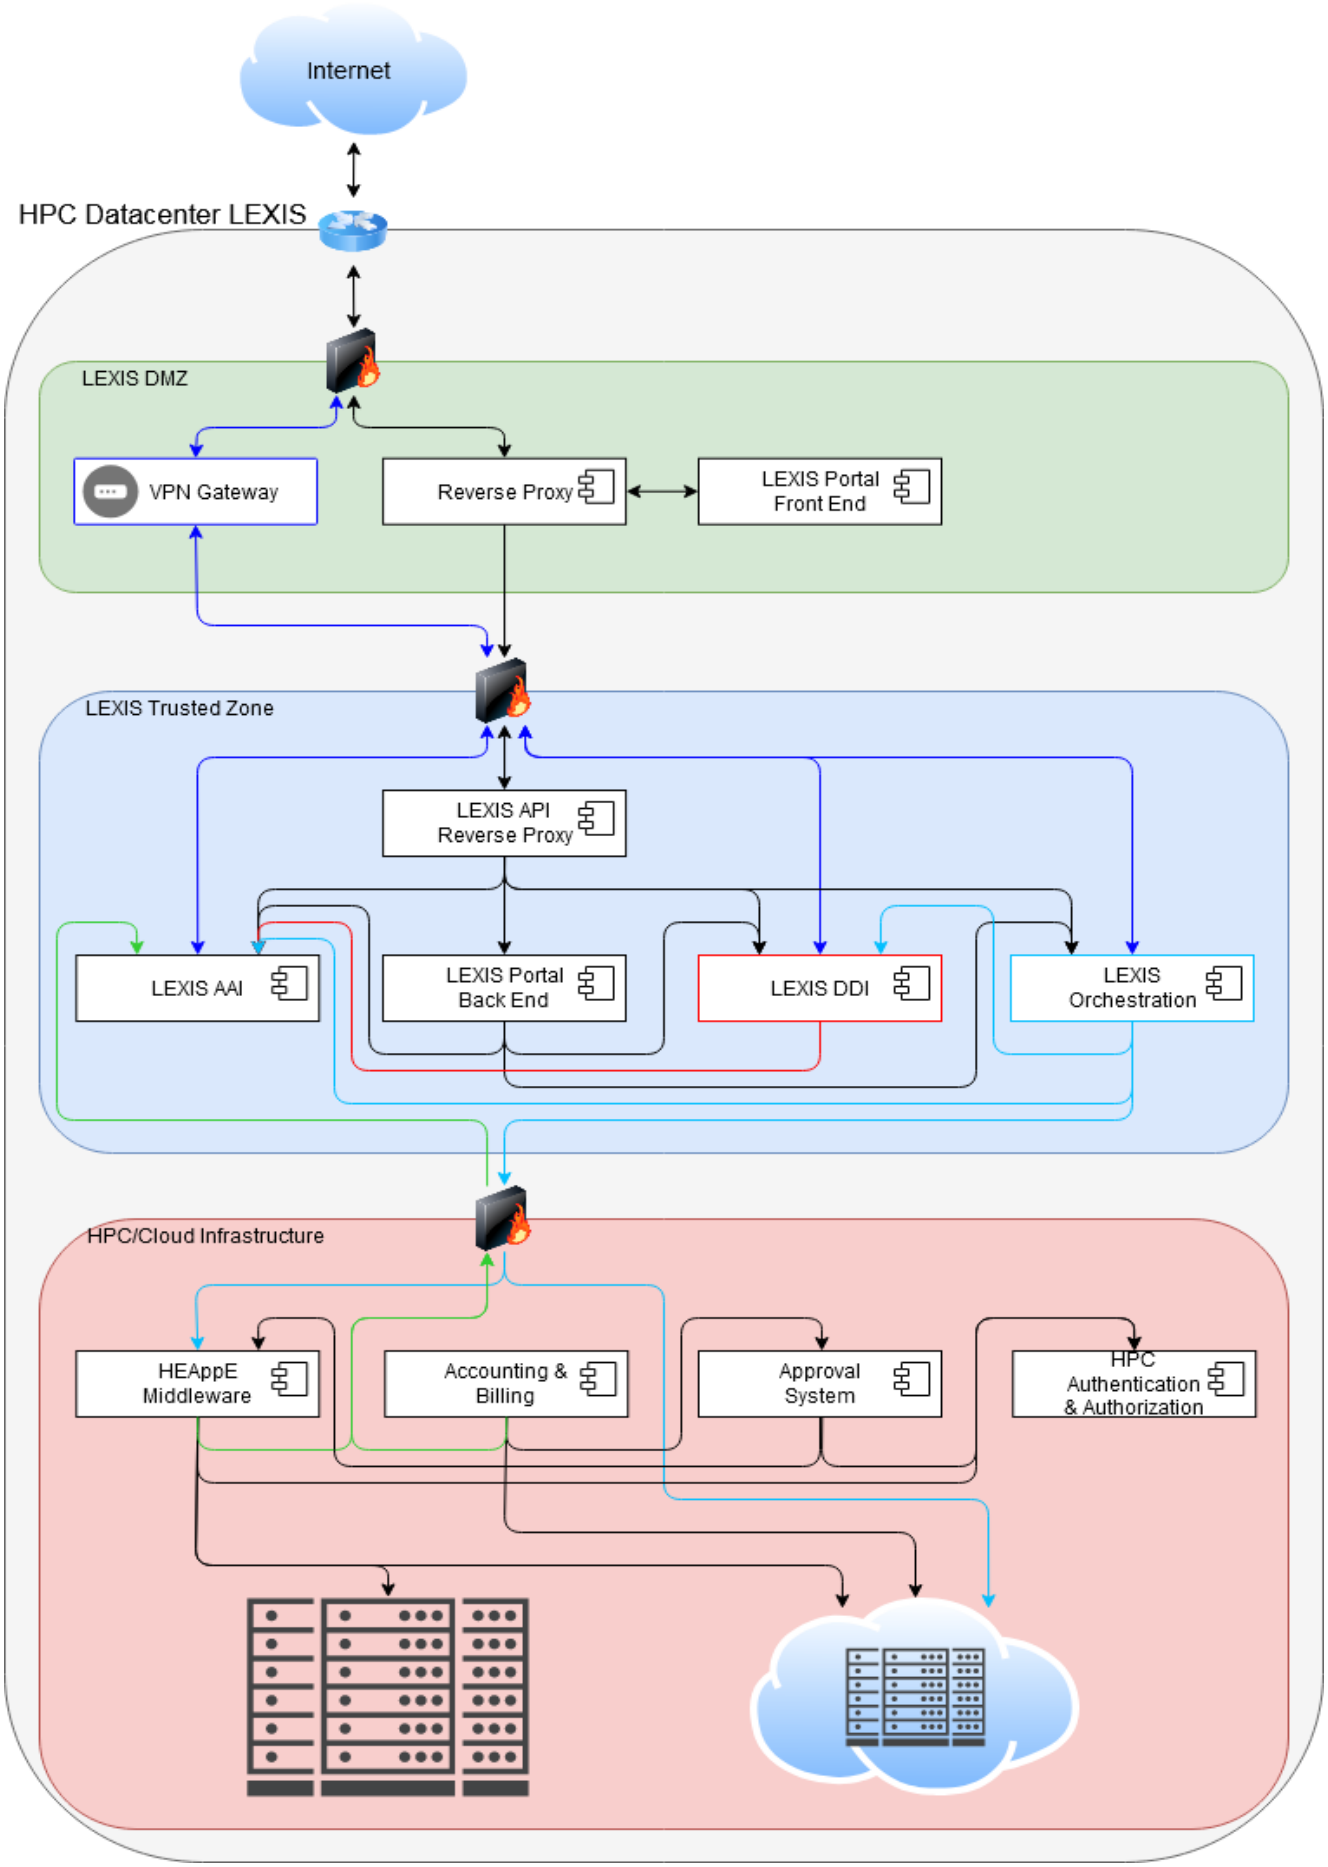

Figure 1 LEXIS High Level Security Architecture Diagram

2.1 AAI SYSTEM AND LEXIS COMPONENTS

From the first technical face-to-face meeting in M2, security has been considered just as important in LEXIS as computing or data-management related features. This means that security is not a concept to be added to the main LEXIS platform functions. Security has instead been taken into account in building the platform from the very beginning. Each and every LEXIS functional service has thus been designed to:

- Be able to delegate the user identity and access management (authentication and authorisation responsibility) to a single system of reference (the LEXIS AAI),
- Have this LEXIS AAI rely on state-of-the-art standards and implement best practises such as role-based and token-based security transactions,
- Work around differences in the security standards among services (typically OAuth2/OpenIDC for LEXIS Portal and LEXIS DDI and SAML for LEXIS Orchestration),
- Work around differences in the security models among LEXIS HPC service providers.

The LEXIS AAI is a federated system with instances at IT4I and LRZ (as well as further Computing Centres prospectively joining LEXIS). It is exclusively accepted and used by LEXIS services in order to verify LEXIS users' identities and obtain authorisation for each access. This clear decision for a central IAM system is the prerequisite for an appropriate design ("Security by Design") of LEXIS platform. It prevents later road blocks and potential security problems caused by divergent developments and consequential architectural instabilities.

As briefly described in the Introduction, the LEXIS infrastructure is based on the principles that LEXIS Services do not trust each other, but rely on LEXIS AAI for authenticating an identity (a user, a process, etc.) and granting access to a resource (compute, data, etc.). The concept of "No Trust" has been taken into account in setting up each LEXIS service.

The modular LEXIS infrastructure design allows to put in place a proper separation of duties, as each LEXIS component is responsible for handling only one service. All components refer to the LEXIS AAI in order to authenticate incoming requests and grant any access to any resource. For instance, the LEXIS DDI will not trust any request from the LEXIS Portal without ensuring that it contains proper authentication and authorisation from the LEXIS AAI, which is accessed here through OpenIDC (based on the OAuth2.0 protocol and using JWT Tokens).

Last but not least, all LEXIS infrastructure is designed with auditing capabilities as per defence in depth principle. As for Keycloak as central LEXIS AAI component in particular, any action can be recorded for further analysis by LEXIS administrators. The auditing capabilities will be used to feed a centralised "Security Event and Information Management" (SIEM) system. With such a system, also correlated auditing events from different parts of the infrastructure can be made use of in order to detect not only attempts to abuse the LEXIS infrastructure, but also misconfiguration leading to security breach.

2.2 RBAC MATRIX

As a result of the co-design phase, authorisations are performed using an RBAC Matrix, keeping in mind both the following aspects:

- Keeping the security simple: only three different level access of any LEXIS Service (list, read, write),
- Respecting the principle of "Least Privileges": only granting access to resource or process to an identity when it is mandatory.

The LEXIS RBAC matrix considers all accesses to resources required by an identity with a fine-grained approach. During the implementation phase, a relatively simple version of the RBAC Matrix has been favoured over alternatives, increasing clarity and speed of implementation from the original RBAC Matrix in order to make things easier and to speed up the implementation. The approach chosen to build the RBAC Matrix (taking into consideration all LEXIS services with all specific use cases), allows to easily add more fine-grained permission if needed later on.

In a few words the LEXIS RBAC Matrix is now taking into consideration five main roles as follows:

- *LEXIS Administrator*: This role allows a LEXIS Administrator to have full management over all Organisations including Identity and Access, Billing, Licensing, Project, Orchestration and Data without giving any access to the resource themselves (including Orchestration information or Data input or output) as they may contain sensitive data from customer perspective. It also implies that the role is not given permission to elevate his privileges to gain access to sensitive data.
- *LEXIS Support*: This role allows a LEXIS Support member to have read access over some Organisations including Identity and Access, Billing, Licensing, Project, Orchestration and Data without giving access to the sensitive resource themselves (same motivation as for LEXIS Administrator role above).
- *LEXIS Organisation Manager*: This role allows a LEXIS User to fully manage an Organisation including Identity and Access, Billing, Licensing (excluding elevating his privileges to gain access to any resource).
- *LEXIS Project Manager*: This role allows a LEXIS User to fully manage a LEXIS Project within an Organisation including read access to Identity and Access, Billing, Licensing, full access to LEXIS Project within the Organisation and read access to Orchestration and Data without giving read access to the sensitive resource themselves (including Orchestration information or Data input or output). The access is restricted by granting access on specific LEXIS Project Organisation.
- *LEXIS User*: This role allows a LEXIS User to fully manage LEXIS Project within an Organization including Project, Orchestration and Data (including access to the resources themselves: Orchestration information or Data input or output). The access is restricted by granting access on specific LEXIS Project within the Organizations.

For a more better visualization and representation of the RBAC Matrix, please refer to the RBAC MATRIX spreadsheet in the ANNEX.

2.3 HIGH AVAILABILITY & RESILIENCE

The LEXIS platform has been designed with high availability and resilience in mind. For instance, we are aiming at deploying each LEXIS Service of LEXIS platform on both HPC Supercomputing centres (IT4I and LRZ) with a high availability or a cluster mode.

This first stage deployment will provide a solid base for the LEXIS platform allowing to extend it with additional computing centre (HPC Supercomputing centres or Cloud platform). Taking into account several constraints (such as security and regulatory policies, qualification process for the additional HPC centre including accounting system, level of autonomy, time constraints, etc.), two levels of integration can be considered:

- **Full integration**: The additional HPC computing centre will need to deploy and administer all LEXIS Services, as described below:
  - *LEXIS Portal*: Further investigation will be needed in order to choose the best deployment mode between active/active or active/passive deployment in term of performance and user interface experience for customer, but there is no blocking point in terms of security.
  - *LEXIS AAI deployment*: An active/passive mode is envisioned as high availability will already be guaranteed by the LEXIS AAI deployed in IT4I and LRZ supercomputing centres. A fully distributed mode between more than two computing centres is very challenging (as a matter of fact, even the Keycloak project does not mentioned such deployment). Even if a solution with an active/passive database cluster and an active/passive Keycloak deployment with load balancer is a valid and secure technical solution, further investigations are required to ensure optimal performance and administration/maintainability of such deployment.
  - *LEXIS DDI deployment*: All LEXIS DDI components have to be deployed in order to create a new "iRODS Zone" that will be federated with existing ones in IT4I and LRZ supercomputing centres.
  - *LEXIS Orchestrator deployment*: A YORC cluster will be deployed and interconnected with the existing YORC clusters already deployed in both T4I and LRZ supercomputing centres.
  - *HEAppE middleware deployment*: HEAppE will be deployed in the additional computing centre in order to allow proper mapping between internal HPC accounts and LEXIS accounts.

- **Client integration**: The additional computing centre will only need to deploy and administer LEXIS DDI components and the HEAppE middleware and allow access to the on-site HPC/Cloud infrastructure as well.

In both cases, the additional computing centre will need to provide full connectivity to both IT4I and LRZ through a site-to-site VPN.

The LEXIS AAI will be deployed in "Cross Datacenter Replication Mode", which allows to deploy a cluster version of Keycloak among two different computing/data centres. On top of this deployment, one web proxy with load balancing feature will be deployed in each HPC supercomputing centre, allowing to redirect any request to the available Keycloak service (local or remote one). This deployment allows to respect "Fail Securely" principle as this will prevent from leaking information on the LEXIS AAI (such as internal IP address or DNS name by using high availability, and redirecting database or configuration error by using web reverse proxy filtering) in case of connectivity issue to local Keycloak.

Figure 2 shows our LEXIS AAI setup in detail. In case the geographically distributed active/active deployment would prove unstable or slow in heavy "real-world" usage, an active/passive deployment could be used instead.

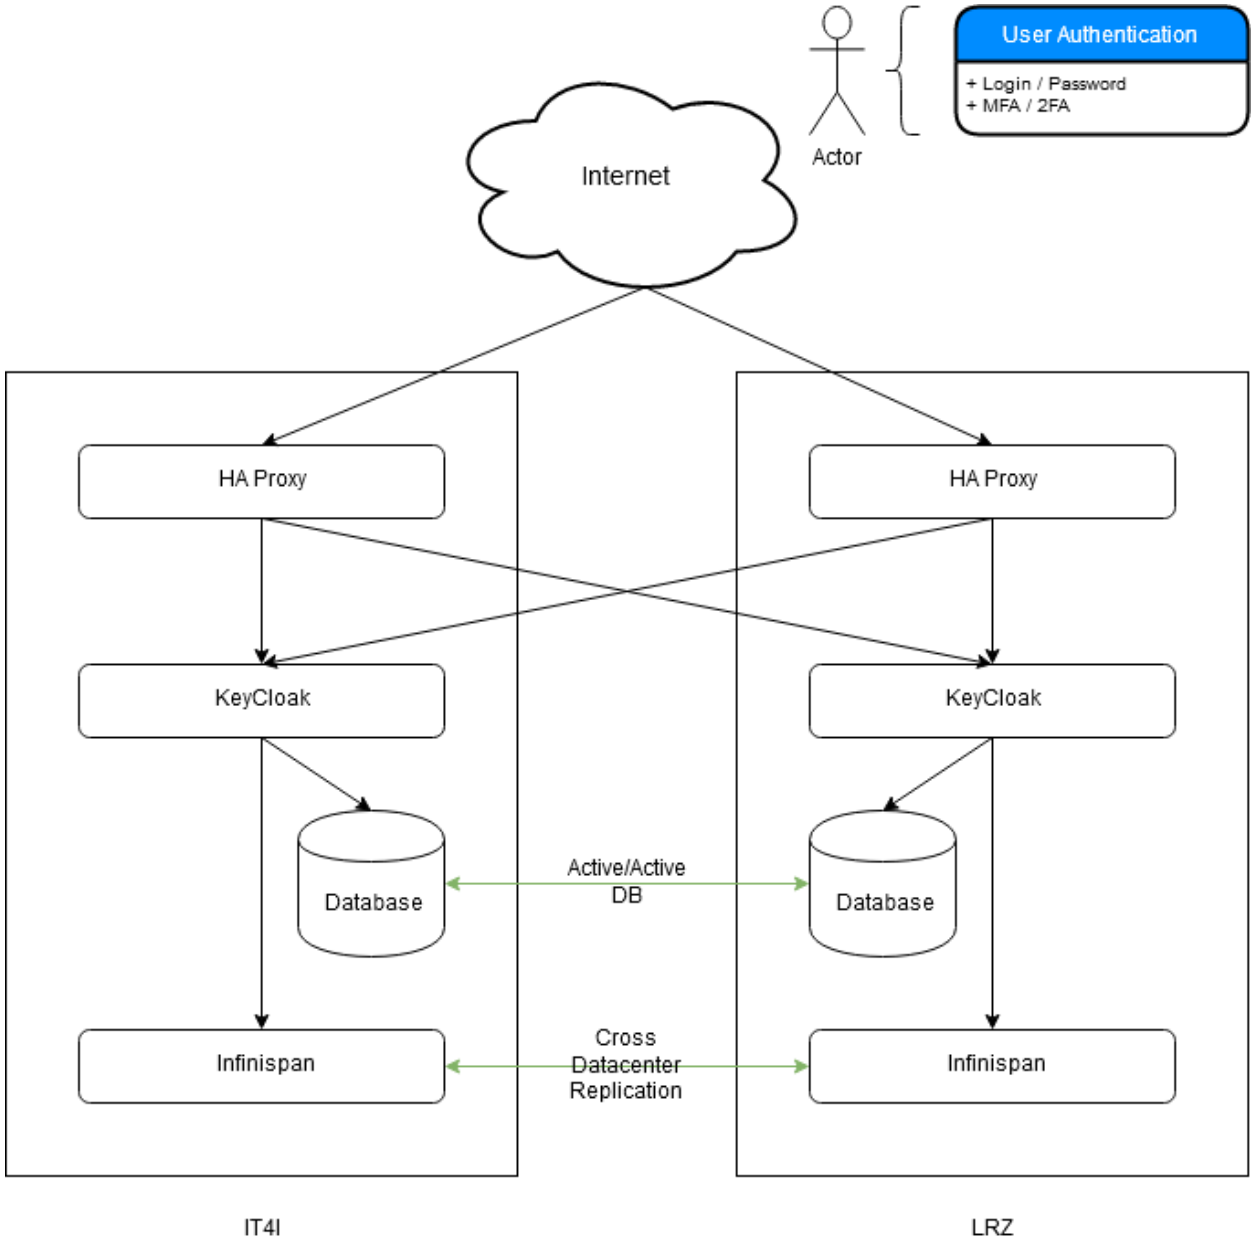

Figure 2 LEXIS AAI High Availability Deployment Diagram

The LEXIS Orchestration layer relaying on YORC is deployed redundantly with one instance per HPC Supercomputing centre (IT4I/LRZ). Additional computing centres included into LEXIS may not be obliged to add a further instance.

The LEXIS DDI relying on iRODS is deployed with two independent zones at IT4I and LRZ (with analogous plans for possible additional LEXIS centres). This splits the iRODS metadata catalogue of the LEXIS DDI (iCAT, comprising iRODS "file-system" and descriptive metadata), decreasing the chance of a complete failure. Each zones' central iCAT server is deployed with a redundant setup (high-availability, redundant virtual machines, or load-balanced, duplicated/synchronised service setup). A research publication on our iRODS resiliency concept, which follows earlier approaches, is currently being prepared [8]. When computing/data centres join or leave the LEXIS platform, a new zone is simply added/removed from the federation. However, leaving LEXIS also implies a separation of research data to be kept at the centre, to be kept in LEXIS and to be deleted, as well as prohibiting further access to the platform. To this purpose, a controlled decommissioning process will be prepared.

3 SEPARATION OF DUTIES

LEXIS provides an easy-to-use HPC/Cloud Computing service for industry and society. This has the objective of enlarging the community which uses computing services to boost their Research and Development.

In order to ensure secure usage in such a demanding context, a clear separation of duties is needed to simplify a sufficient "hardening" of each part of LEXIS services. The LEXIS platform is thus built on top of existing HPC Supercomputing centres services, reusing their intrinsic security layer, adding its own security layer in order to enhance the overall security. A similar relation will ensue between our LEXIS platform security and the "native" security measures in further computing centres who like to join LEXIS. For instance, the computing centres inter-connexion will be achieved using site-to-site VPN from additional computing centre to both IT4I and LRZ HPC Supercomputing centres.

In Figure 3, we detail security measures within LEXIS and "native" infrastructure at IT4I and LRZ.

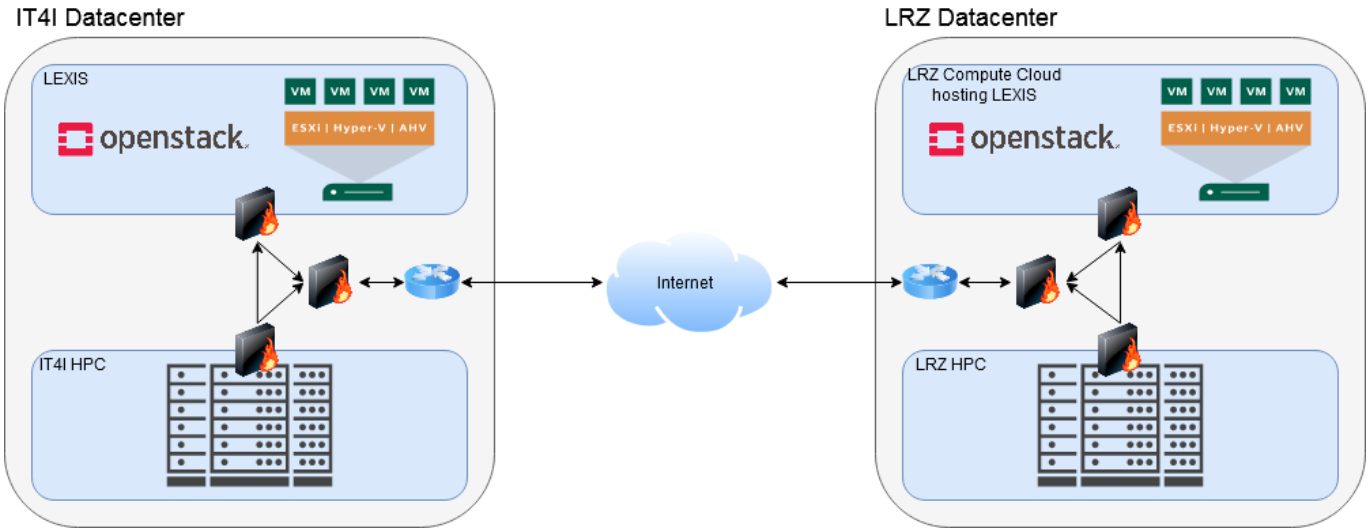

Figure 3 LEXIS Deployment on both IT4I and LRZ Supercomputing Centres

3.1 NETWORK ISOLATION [STATUS QUO]

The LEXIS project uses existing networks at IT4I and LRZ, as well as newly added or configured virtual or physical networks. We describe some particularly important aspects of this below, in separate sections for LRZ and IT4I. Figure 3 sketches the relation between internal firewalling in each computing/data centre and the interconnectivity between the centres (details see Section 3.2) over a firewall/routing endpoint at each site.

3.1.1 At IT4I

At IT4I, the experimental LEXIS infrastructure and production HPC systems (under ISO 27001 certification) are separated on the administrative level and on the network level. The experimental infrastructure shares only boundary elements of the physical network with the production HPC landscape.

Experimental LEXIS Infrastructure

The experimental systems are separated into their own private network subnet accessible only via VPN. Connection to the WAN is realised via boundary network elements of the IT4I computing/data centre, and a public IP subnet is served through a dedicated Virtual Local Area Network (VLAN). The infrastructure is operated by a team of researchers assigned to the LEXIS project and thanks to its separation from the production systems it does not have to be under supervision of the IT4I Supercomputing Services (SCS). Internally, the network of the experimental infrastructure is further divided to several VLANs. One VLAN is used for infrastructure services (Object Storage endpoints, Directory Services, VPN servers, etc.). Servers which are part of the LEXIS platform (YORC, iRODS, etc.) are separated into their own VLAN. Project partners which participate in development of the platform can access the platform servers through the dedicated VPN access. Secure remote administration consoles are accessible only in this internal VLAN and through the VPN. Some servers (e.g. iRODS) require public IP addresses, which are available in their own VLAN and only selected ports are opened on the servers' firewalls.

HPC Systems

Connection to the HPC systems' fabric is realised through a set of dedicated gateway nodes, which have NICs connected to both the cluster InfiniBand network and the LEXIS infrastructure Ethernet. The gateway nodes are operated by the IT4I operations team (SCS) and are not directly accessible by the LEXIS team. This separation of concerns allows us to maintain a very high flexibility in terms of possible configuration without disrupting main IT4I operations or create security vulnerabilities.

3.1.2 At LRZ

At LRZ, the Compute Cloud and Virtualisation infrastructures for running virtual machines are a part of the production systems, as are the high-performance computing systems (Linux Cluster, SuperMUC-NG). The security concept of those systems is documented in the LRZ IT Security Management System (ISO 27001 compliant ISMS part of the LRZ I/SMS). In terms of network connectivity, virtual machines (which are mainly used to run persistent and higher-availability LEXIS services) and the more experimental systems (Burst Buffers, Experimental Storage System) are partially under the control of the LRZ LEXIS Team: they can, for example be flexibly put into different or also newly-allocated VLANs on request. In the following, we first give a basic overview over Compute-Cloud, Linux-Cluster and SuperMUC-NG network connectivity, and then describe the connectivity for VMs and experimental systems used in the LEXIS infrastructure.

Compute Cloud

The LRZ Compute Cloud is used to dynamically run and stop virtual machines via a web interface or through the OpenStack API. It has been in production for a few years, and is managed by the ITS group, which also controls the internal network setup of the machines, minimizing the attack surface. For reachability from the outside, the user can (does not have to) "attach an IP address" to a VM, which is - according to his/her choice - routed either within the Munich Scientific Network of LRZ & universities (MWN), or world-wide. Within OpenStack, the VMs are reached via a NAT- and firewalling-like mechanism in OpenStack. This allows the user (whether an LRZ project member or a partner) to set security groups for a specific machine, where the reachability e.g. of ports is defined.

Linux Cluster

The LRZ Linux Cluster has login nodes which are reachable from the outside through security equipment such as firewall with relatively strict limitation of connection attempts per unit time. It is a production system administrated by LRZ's HPC group.

SuperMUC-NG

The top-level HPC machine at LRZ, SuperMUC-NG, is managed by LRZ's HPC group as well. In line with computational power and possible abuse potential, this machine is strongly secured. The login nodes have public IP addresses, but only port 22 is open (incoming) to fixed IPs which have to be whitelisted by the user (all login nodes have the second interface to the PRACE internal network between European HPC machines, which is somewhat less restrictive). User accounting follows a procedure based on a strong and competitive scientific and administrative eligibility check. No outgoing traffic initiation from the login nodes is allowed. It is further made impossible to connect from login nodes to compute nodes via the OmniPath inter-node connection; only the SLURM system can connect to the compute nodes.

Virtualisation system

All Virtual Servers which are used for the Infrastructure of the LEXIS platform are put in a separate VLAN, routed through a BSD Packet Filter based "Virtual Firewall", which LRZ offers as a service. The machines used in for LEXIS infrastructure have local firewall based on "netfilter/iptables" settings as strict as possible. Specific ports on the machines related to some applications can be opened and accessed by partners. Settings of the Virtual Firewall and the internal firewall (based on netfilter/iptables) are administrated by the LRZ LEXIS team in this section.

Burst buffers and experimental storage system

At the time of writing, the burst buffers are in a preliminary setup, in an MWN-routable VLAN. This provides some reduction of the attack surface, as access is only possible from within MWN or via VPN. Iptables on the buffers are configured as restrictive as possible. The experimental storage system right now has an internet-routable IP address and a "netfilter/iptables" based local firewall. As soon as this system is used for somewhat sensitive data or services, this setup will be further hardened.

3.2 SECURE LEXIS NETWORK ARCHITECTURE

The complete LEXIS network architecture, interconnecting the different participants, will rely on Network access restrictions following a concept of "network security zones".

As shown in Figure 4, DMZ will be set up to provide LEXIS services to internet and protect them from unwanted access (basically allowing HTTPS and FTPS/SFTP incoming traffic and blocking outgoing internet traffic except HTTPS, FTPS/SFTP, DNS, NTP, SMTPS). This zone will contain some web proxies with web filtering capabilities that could be enriched at a later point in order to properly filter incoming internet traffic, audit and block "malicious" or suspicious incoming traffic. Basically, the traffic can be monitored for improper HTTPS requests (incorrect protocol, SQL injection, XSS, etc.) as described by OWASP in their "Top 10 Web Application Security Risks" [9, 10].

This DMZ will also be the entry point for the VPN interconnection of IT4I and LRZ for any non-internet facing traffic (such as any traffic between the servers of the LEXIS AAI, Orchestrator and DDI systems). As reflected in Figure 4, the VPN service will be isolated from other LEXIS Services in the DMZ.

A "Trusted Zone" will also be set up in order to deploy LEXIS Services and allow communication from LEXIS Services (LEXIS Portal back end, LEXIS Orchestrator, and LEXIS DDI) to the LEXIS AAI wherever appropriate. An additional local DNS server will be deployed at each site in order to ensure redundant domain name resolution, such that our highly available services are always reachable under their FQDN. In order to properly ensure secure communication (using SSL/TLS or HTTPS, or any other secure protocol), an NTP synchronisation will be put in place. The firewall

protecting the “Trusted Zone” will be configured accordingly to the firewall protecting the “DMZ”, allowing incoming and outgoing traffic from “DMZ” (basically HTTPS, FTPS/SFTP, DNS and VPN) to the dedicated server.

The VPN connection between the HPC Supercomputing centres, through which all traffic between the respective “Trusted Zones” must be routed, may introduce some latency for synchronizing LEXIS services (LEXIS AAI, LEXIS Portal BackEnd, LEXIS DDI, LEXIS Orchestrator). If any impact on performance (due to any overhead introduced by the VPN) is observed, we may put in place counter measures such as:

- Link aggregation with Network Bonding (aggregating bandwidth across multiple physical links by combining several network interfaces together into one single interface in order to improve performance),
- Traffic identification and redirection (identifying, e.g., LEXIS DDI traffic dealing with “public data” that can be transmitted over the internet without a secure communication channel).

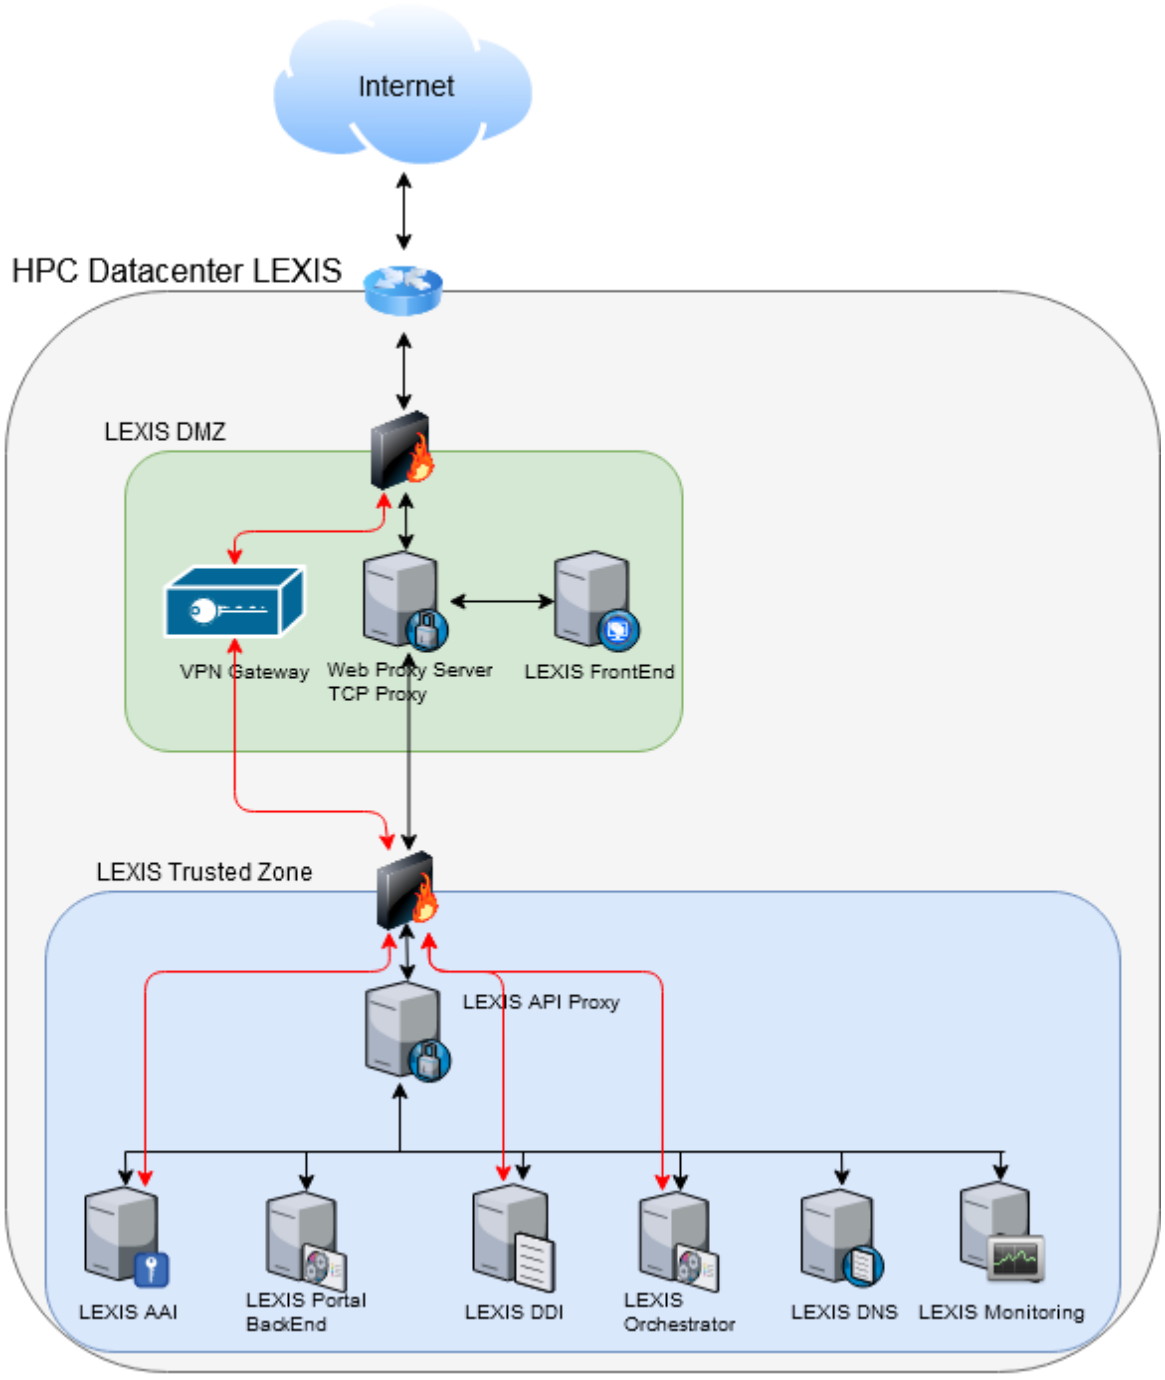

Figure 4 LEXIS Network Security Zones

3.3 INTERCONNECTION BETWEEN LEXIS COMPUTING CENTRES

The interconnection between LEXIS Supercomputing centres, initially based on white listing the respective IPs in machine or dedicated firewalls, will soon be based on the VPN network bridge discussed. It is managed by the LEXIS teams of IT4I and LRZ. With the VPN, sniffable traffic on the internet would be avoided even if insecure protocols were used. Also, the use of internet routable IP addresses (cf. Section 3.1) will be gradually minimised.

In case of additional HPC Supercomputing centres joining the LEXIS platform, a site-to-site VPN capability is mandatory. Performance tests will be required in order to decide whether all traffic can be routed over a VPN connection to one of the existing centres or multiple connections have to be set up.

4 SOFTWARE SECURITY

Finally, any piece of software used to set up the LEXIS service can potentially bear security risks. This can be due to intrinsic vulnerabilities, design flaws or a misconfiguration. This section of the present deliverable deals with this aspect in general, and then gives some more details on security aspects of the orchestration system, the LEXIS DDI and the LEXIS portal as central parts of the infrastructure.

Any software used in LEXIS, be it included within a framework or just a single piece of code, must closely follow security standards as already described in the introduction:

- The attack surface area must be minimised; for instance, no unnecessary service/feature should be exposed to any user.
- Default configurations should be conservative and safe, for instance, minimal privilege should be first granted to any user and privilege elevation should be audited. The password policy must also be strong. If software comes with unsafe defaults, an automatised setup with safe defaults will be taken care of within LEXIS, e.g. via appropriate ansible scripts.
- Configuration and functionality must respect the principle of least privileges. A user should be granted access to any service (software part, hardware, process, job, task, daemon) only as far as needed to accomplish his task.
- Configuration and software have to ensure secure failure. If a process or a transaction cannot be performed successfully, this must not result in additional privileges or the provision of additional data (e.g. sensitive logs) to the user.

In order to ensure the security of the whole LEXIS platform, regular security assessments will be conducted by O24 on LEXIS software components. These will follow best practices in software security, such as the ones described by OWASP for Web Applications and APIs [9, 10].

4.1 ORCHESTRATOR (YORC) SECURITY

The orchestration system allows to run jobs on LEXIS HPC infrastructures, it allocates resources, then deploys software on LEXIS Cloud infrastructures, and finally runs workflows - sequences of operations that can be executed on these platforms. Its safe operation is thus a central prerequisite for the LEXIS platform being secure. Its core components are:

- Front end, Alien4Cloud.
- Back end, YORC (Ystia Orchestrator).

The Alien4Cloud front end requires user authentication. In order to perform this as safe as possible, it is configured to delegate its user authentication to the LEXIS AAI, i.e. Keycloak. In our Keycloak configuration, a new “Keycloak Client” is dedicated to Alien4Cloud. This client is using the SAML 2 protocol, as this is currently supported by A4C.

Role mapping is currently not performed automatically between Keycloak and Alien4Cloud. Thus, in the meantime, roles must be assigned explicitly to users within Alien4Cloud in order to appropriately restrict their rights.

The following roles are supported by Alien4Cloud:

- *Administrator*, who can configure which resources which user can create on each infrastructure,
- *Architect*, who can add new components and application templates in Alien4Cloud catalogue of components, but cannot deploy applications,
- *Application manager*, who cannot add new components in A4C catalogue, but can create new applications from these components, deploy/undeploy applications and run workflows.

To deploy applications or run workflows, Alien4Cloud communicates with its back end orchestrator YORC using TLS (Transport Layer Security) with mutual authentication through certificates.

YORC itself communicates with infrastructure using HEAppE described below (Section 4.2).

Configuration settings can be stored in a Vault (HashiCorp Vault) coming with the orchestrator, so that configuration files do not contain plain text values but keys referencing secrets in the vault.

For software deployments on Cloud compute instances, YORC is relying on “Ansible”, an open-source software managing provisioning, configuration, and deployment of applications. Ansible relies on “SSH” to perform operations on remote hosts. The “Ansible” vault feature is used. It allows to store connection secrets encrypted on the orchestrator host for the time of the execution of the operation on the remote host; afterwards, those files are deleted.

4.2 HEAPPE SECURITY

HEAppE performs the mapping of LEXIS users to functional (non-privileged) accounts for the HPC and OpenStack infrastructures in each centre. Due to this security-critical functionality, it is always deployed in a centre´s private network.

The external user accounts and internal cluster accounts will not be stored in HEAppE’s internal database. Instead, external (LEXIS) user accounts are stored in the LEXIS AAI (Keycloak), whereas internal accounts are kept in an SSH Agent instance. The PI of each computational project is responsible for adding internal cluster accounts into the specific SSH Agent instance, which he can and should run separately for each project. The HEAppE middleware architecture, shown in Figure 5, reflects this.

The HEAppE Middleware only enables the users to run a prepared set of so-called Command Templates. Each template practically includes an arbitrary script or executable file that will be executed on the cluster, any dependencies or third-party software it might require, and the queue/partition that should be used for processing (type of computing nodes to be used on the cluster).

HEAppE Middleware was primarily developed for an HPC infrastructure but in the scope of the LEXIS project the same security mechanism can be utilized also for an OpenStack environment. While the HEAppE Middleware encapsulates most of the HPC-related functionality for an HPC infrastructure the same authentication mechanism and the mapping functionality can be easily extended also for an OpenStack environment. For the OpenStack the HEAppE will use the same Keycloak authentication method but instead of mapping the LEXIS user accounts to internal HPC cluster accounts it will provide the LEXIS users with a valid Keystone token to be used via standard OpenStack APIs.

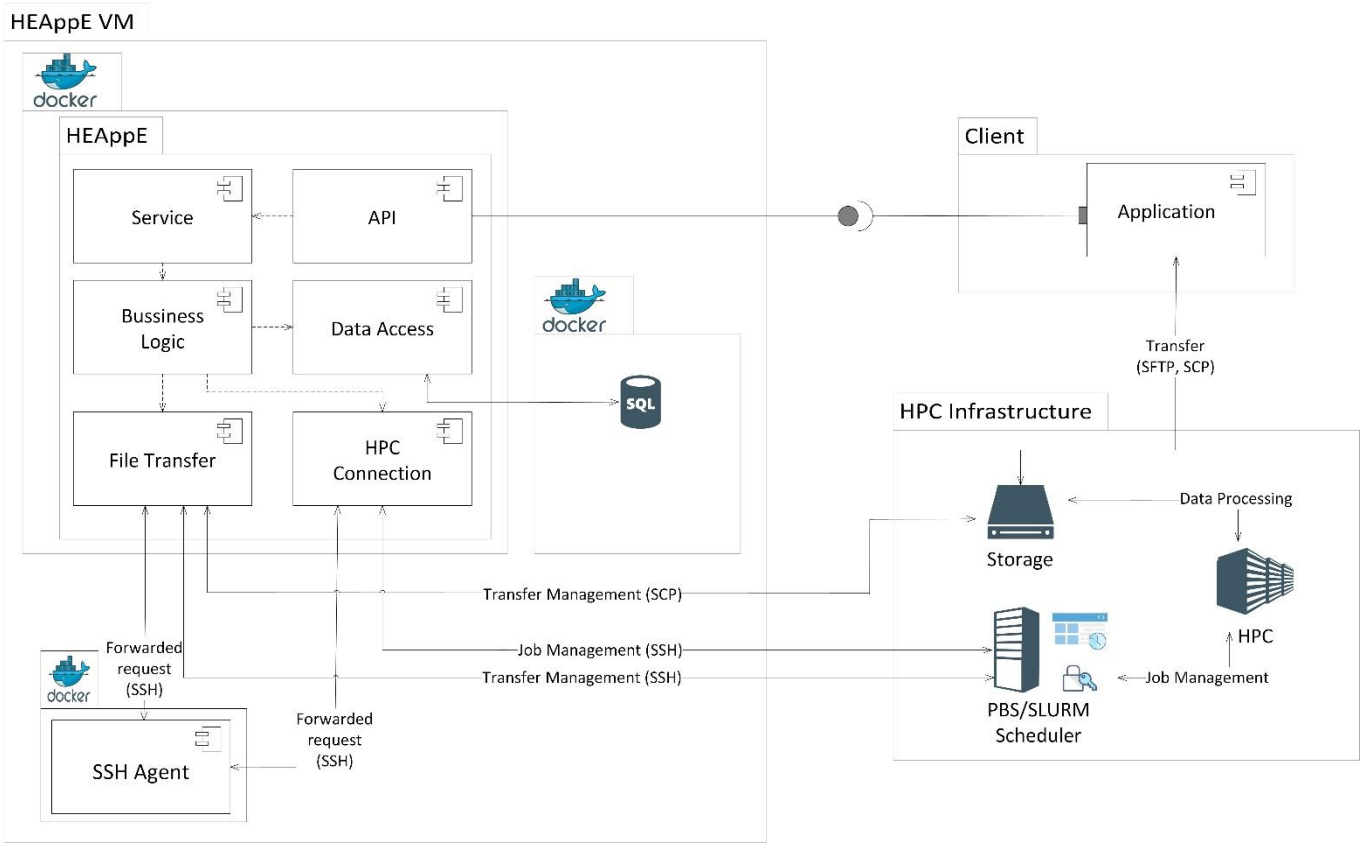

Figure 5 HEAppE Middleware Architecture

4.3 DDI SECURITY

The LEXIS DDI is based on the iRODS system, as discussed in [8]. Criteria and reasons for choosing this system, with a focus on security aspects, have been laid out in [1], and the actual setup is described in [6]. Before describing the restricted access to iRODS within the LEXIS DDI concept (Section 4.3.3), we give some information about iRODS security in general (Section 4.3.1) and about security in the LEXIS DDI configuration (Section 4.3.2).

4.3.1 General iRODS Security Features

iRODS generally features encrypted data traffic (based on certificates/keys generated when installing the system). In a federated setup, the administrators of two iRODS zones share a set of keys to federate. This ensures a secure connection between infrastructure at IT4I and LRZ in our LEXIS use case.

Login to iRODS is per default username/password based, and the connectivity is kept utilizing and keeping an access token in a local file. Alternative mechanisms, avoiding this relatively weak concept, are available. For LEXIS, we accordingly use an OpenID plugin for iRODS somewhat extended by the LEXIS WP3 team.

For logged-in users, iRODS uses strict permissions control with Access-Control Lists (ACLs) similar to a POSIX file system. One of three access levels (“own”, “write”, “read” in order of decreasing “power”) can be granted to any user or group in the ACL of a Data Object (“file”) or Collection (“folder”). Group membership is defined by the administrator(s) of a zone.

4.3.2 DDI configuration - security and data safety

Access to the LEXIS DDI is effected via OpenID authentication, using the central LEXIS AAI systems. The lifetime of LEXIS AAI and possible local tokens is appropriately limited. This fulfils the basic requirements to the LEXIS DDI

mentioned in [1], where design guidelines with regards to security were devised. Besides security, data safety is a major concern in the LEXIS DDI. At LRZ, a redundant setup of the complete iRODS system (see Figure 6 below and [11]) has been chosen on non-redundant, but automatically recovered/migrated virtual machines. At IT4I, redundancy is guaranteed directly on the virtual-machine level.

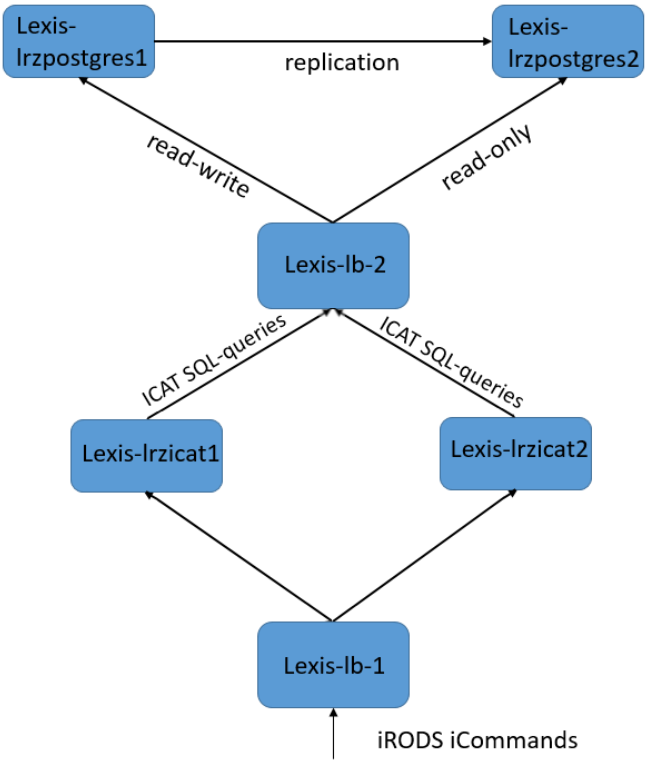

Figure 6 High-Availability iRODS concept at LRZ

Finally, data availability and disaster safety can be increased on demand of the LEXIS users even further: They are able to choose (at some higher cost in the LEXIS accounting system) to have their data duplicated in the IT4I and LRZ zones (see Section 2.3) of the Distributed Data Infrastructure. As the metadata catalogues of the different iRODS zones are separate, this further increases resiliency against failure, attacks and disasters.

4.3.3 Actions on data via DDI APIs

Access to the LEXIS DDI is provided exclusively through REST APIs (see Section 3.2 of [6]), which are right now implemented for data listing/search, data staging, user and rights management, and monitoring. This constrains and sanitises DDI usage patterns (e.g. only certain data movement schemes and source/target systems are possible). The LEXIS DDI APIs are open to the LEXIS Portal and the LEXIS Orchestrator; accessibility from outside the LEXIS ecosystem is not envisaged.

The APIs themselves run on dedicated virtual machines, where the necessary server software and tools (Apache/Nginx, uWSGI, Django, etc.) are installed, and undergo automatised updates when security fixes are available. Ports for new traffic to the APIs are exclusively opened for packets from the portal/orchestrator machines.

HTTPS requests for the APIs need to contain a token (see Section 3.2 of [6]) from the LEXIS AAI. The validity of this token is then verified by each of the LEXIS DDI APIs before actually inducing actions using this token on the iRODS-based distributed data infrastructure. The mechanisms for the LEXIS DDI to use OpenID tokens from Keycloak have been described in [6] and presented on a CS3 workshop in Denmark.

Generally, firewalling in LEXIS will prohibit direct access to any DDI machines. Only the access via higher-level systems (Portal, Orchestrator) through the DDI APIs will be allowed to the user, appropriately reducing the attack surface.

4.4 PORTAL SECURITY

Security is an important consideration for the LEXIS Portal as it will be exposed to the Internet. While general best practices have been employed within the development process and mature, trusted libraries and systems are used as the basis of the development, the complete solution is still evolving. Some questions regarding the deployment configuration are therefore not finally answered yet. Therefore, not all security related aspects of the LEXIS Portal can be addressed yet. However, we can already comprehensively describe how security has been considered in the system design, with some information on the envisaged deployment context.

4.4.1 Portal Security Considerations - Design and Development Perspective

The components of the LEXIS Portal [12] are shown in context in Figure 7. As can be seen, the LEXIS Portal comprises of a number of components, for each of which security has been considered.

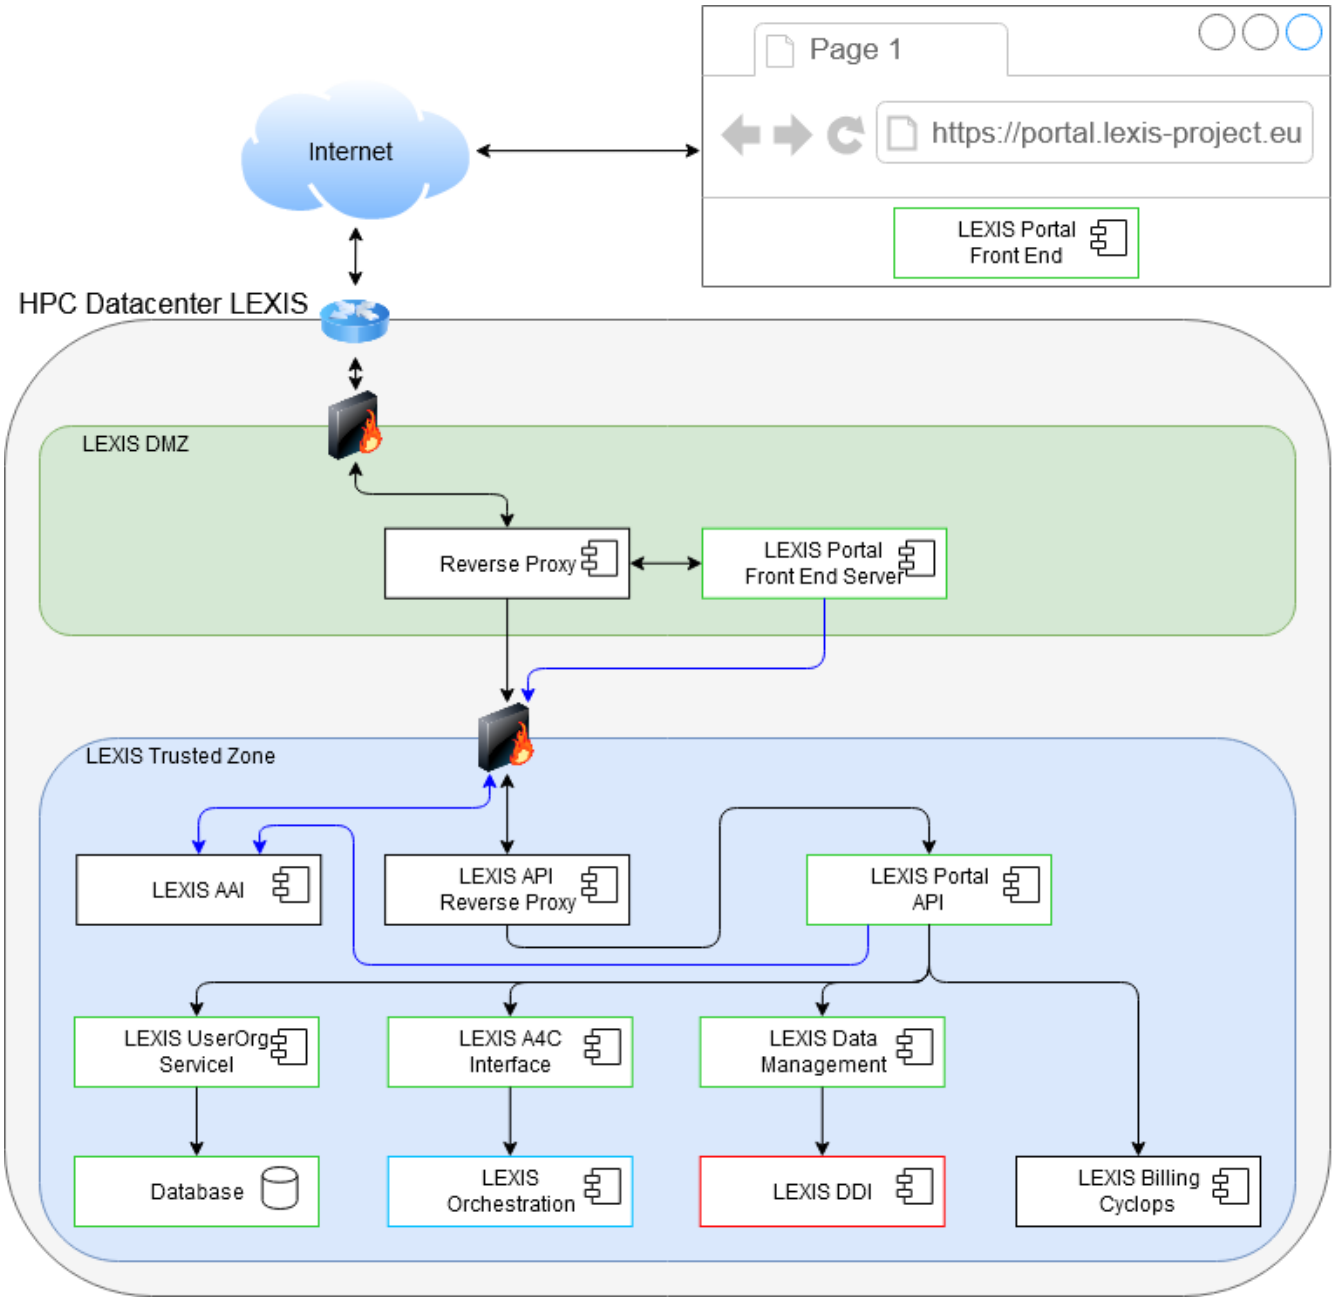

Figure 7 Security architecture of LEXIS Portal

The LEXIS Portal front end (FE) is a rich "React.js" application which runs in the browser. This is served from the LEXIS Portal FE server which offers a HTTP and HTTPS interface; the HTTPS interface is secured by a priori generated certificates. The FE itself comprises of a set of minified JavaScript components embedded in a simple HTML page; this is served as static content from the FE server.

The LEXIS Portal FE and FE server operate an OpenID Connect (OIDC) workflow to support log in to the LEXIS Portal; Keycloak (LEXIS AAI) operates as the OIDC server. To enable this, the LEXIS Portal FE Server is registered with Keycloak in the common realm for the LEXIS platform as a client - the Client ID and Client Secret are provided in the LEXIS Portal FE Server configuration file. The client configuration within Keycloak includes a callback which gets called following a successful login.

Log in follows the standard OIDC flow:

- Login endpoint is provided by the LEXIS Portal FE server, this returns a redirect to Keycloak which contains a specific state parameter,
- The browser is redirected to the login page for the appropriate Keycloak realm,
- The user authenticates within Keycloak,
- Keycloak redirects back to the a priori configured callback on the LEXIS Portal FE Server with a token,
- The LEXIS Portal FE Server calls an endpoint on Keycloak with the token to obtain the user information and a longer-lived bearer token for the user,
- This bearer token is provided to the FE via HTTPS - a specific endpoint is available within the LEXIS Portal FE Server which the FE calls,
- The FE can finally use this bearer token to communicate with the Portal API.

The Portal API is the single-entry point which the FE communicates with. It offers a REST based API defined using a Swagger 2 definition. All calls to this REST endpoint must provide a bearer token or they will be rejected as unauthorised. The LEXIS Portal API performs a token validation check with Keycloak for each token that is provided; if the token is valid, then the logic associated with that API Endpoint is executed.

The Portal API has some further security-oriented logic for some endpoints. Specifically, endpoints associated with obtaining accounting and billing information contain authorisation checks to ensure that the user has appropriate authorisations to obtain this information from the Cyclops accounting and billing system. More specifically, it checks that the user has the role ‘org_admin’ defined within Keycloak.

The LEXIS Portal back end services - UserOrg Service, Dataset Management Interface and Orchestration Interface - operate in a similar way to the Portal API: a Swagger 2 based API is defined together and authentication checks are performed on each REST call. Authorisation capabilities are being added to these services, as defined by the LEXIS RBAC matrix, but this work is not completed yet. The basic premise is that each service is responsible for ‘protecting’ itself and its data - as such, authentication and authorisations operate at multiple levels within the LEXIS Portal.

A major part of the LEXIS Portal back end services has been written in Go - a language that has been proven to be efficient and trustworthy from a security perspective.

4.4.2 Portal Security Considerations - Deployment Perspective

Out of the multiple services the LEXIS Portal is made of, two must be accessible via the Internet - the LEXIS Portal API and the LEXIS Portal FE server. The LEXIS AAI (Keycloak) instance must also be publicly accessible.

With the LEXIS Portal API as the main entry point for use of services, the services which it calls do not need to be exposed to the Internet - it is planned that these will operate in a trusted network zone not accessible via the Internet. Network based restrictions can thus block the access to the lowest-level LEXIS services from the outside of a computing centre without any functionality problems (apart possibly from insignificant exceptions, e.g., for cross-site testing/monitoring). This, again, helps to reduce the attack surface of the LEXIS platform.

5 SUMMARY

Security has been one of the main focus points within LEXIS since the early co-design phase of the platform.

The “Security by Design” principle has been taken into account in setting up every LEXIS component, and partners have proven very active in providing or adopting respective solutions. Thus, the LEXIS Portal, the LEXIS Orchestrator and HEAppE Middleware, and the LEXIS DDI in particular have reached a high degree of security. In this area, the experience from both IT4I and LRZ, combined with architectural guidance from Atos and Security guidance from O24 have been the key to our achievements. The current LEXIS infrastructure is a development infrastructure that will evolve to a production infrastructure, and in the course of this all the measures described in this document will be put in place.

A network security assessment will be conducted on the final version of the LEXIS platform in order to properly ensure the attack surface area has been minimised and no unnecessary services have been exposed. An application and web application security assessment will also be conducted specific focus on default configuration, separation of duties and failing securely.

On the other hand, a federated authentication and authorisation system has been provided in the context of federating tier 0 and tier 1 European computing/data centres within LEXIS. Clearly, with the different security policies already in place locally, this has been a challenge. We mastered this by separating the LEXIS user administration from the local one, with HEAppE providing a secure and flexible mapping and access mechanism for computing jobs. Authorisation within LEXIS is properly handled with a state-of-the-art RBAC Matrix. The implementation we are aiming at is aligned with the security standards used in Cloud Computing with minimal permissions granted to each user according to the access policies. The RBAC matrix used at this point of time is relatively simple in order to ease secure and clear implementation. Later, this can be extended by adding automation to split privileges while updating the user access rights.

Adding additional computing centres in LEXIS platform has already been taken into account in the design of the security architecture. As a matter of fact, any computing site in LEXIS can keep its own implementation of security, regulatory and sovereignty policies. Depending on the strategy, wishes and possibilities deriving from these policies, different degrees of integration with the LEXIS platform (full integration or integration as client) can be considered. The main security aspects are already taken care of, but of course the integration of further centres with LEXIS will be accompanied by the “Lessons Learned” process set up by LEXIS WP2. Likewise, decommissioning procedures for centres potentially leaving LEXIS will be devised and take care of aspects such as data deletion/preservation, and blocking access by invalidating tokens and passwords.

REFERENCES

[1] LEXIS Deliverable, *D4.1 Analysis of Mechanism for Securing Federate Infrastructure*.

[2] NCSC, “Secure design principles,” [Online]. Available: https://www.ncsc.gov.uk/collection/cyber-security-design-principles/cyber-security-design-principles. [Accessed April].

[3] NCSC, “Secure design principles: Guides for the design of cyber secure systems,” [Online]. Available: https://www.ncsc.gov.uk/collection/cyber-security-design-principles. [Accessed April 2020].

[4] INFOSEC, “Secure System Design Principles And The CISSP,” [Online]. Available: https://resources.infosecinstitute.com/category/certifications-training/cissp/domains/security-engineering/secure-system-design-principles/. [Accessed April 2020].

[5] OWASP, “Security by Design Principles,” [Online]. Available: https://wiki.owasp.org/index.php/Security_by_Design_Principles. [Accessed April 2020].

[6] LEXIS Deliverable, *D3.3 Mid-Term Infrastructure (Deployed System Hard/Software).*

[7] LEXIS Deliverable, *D2.3 Report of LEXIS Technology Deployment - Intermediate Co-Design.*

[8] LEXIS Deliverable, *D2.2 Key parts LEXIS Technology Deployed on Existing Infrastructure and Key Technologies Specification.*

[9] “OWASP Top Ten,” [Online]. Available: https://owasp.org/www-project-top-ten/. [Accessed April 2020].

[10] „OWASP API Security Project,“ April 2020. [Online]. Available: https://owasp.org/www-project-api-security/.

[11] “Configuring iRODS for High Availability,” iRODS, 7 July 2015. [Online]. Available: https://irods.org/2015/07/configuring-irods-for-high-availability/. [Accessed April 2020].

[12] LEXIS Deliverable, *D8.1 First Release of LEXIS Portal (will include report).*

A RBAC MATRIX

The first RBAC Matrix used for LEXIS AAI has been made taking into account all layers that could be provided within LEXIS Services (LEXIS Portal, LEXIS Orchestrator, and LEXIS DDI). As a matter of fact, this could be simplified for an easier implementation better fitting LEXIS project timeframe. A simplified version has been derived from full and detailed RBAC Matrix as this does not disallow to re-introduce changes later.

Please find the simplified version of RBAC Matrix below, with the following legend:

F: Full Access
P: Partial Access restricted by attributes based on the Organisation/Project/Workflow
P: Partial Access restricted by attributes based on the Organisation/Project/Workflow

PO: Partial Access restricted by Organisation
PP: Partial Access restricted by Project
PW: Partial Access restricted by Workflow

LEXIS User	LEXIS Workflow Manager	LEXIS Project Manager		LEXIS License Manager	LEXIS Financial Manager	LEXIS IAM Manager	LEXIS Organization Manager		LEXIS Support	LEXIS Administrator	LEXIS ROLES		
end_usr	wfl_mgr	prj_mgr		lic_mgr	fin_mgr	iam_mgr	org_mgr		lex_sup	lex_adm			LEXIS PERMISSIONS
	P (PO, PP)			P (PO)					P (PO)	F	Organization Management	org_list	List Organizations
												org_read	Read details of an
											Identity & Access Management	iam_list	List Users
												iam_read	Read details of a User
											Billing Management	bil_list	List Billing&Payment
												bil_read	Read details of a
											Licensing Management	lic_list	List Licensing informations
												lic_read	Read details of a Licensing
				P (PO)						F	Identity & Access	iam_write	Create/Delete/Update a User
											Organization	org_write	Create/Delete/Update an
											Billing Management	bil_write	Create/Delete/Update a
											Licensing Management	lic_write	Create/Delete/Update a
P (PO, PP)	P (PO, PP)			P (PO)					P (PO)	F	Project Management	prj_list	List Projects
											Workflow Management	wfl_list	List Workflows
P (PO, PP)	P (PO, PP)			P (PO)					P (PO)	F	Project Management	prj_read	Read details of a Project
											Workflow Management	wfl_read	Read details of a Workflow
											Project Management	prj_write	Create/Delete/Update a
											Workflow Management	wfl_write	Create/Delete/Update/Start/St
P (PP)	P (PP)								P (PP)	F	Computation Management (jobs, tasks of differents systems OpenStack/YORC/HEApp E)	cpu_list	List Computations
P (PP)												cpu_read	Read details of a Computation
P (PP)										F		cpu_write	Create/Delete/Update/Start/St op a Computation
P (PP)	P (PP)								P (PP)	F	Data Management (iRODS DDI and WCDA)	dat_list	List Datasets
P (PP)												dat_read	Read details of a Dataset
P (PP)										F		dat_write	Create/Delete/Update/Import/ Export a Dataset