Large-scale EXecution for Industry & Society

Deliverable D3.4

Monitoring System

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	D3.4 | Monitoring System
RESPONSIBLE AUTHOR	Martin Golasowski (IT4I)
WORKPACKAGE ID | TITLE	WP3 | LEXIS Data System
WORKPACKAGE LEADER	LRZ
DATE OF DELIVERY (CONTRACTUAL)	30/06/2021 (M30)
DATE OF DELIVERY (SUBMITTED)	12/07/2021 (M31)
VERSION | STATUS	V1.0 | Final
TYPE OF DELIVERABLE	O (Other)
DISSEMINATION LEVEL	PU (Public)
AUTHORS (PARTNER)	IT4I, LRZ
INTERNAL REVIEW	Piyush Harsh (CYC); Alberto Scionti (LINKS)

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	Initial version and table of contents created	14/06/2021	Martin Golasowski (IT4I)
0.2	Text and screenshots added	17/06/2021	Martin Golasowski (IT4I), Mohamad Hayek (LRZ)
0.21	Review	22/06/2021	Piyush Harsh (CYC); Alberto Scionti (LINKS)
0.3	Reviewer's comments addressed	22/06/2021	Martin Golasowski (IT4I)
0.4	Final version	28/06/2021	Martin Golasowski (IT4I)
1.0	Final check	12/07/2021	Kateřina Slaninová (IT4I)

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
EURAXENT	MARC DERQUENNES
GFZ	HELMHOLTZ ZENTRUM POTSDAM DEUTSCHESGEOFORSCHUNGSZENTRUM GFZ
ICHEC	NATIONAL UNIVERSITY OF IRELAND GALWAY / Irish Centre for High-End Computing
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 deliverable describes monitoring system which purpose is to collect performance metrics, provide visualisation and alerting and conduct health checks of critical components of the LEXIS infrastructure.

Position of the deliverable in the whole project context

The monitored infrastructure related to the LEXIS DDI is described by the previous deliverables D3.2 [1] and D3.3 [2]. Even though this deliverable is mainly outcome of task T3.4 positioned in WP3 - LEXIS Data System, it also describes monitoring of other elements of the infrastructure used in IT4I and LRZ through all WPs of the LEXIS project.

1 INTRODUCTION

Purpose of the Monitoring System is to provide performance metrics, health-checks and alerting capabilities for all critical modules of the LEXIS platform. This document briefly describes its main design ideas, tools used and primarily documents its deployment in LRZ and IT4I as the two HPC centres where the federation through the LEXIS platform is tested. The system can be deployed to other centres as needed.

In terms of the observability theory, this Monitoring System focuses mainly on collecting performance metrics from the individual parts of the infrastructure, performing health checks, and sending alerts to the operations team via e-mail, chatbots and other channels. Part of the system is also the infrastructure for collecting selected log messages and events, primarily for DDI auditing purposes.

Main assets created in the Task 3.4 – LEXIS Monitoring System are the deployment of the Monitoring System components in both LRZ and IT4I, configuration files and scripts, customised Grafana dashboards and collection of Robot test cases and their descriptions. Selected components of the Monitoring System are published on GitHub[1]

2 ARCHITECTURE

The services which implement the LEXIS platform (DDI, AAI, Orchestration, Portal) are predominantly deployed in virtual and bare-metal machines with containers in selected places. The monitoring systems must be able to monitor all levels of the infrastructure spanning from the machines themselves, through the support services of the infrastructure, up to the LEXIS services. It also must account for slight differences in the physical infrastructure provided by IT4I and LRZ. Figure 1 shows the individual elements of the system with emphasis on various levels of the infrastructure.

[1] Monitoring system at GitHub: https://github.com/lexis-project/monitoring-system

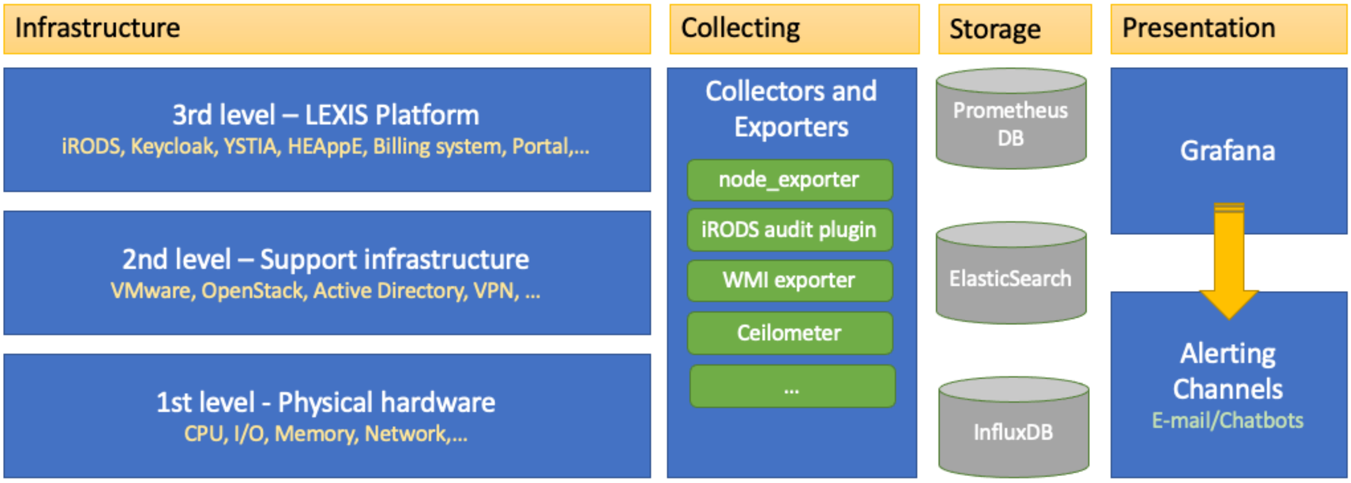
Infrastructure
Collecting
Storage
Presentation
3rd level – LEXIS Platform
iRODS, Keycloak, YSTIA, HEAppE, Billing system, Portal,...
2nd level – Support infrastructure
VMware, OpenStack, Active Directory, VPN, ...
1st level - Physical hardware
CPU, I/O, Memory, Network,...
Collectors and Exporters
node_exporter
iRODS audit plugin
WMI exporter
Ceilometer
...
Prometheus DB
ElasticSearch
InfluxDB
Grafana
Alerting Channels
E-mail/Chatbots

Figure 1: Levels of the LEXIS Monitoring System

2.1 PERFORMANCE METRICS - PROMETHEUS/GRAFANA

Performance metrics are collected in Prometheus timeseries DB[2] deployed in both LRZ and IT4I centres. Each instance then periodically collects metrics from configured endpoints - namely exporters. For example, basic metrics like CPU load or RAM usage in Linux are collected by a dedicated node_exporter running in each instance of Linux OS in the infrastructure. Some services like CEPH or HAProxy expose their own metric endpoint which can be scraped by Prometheus. The complete list of deployed exporters is presented in Section 3.

Grafana[3] is a web based timeseries visualisation tool capable of working with various data sources such as timeseries, document or SQL databases. It is used by the Monitoring System as main tool for visualisation and alerting. Each centre (LRZ and IT4I) which has Grafana deployed as part of the Monitoring System, implements its own set of dashboards. Some of them are shared between the centres, some are specialised to account for differences in the infrastructures. Grafana is also used to implement the alerting pipeline.

2.2 HEALTH CHECKS

Purpose of health checks is to periodically monitor correct behaviour of a service and to raise an alarm in case the behaviour changes (i.e., HTTP response code is not 200, etc.). In the LEXIS Monitoring System, we are using two types of health checks:

- Basic HTTP, TCP and ICMP checks using blackbox exporter and HAProxy backends,
- Continuously executed Robot test suites.

The first component is used to verify that a particular web server is running and responding correctly, or a machine responds to ping. The solution is based on blackbox exporter and on metrics exported by HAProxy for endpoints which are deployed in High Availability (HA) mode.

We have collected a set of test cases from all the LEXIS platform components developed in other technical WPs and implemented them using the two solutions described above. The collected test cases are available on GitHub[4].

2.2.1 Robot test suites

The second component — Robot test suites — is used to implement more complex tests which verify various functionalities of the platform components. The Monitoring System runs these tests continuously and sends error

[2] Prometheus - https://prometheus.io

[3] Grafana - https://grafana.com

[4] Monitoring system at GitHub: https://github.com/lexis-project/monitoring-system

messages on several alerting channels in case the tests are failing. The results of each run are stored in a SQL database using TestArchiver[5]. The SQL database is connected to the Grafana as data source and the test results are visualised on a custom dashboard.

List of implemented test suites:

- iRODS Local - basic test of iRODS zone, obtains session with password and OpenID token and performs a file transfer,
- iRODS Federation - tests file transfer between federated zones with OpenID token,
- Auth test - obtains and validates OpenID token from LEXIS AAI,
- API test - verifies function of DDI APIs,
- Handle test - verifies B2SAFE/HANDLE.NET server function.

2.3 LOG AUDIT - ELASTICSEARCH STACK

LEXIS DDI is based on iRODS which provides auditing capabilities through a plugin. The plugin emits many AMQP messages as the individual requests are processed by the iRODS iCAT servers. The messages are sent to the AMQP broker RabbitMQ and consumed by Logstash instance which stores them to the ElasticSearch database. A set of custom dashboards has been created in Grafana to monitor various metrics such as:

- Data throughput read/write,
- I/O operations count,
- Statistics per project/user/organization.

These data can be used to obtain an audit trail for all datasets stored in the LEXIS DDI.

2.4 PUBLIC HEALTH STATUS PAGE

LEXIS platform operational status is reported to the users by embedding a set of Grafana dashboards in the LEXIS portal. These dashboards are based on the health checks described in Section 2.2 and provide basic overview of the operational status for the platform users. Figure 2 shows the dashboards for services deployed at LRZ, this dashboard will be available in the portal.

[5] TestArchiver: https://github.com/salabs/TestArchiver

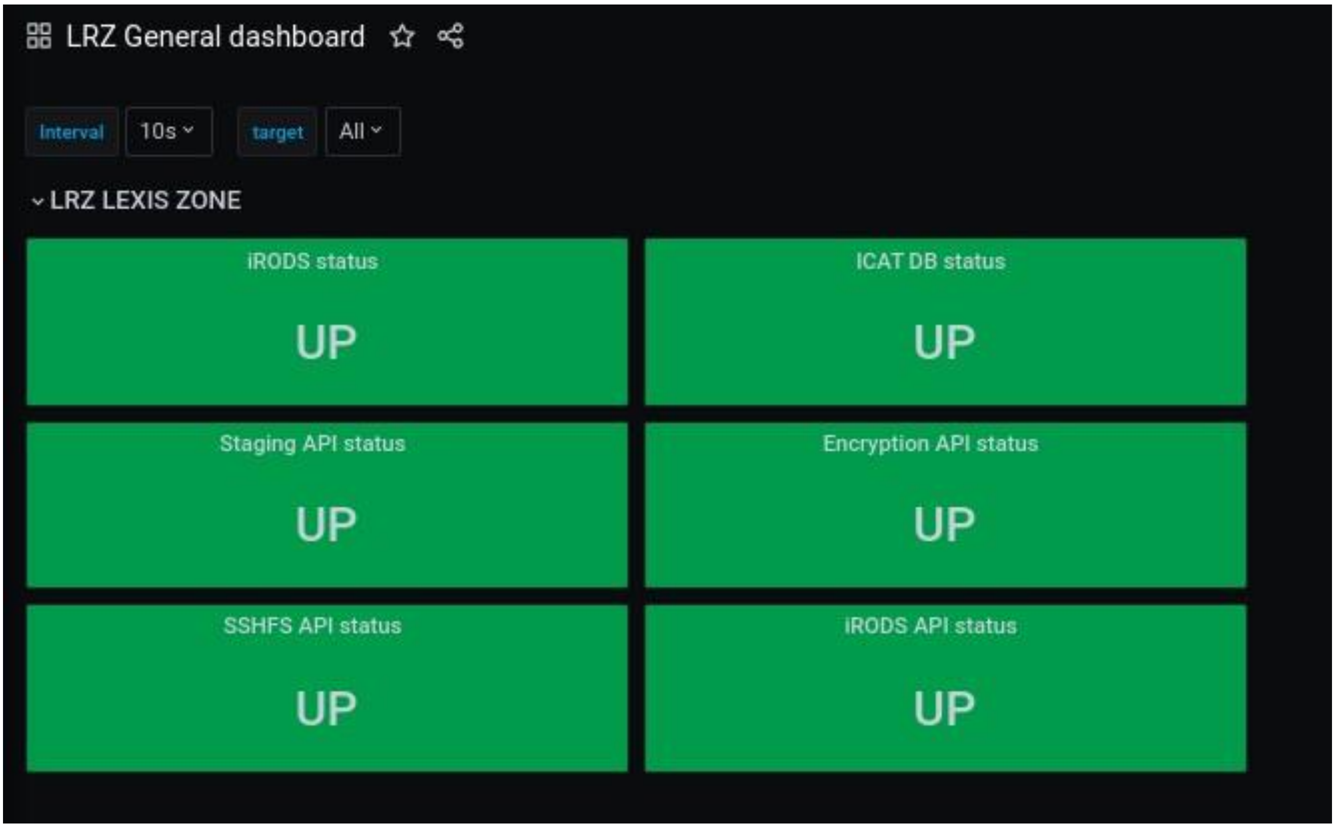

LRZ General dashboard

Interval 10s

target All

LRZ LEXIS ZONE

iRODS status UP	ICAT DB status UP
Staging API status UP	Encryption API status UP
SSHFS API status UP	iRODS API status UP

Figure 2: General LEXIS dashboard at LRZ

3 IMPLEMENTATION AND DEPLOYMENT

This section lists the individual Monitoring System components deployed at IT4I and LRZ. Screenshots of selected important Grafana dashboards are provided as well. Thanks to the flexibility of Grafana, we were able to reuse a lot of detailed dashboards for exporters that provide a large number of metrics. Example of such dashboard is in Figure 3 which shows dashboard provided by the community at https://grafana.com/grafana/dashboards/12566 imported in IT4I and used to monitor the Windows based instances.

Example of a custom-made dashboard is in Figure 5 where it shows network utilization on the VPN servers. This flexibility was one of the main points why we have selected Grafana as the solution for metrics visualisation and alerting.

3.1 IT4INNOVATIONS

List of Monitoring System components endpoints deployed at IT4I.

- Grafana: `https://monitoring.msad.it4i.lexis.tech`,
- Prometheus: `https://monitoring.msad.it4i.lexis.tech:9090`.

3.1.1 Prometheus targets

List of target types configured at IT4I Prometheus DB instance:

- Node exporter (https://github.com/prometheus/node_exporter) - system metrics for all Linux OS instances, virtualized and bare-metal,
- Blackbox exporter (https://github.com/prometheus/blackbox_exporter) - HTTP, ICMP and TCP health checks,
- Celery exporter (https://github.com/OvalMoney/celery-exporter) - monitoring of DDI API Celery workers,
- HAProxy (https://www.haproxy.org) - traffic metrics provided by HAProxy instances,
- Windows exporter (https://github.com/prometheus-community/windows_exporter) - system metrics for all Windows OS instances,
- Keycloak (https://github.com/aerogear/keycloak-metrics-spi) - metrics exported by Keycloak AAI,
- CEPH (https://ceph.io/en) - metrics exported by CEPH cluster (IO, storage, cluster status, etc.).

3.1.2 Grafana dashboards

List of the most important Grafana dashboards used at IT4I.

- 1st level - basic hardware and OS metrics
  - Base metrics like CPU load, RAM usage or disk space used and many other metrics related to the machine state. Used to spot performance issues and raise alarms when an instance gets overloaded or there is not enough free disk space or RAM.

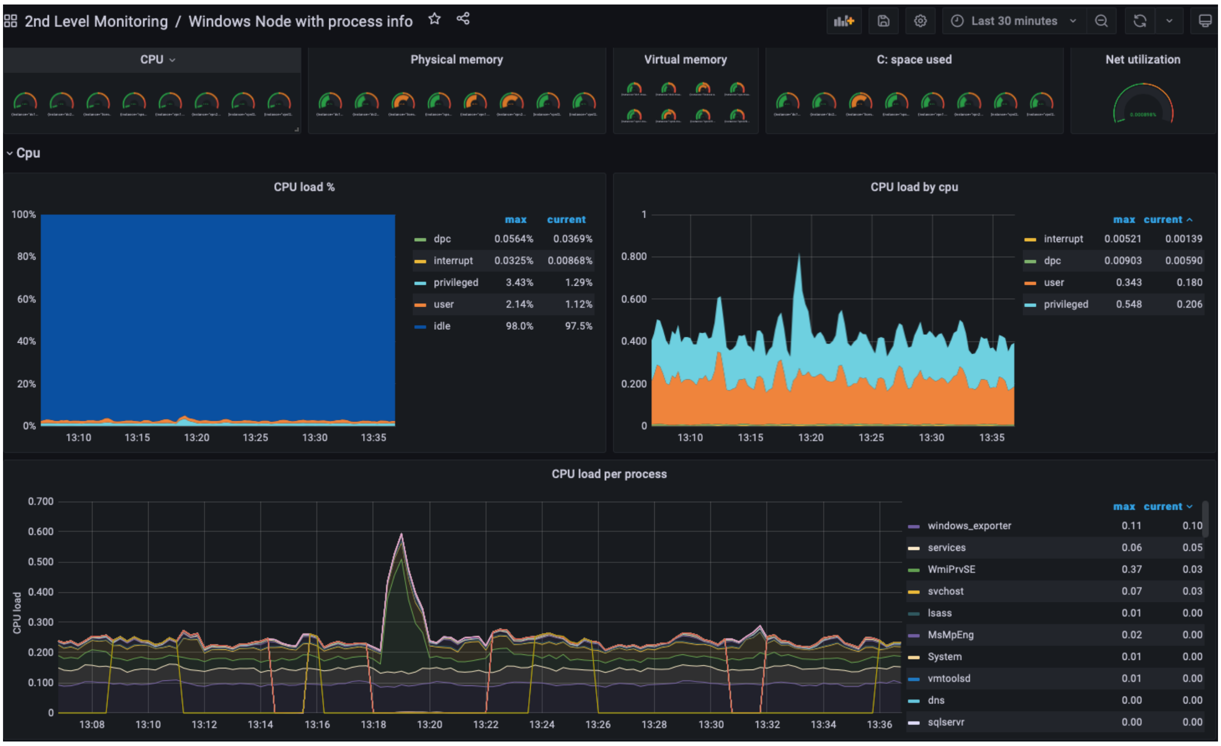

	max	current
dpc	0.0564%	0.0369%
interrupt	0.0325%	0.00868%
privileged	3.43%	1.29%
user	2.14%	1.12%
idle	98.0%	97.5%

	max	current ^
interrupt	0.00521	0.00139
dpc	0.00903	0.00590
user	0.343	0.180
privileged	0.548	0.206

	max	current ⌄
windows_exporter	0.11	0.10
services	0.06	0.05
WmiPrvSE	0.37	0.03
svchost	0.07	0.03
lsass	0.01	0.00
MsMpEng	0.02	0.00
System	0.01	0.00
vmtoolsd	0.01	0.00
dns	0.00	0.00
sqlservr	0.00	0.00

Figure 3: Dashboard for Windows metrics at IT4I

- 2nd level - infrastructure services
  - CEPH cluster monitoring (see Figure 4) provides an overview of the state and current utilization of the CEPH cluster operated by IT4I within the experimental LEXIS infrastructure.

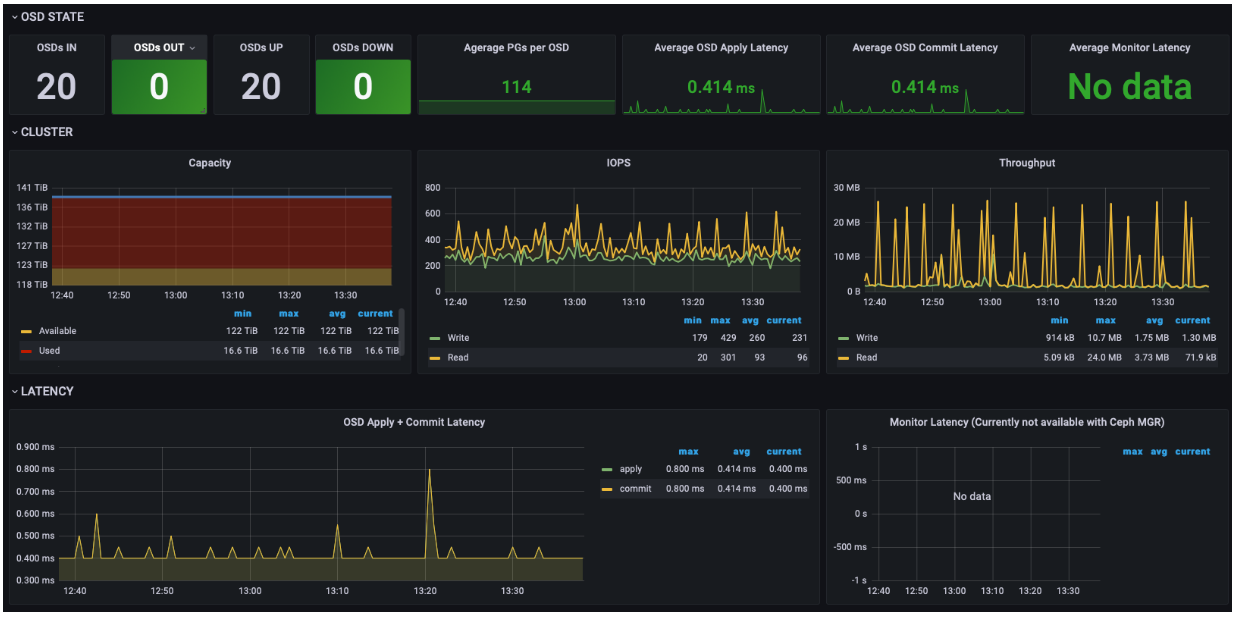

	min	max	avg	current
Available	122 TiB	122 TiB	122 TiB	122 TiB
Used	16.6 TiB	16.6 TiB	16.6 TiB	16.6 TiB

	min	max	avg	current
Write	179	429	260	231
Read	20	301	93	96

	min	max	avg	current
Write	914 kB	10.7 MB	1.75 MB	1.30 MB
Read	5.09 kB	24.0 MB	3.73 MB	71.9 kB

	max	avg	current
apply	0.800 ms	0.414 ms	0.400 ms
commit	0.800 ms	0.414 ms	0.400 ms

Figure 4: Dasboard for CEPH cluster metrics at IT4I

  - pfSense Firewall logs, System dashboards, WAN traffic,
  - HAProxy dashboards,
  - VPN L2L Servers, tunnel status to IT4I and LRZ.
- 3rd level
  - iRODS iCAT, network traffic,
  - Keycloak,
  - MQTT monitoring,
  - Robot DDI tests,
  - Celery DDI workers,
  - LEXIS endpoints health check.

3.1.3 pfSense Firewall monitoring

PfSense instances at IT4I are monitored using dedicated database instances. Performance metrics and other stats are stored in the InfluxDB, firewall, intrusion detection and prevention system (IDS/IPS) logs are stored in ElasticSearch. Both of these instances are added to the Grafana instance at IT4I as additional data sources and dashboards are available. Most important metrics collected are:

- Network bandwidth utilization and packet count per VLAN and IP protocol,
- Machine metrics like CPU and RAM load, free disk space,
- Suspicious transmissions, firewall hits.

3.1.4 VPN tunnel monitoring

There are two VPN services provided in IT4I LEXIS infrastructure:

- `vpn.it4i.lexis.tech`: The developer's VPN to reach machines used to host the LEXIS platform services at IT4I (dashboard in Figure 5).

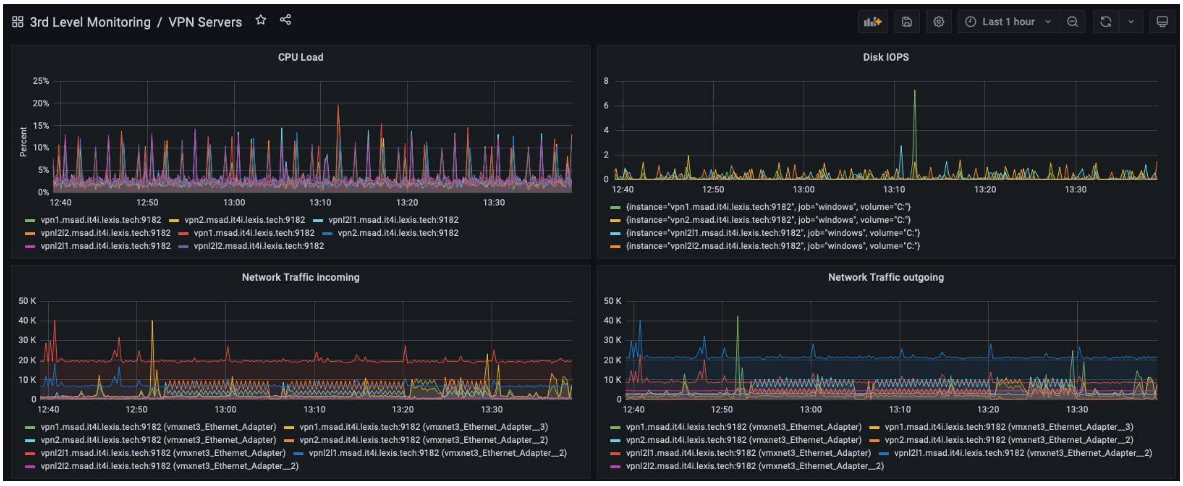
3rd Level Monitoring / VPN Servers
Last 1 hour
CPU Load
Percent
25%
20%
15%
10%
5%
0%
12:40
12:50
13:00
13:10
13:20
13:30
vpn1.msad.it4i.lexis.tech:9182
vpn2.msad.it4i.lexis.tech:9182
vpnl2l1.msad.it4i.lexis.tech:9182
vpnl2l2.msad.it4i.lexis.tech:9182
vpn1.msad.it4i.lexis.tech:9182
vpn2.msad.it4i.lexis.tech:9182
vpnl2l1.msad.it4i.lexis.tech:9182
vpnl2l2.msad.it4i.lexis.tech:9182
Disk IOPS
8
6
4
2
0
{instance="vpn1.msad.it4i.lexis.tech:9182", job="windows", volume="C:"}
{instance="vpn2.msad.it4i.lexis.tech:9182", job="windows", volume="C:"}
{instance="vpnl2l1.msad.it4i.lexis.tech:9182", job="windows", volume="C:"}
{instance="vpnl2l2.msad.it4i.lexis.tech:9182", job="windows", volume="C:"}
Network Traffic incoming
Network Traffic outgoing
50 K
40 K
30 K
20 K
10 K
0
vpn1.msad.it4i.lexis.tech:9182 (vmxnet3_Ethernet_Adapter)
vpn1.msad.it4i.lexis.tech:9182 (vmxnet3_Ethernet_Adapter__3)
vpn2.msad.it4i.lexis.tech:9182 (vmxnet3_Ethernet_Adapter)
vpn2.msad.it4i.lexis.tech:9182 (vmxnet3_Ethernet_Adapter__2)
vpnl2l1.msad.it4i.lexis.tech:9182 (vmxnet3_Ethernet_Adapter)
vpnl2l1.msad.it4i.lexis.tech:9182 (vmxnet3_Ethernet_Adapter__2)
vpnl2l2.msad.it4i.lexis.tech:9182 (vmxnet3_Ethernet_Adapter__2)

Figure 5: VPN Traffic dasboard

- `vpnsites.it4i.lexis.tech`: The Lan2Lan (L2L) VPN which provides secure tunnel for HA synchronization of various services.

Each service is backed by two servers deployed in high availability mode. The VPN is provided by the Microsoft Routing and Remote Access Services (RRAS), logs from these services are exported by a PowerShell script to a MSSQL database. The logs contain connect/disconnect events and other metrics provided by the service. Custom dashboard and alerting rules have been created in IT4I Grafana.

The L2L VPN tunnel is monitored using ICMP checks of the remote machines with blackbox exporter with alerting rules to ensure correct operation of the tunnel.

3.1.5 iRODS auditing - ElasticSearch

The iRODS auditing infrastructure is used to collect and display audit messages produced by the iRODS iCAT servers as they are accessed by the iRODS clients. Detailed information about the iRODS auditing can be found in the LEXIS Deliverable D3.3 [2, 3].

Components of the iRODS auditing infrastructure deployed at IT4I:

- ElasticSearch + Logstash - `elastic.msad.it4i.lexis.tech`
- RabbitMQ broker - `rabbitmq.msad.it4i.lexis.tech`

This ELK instance is added in Grafana as data source and custom dashboards have been created in Grafana to report various metrics.

3.1.6 Robot tests

- Continuous execution of test suites - `robot.msad.it4i.lexis.tech`
- SQL Database with test results - `postgres-lexis.msad.it4i.lexis.tech`

3.1.7 Raspberry Pi Monitoring screen kiosk

We have also created an image for the Raspberry Pi computer which automatically connects to the LEXIS VPN and displays a selected playlist of rotating dashboards from Grafana. It is used by IT4I staff responsible for the LEXIS infrastructure operations to display various metrics on big screens to provide overview and quick situation awareness. One of the screens used to display the dashboards is in Figure 6 along with the Raspberry Pi computer which drives the display in Figure 7

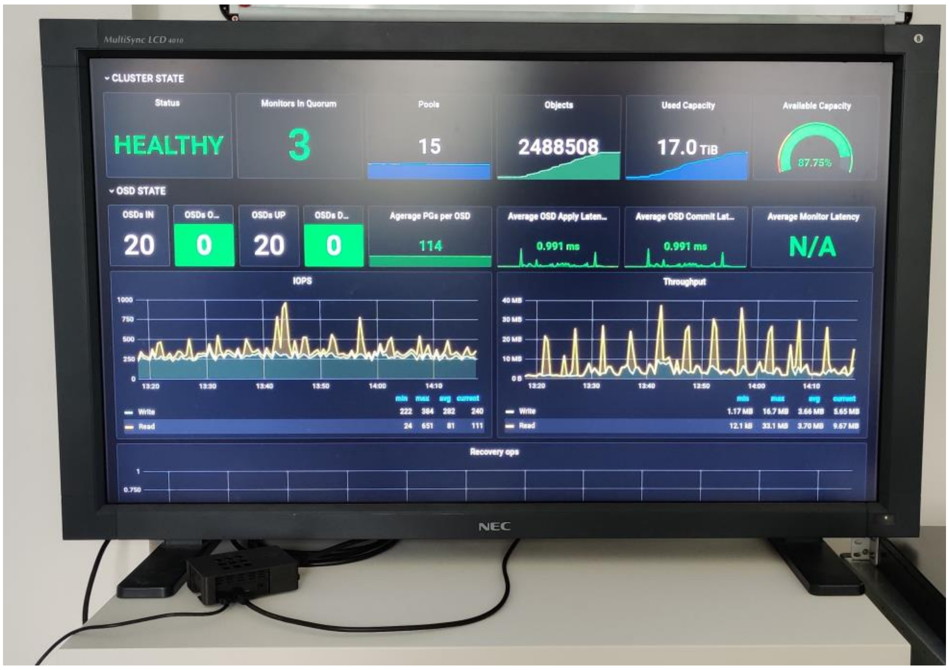

Figure 6: CEPH Dahboard displayed on large screen with Raspberry Pi

Figure 7: Raspberry Pi computer with custom image for monitoring display

3.2 LRZ

List of Monitoring System components deployed at LRZ.

- Grafana: `https://sikplrz-lexis-elasticsearch.srv.mwn.de/grafana`
- Prometheus: `https://sikplrz-lexis-elasticsearch.srv.mwn.de/prometheus/`

3.2.1 Prometheus targets

List of target types configured at LRZ Prometheus DB instance:

- Node exporter - system metrics for all Linux OS instances (see Figure 8)

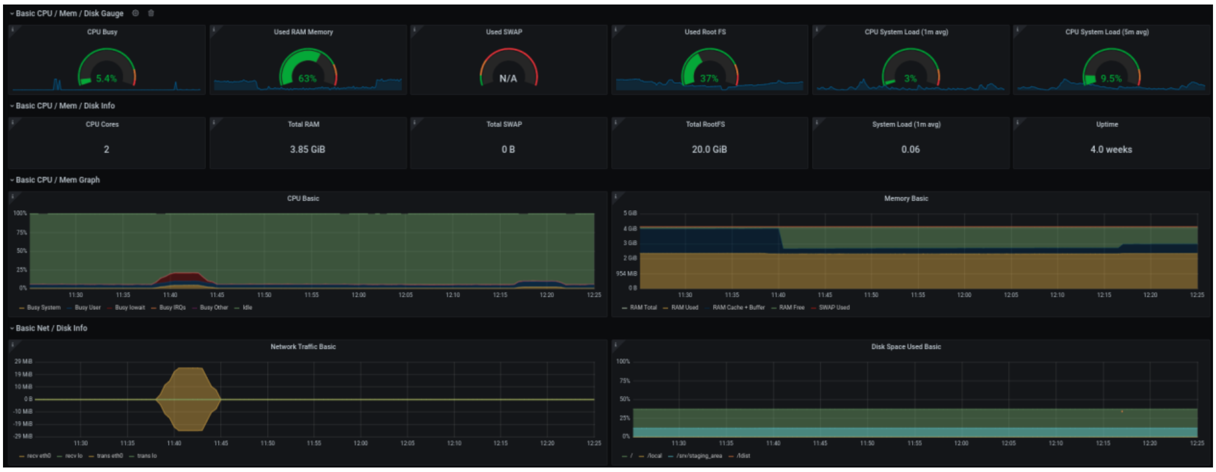

Figure 8: Linux OS metrics dashboard at LRZ

- Blackbox exporter - HTTP, ICMP and TCP healthchecks (see Figure 9)

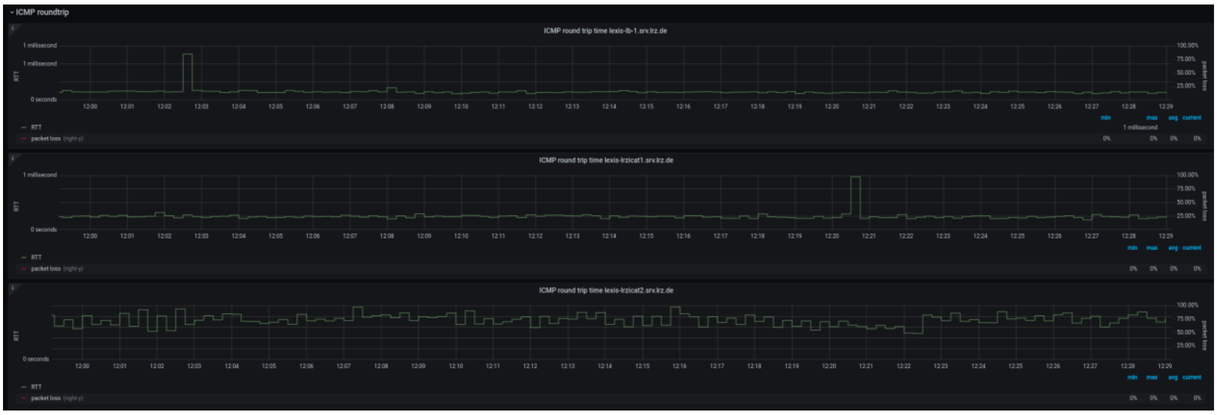
ICMP roundtrip
ICMP round trip time lexis-lb-1.srv.lrz.de
ICMP round trip time lexis-lrzicat1.srv.lrz.de
ICMP round trip time lexis-lrzicat2.srv.lrz.de
RTT
packet loss (right-y)
min
max
avg
current

Figure 9: Healthcheck dashboard at LRZ

- Celery - monitoring of DDI API Celery workers (see Figure 10)

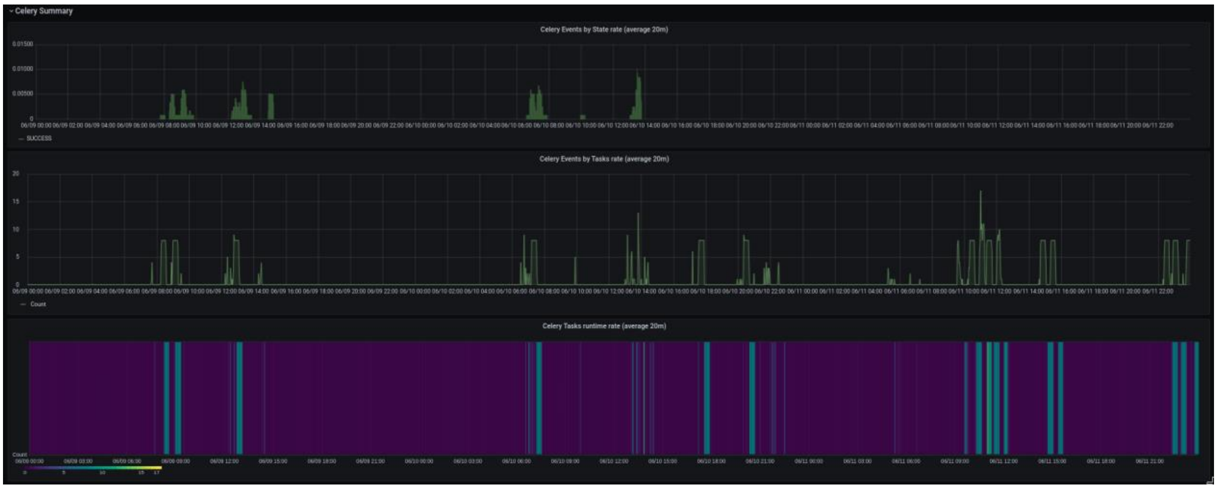
Celery Summary
Celery Events by State rate (average 20m)
SUCCESS
Celery Events by Tasks rate (average 20m)
Count
Celery Tasks runtime rate (average 20m)

Figure 10: Celery worker dashboard at LRZ

- HAProxy - traffic metrics provided by HAProxy instances (see Figure 11)

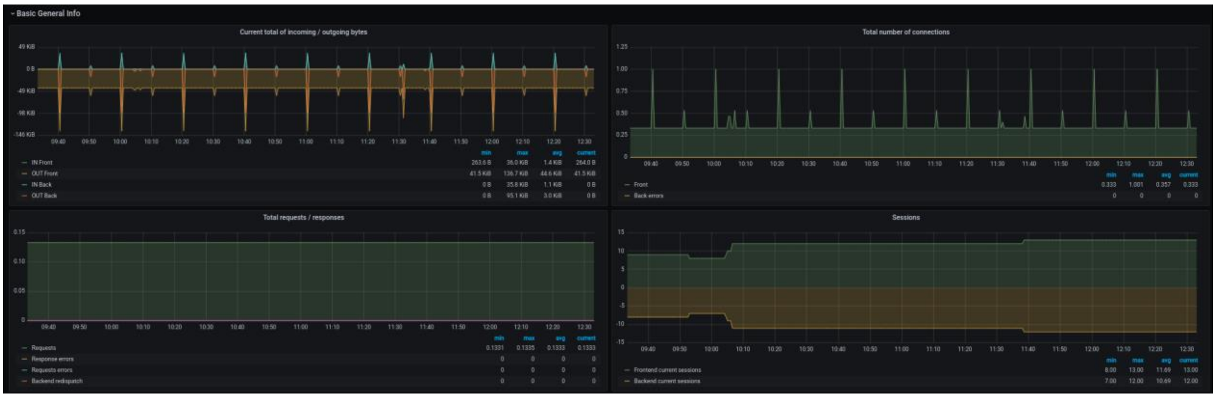

Figure 11: HAProxy dashboard at LRZ

- PostgreSQL exporter - monitor the database access and load (see Figure 12)

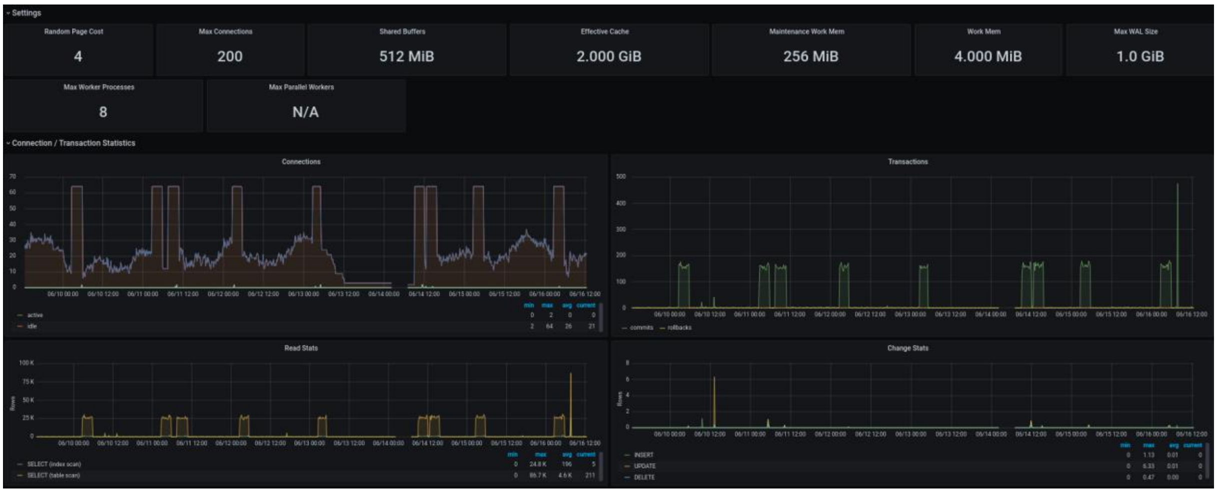
Settings

Random Page Cost 4

Max Connections 200

Shared Buffers 512 MiB

Effective Cache 2.000 GiB

Maintenance Work Mem 256 MiB

Work Mem 4.000 MiB

Max WAL Size 1.0 GiB

Max Worker Processes 8

Max Parallel Workers N/A

Figure 12: PostgreSQL server dashboard at LRZ

- MYSQL exporter - monitor the keycloak database access and load (see Figure 13)

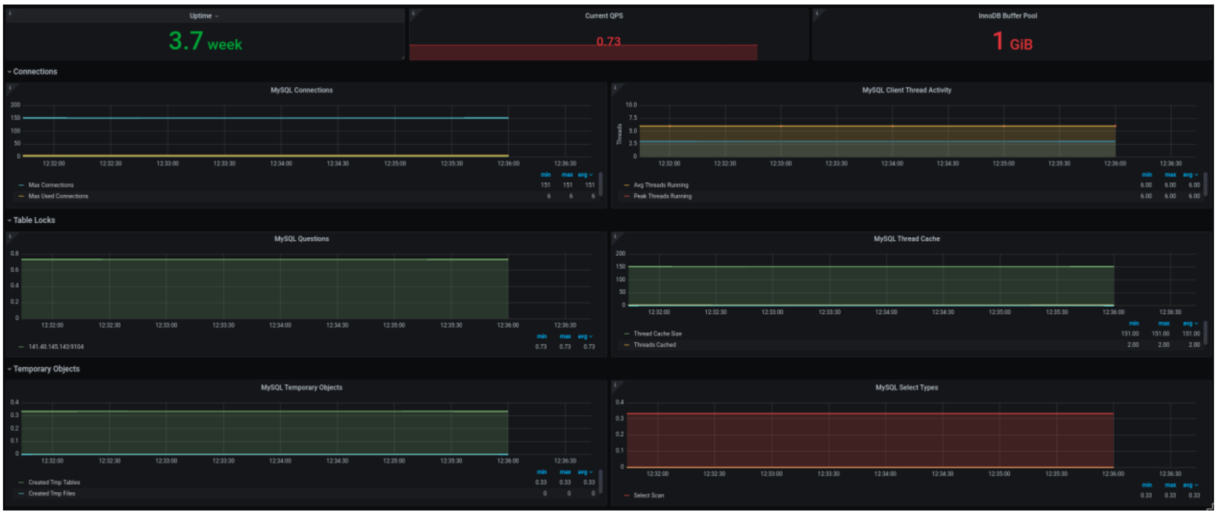
Uptime 3.7 week

Current QPS 0.73

InnoDB Buffer Pool 1 GiB

Figure 13: MySQL server dashboard at LRZ

3.2.2 Telegraph and InfluxDB targets

List of target types configured at LRZ InfluxDB instance:

- Nginx monitoring - shows the geolocation of requests coming into the staging API. An example is depicted in Figure 14.
- Pfsense firewall monitoring - monitors the incoming and outgoing traffic through the firewall. An example is depicted in Figure 15.

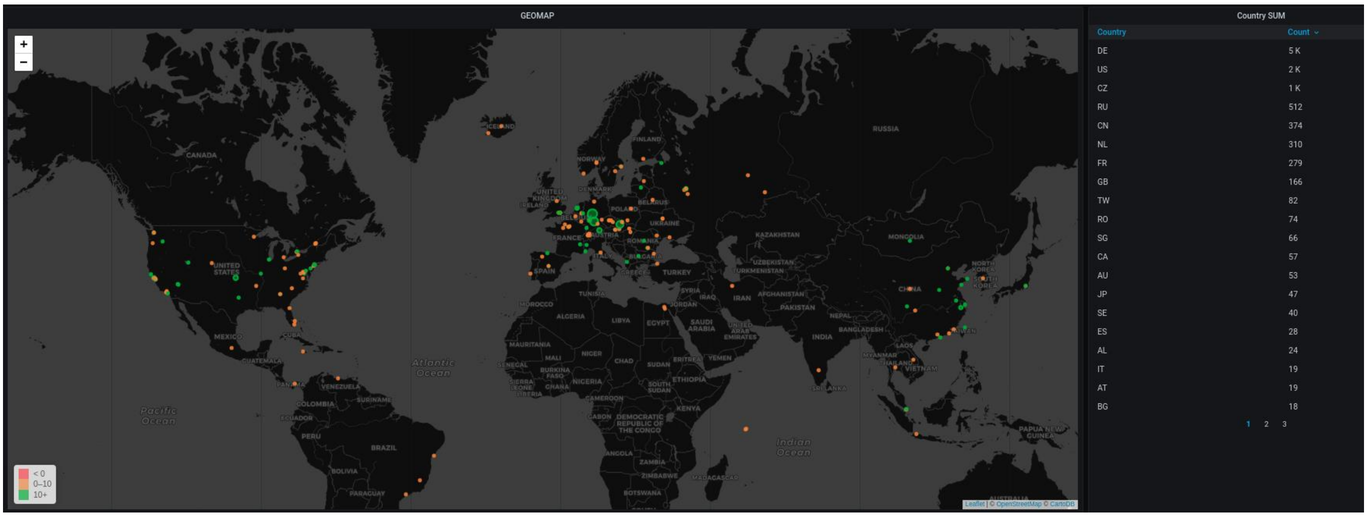

Country	Count
DE	5 K
US	2 K
CZ	1 K
RU	512
CN	374
NL	310
FR	279
GB	166
TW	82
RO	74
SG	66
CA	57
AU	53
JP	47
SE	40
ES	28
AL	24
IT	19
AT	19
BG	18

Figure 14: Geolocation dashboard for nginx server at LRZ

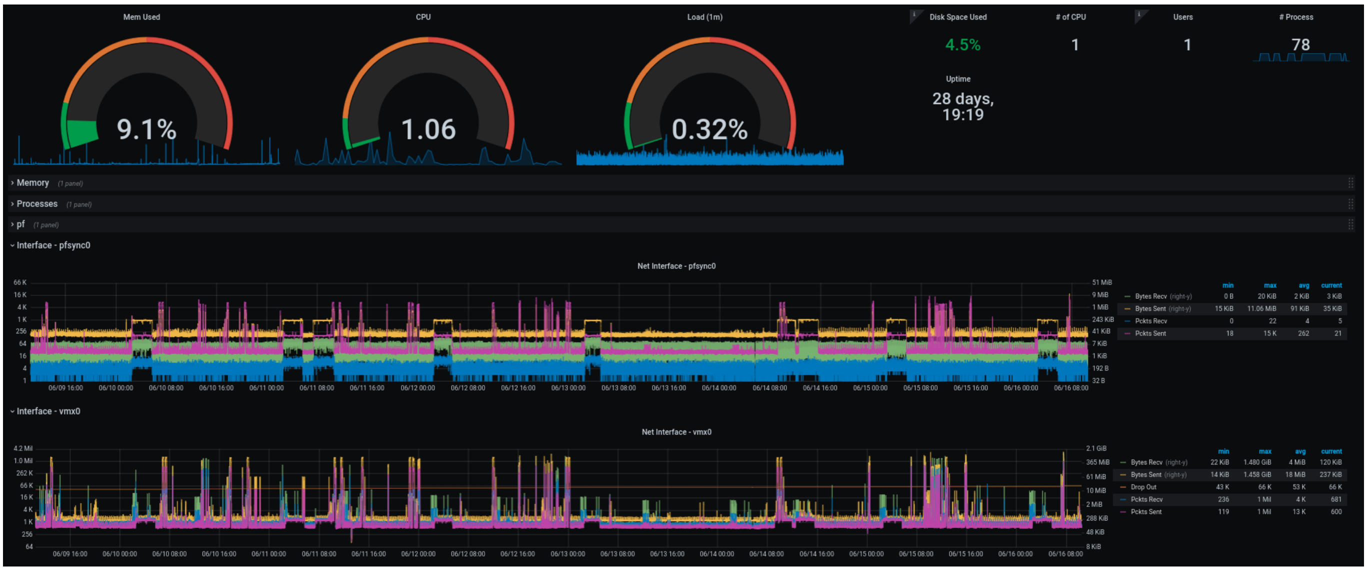

	min	max	avg	current
Bytes Recv (right-y)	0 B	20 KiB	2 KiB	3 KiB
Bytes Sent (right-y)	15 KiB	11.06 MiB	91 KiB	35 KiB
Pckts Recv	0	22	4	5
Pckts Sent	18	15 K	262	21

	min	max	avg	current
Bytes Recv (right-y)	22 KiB	1.480 GiB	4 MiB	120 KiB
Bytes Sent (right-y)	14 KiB	1.458 GiB	18 MiB	237 KiB
Drop Out	43 K	66 K	53 K	66 K
Pckts Recv	236	1 Mil	4 K	681
Pckts Sent	119	1 Mil	13 K	600

Figure 15: pfSense system dashboard at LRZ

3.2.3 Grafana dashboards

List of the most important Grafana dashboards used at LRZ:

- 1st level - basic hardware and OS metrics
  - Node stats.
- 2nd level - infrastructure services
  - pfSense Firewall logs, System dashboards, WAN traffic,
  - HAProxy dashboards,
  - VPN L2L client, tunnel status to IT4I and LRZ.

- 3rd level
    - iRODS iCAT, network traffic,
    - PostgreSQL DB,
    - Keycloak and keycloak DB,
    - Robot DDI tests,
    - Celery DDI workers,
    - API monitoring (Staging API, Gridmap API, and SSHFS API),
    - LEXIS endpoints healthcheck.

3.2.4 iRODS auditing - ElasticSearch

The iRODS auditing with ElasticSearch:

- ElasticSearch + Logstash - `sikplrz-lexis-elasticsearch.srv.mwn.de`
- RabbitMQ broker - `lexis-lrzdata-steering.srv.lrz.de`

This ELK instance is added in Grafana as data source and custom dashboards have been created in Grafana to report various metrics.

3.2.5 Robot tests

- Continuous execution of test suites - `sikplrz-lexis-elasticsearch.srv.mwn.de`
- SQL Database with test results - `lexis-lrzdata-steering.srv.lrz.de`

4 CONCLUSION

This short report provides an overview of the LEXIS Monitoring System with a focus on the parts deployed at IT4I and LRZ. It is based on technologies widely used in the industry, where the large fraction of the task T3.4 effort was devoted to design, deploying and configuring the various components of the system. Components and services of the LEXIS platform deployed in both centres are the same; however, the underlying infrastructure differs in minor details. Therefore, we listed the deployed components of the Monitoring System in two separate sections for each centre.

REFERENCES

[1] LEXIS Deliverable, *D3.2 Mid-Term infrastructure.*

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

[3] H. Xu, J. Coposky, D. Bedard and et al., "A Method for the Systematic Generation of Audit Logs in a Digital Preservation Environment and Its Experimental Implementation In a Production Ready System," *iPRES,* p. 201, 2015.