



European
Commission

Horizon 2020
European Union funding
for Research & Innovation

Big Data technologies and extreme-scale analytics



Multimodal Extreme Scale Data Analytics for Smart Cities Environments

D5.3: HPC infrastructure and resource management for audio-visual data analytics – initial version [†]

Abstract: The purpose of this deliverable is to present respective activities to analyse and establish properly the architecture for initial MARVEL services deployment. This document details resources provided by PSNC to the consortium within the scope of processing and storing data. Computing resources for running applications come from domains like HPC and Cloud where virtual machines with framework services are set-up according to pilot demands. Another crucial module of MARVEL's architecture is Data Corpus, which is intended to store and effectively process data in open science and open data manner. Moreover, the report presents organisational and operational aspects related to the procedures for accessing resources and their proposer use. Presented organisation of the infrastructure allocated for the purposes of the project is of a preliminary nature and will be adapted in the future in accordance with the needs reported by the users.

Contractual Date of Delivery	30/06/2022
Actual Date of Delivery	30/06/2022
Deliverable Security Class	Public
Editor	<i>Marcin Lawenda (PSNC)</i>
Contributors	PSNC, FORTH, IFAG, AU, ATOS, CNR, FBK, TAU, STS, ITML, GRN, ZELUS
Quality Assurance	<i>Manolis Falelakis (INTRA)</i> <i>Tomas Pariente Lobo (ATOS)</i>

[†] The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 957337.

The *MARVEL* Consortium

Part. No.	Participant organisation name	Participant Short Name	Role	Country
1	FOUNDATION FOR RESEARCH AND TECHNOLOGY HELLAS	FORTH	Coordinator	EL
2	INFINEON TECHNOLOGIES AG	IFAG	Principal Contractor	DE
3	AARHUS UNIVERSITET	AU	Principal Contractor	DK
4	ATOS SPAIN SA	ATOS	Principal Contractor	ES
5	CONSIGLIO NAZIONALE DELLE RICERCHE	CNR	Principal Contractor	IT
6	INTRASOFT INTERNATIONAL S.A.	INTRA	Principal Contractor	LU
7	FONDAZIONE BRUNO KESSLER	FBK	Principal Contractor	IT
8	AUDEERING GMBH	AUD	Principal Contractor	DE
9	TAMPERE UNIVERSITY	TAU	Principal Contractor	FI
10	PRIVANOVA SAS	PN	Principal Contractor	FR
11	SPHYNX TECHNOLOGY SOLUTIONS AG	STS	Principal Contractor	CH
12	COMUNE DI TRENTO	MT	Principal Contractor	IT
13	UNIVERZITET U NOVOM SADU FAKULTET TEHNICKIH NAUKA	UNS	Principal Contractor	RS
14	INFORMATION TECHNOLOGY FOR MARKET LEADERSHIP	ITML	Principal Contractor	EL
15	GREENROADS LIMITED	GRN	Principal Contractor	MT
16	ZELUS IKE	ZELUS	Principal Contractor	EL
17	INSTYTUT CHEMII BIOORGANICZNEJ POLSKIEJ AKADEMII NAUK	PSNC	Principal Contractor	PL

Document Revisions & Quality Assurance

Internal Reviewers

1. *Manolis Falelakis (INTRA)*
2. *Tomas Pariente Lobo (ATOS)*

Revisions

Version	Date	By	Overview
1.0	29/06/2022	Marcin Lawenda	PC and TC review comments addressed
0.8	28/06/2022	Despina Kopanaki	PC and TC review comments
0.8	24/06/2022	Marcin Lawenda	Comments from feedback implemented
0.7	23/06/2022	Manolis Falelakis Tomas Pariente Lobo	Feedback on revised document
0.7	22/06/2022	Marcin Lawenda	Revised document after internal reviews
0.6	16/06/2022	Manolis Falelakis Tomas Pariente Lobo	Review process
0.60	14/06/2022	Marcin Lawenda	First full draft
0.30	03/06/2022	All partners	Contributions to the documents
0.20	08/05/2022	Marcin Lawenda	Revised ToC
0.15	06/05/2022	Dragana Bajovic (STPM)	Comments on ToC
0.10	29/04/2022	Marcin Lawenda	Initial ToC.

Disclaimer

The work described in this document has been conducted within the MARVEL project. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 957337. This document does not reflect the opinion of the European Union, and the European Union is not responsible for any use that might be made of the information contained therein.

This document contains information that is proprietary to the MARVEL Consortium partners. Neither this document nor the information contained herein shall be used, duplicated or communicated by any means to any third party, in whole or in parts, except with prior written consent of the MARVEL Consortium.

Table of Contents

LIST OF TABLES.....	5
LIST OF FIGURES.....	6
LIST OF ABBREVIATIONS.....	7
EXECUTIVE SUMMARY	9
1 INTRODUCTION.....	10
1.1 PURPOSE AND SCOPE OF THIS DOCUMENT	10
1.2 INTENDED READERSHIP.....	10
1.3 CONTRIBUTION TO WP5 AND PROJECT OBJECTIVES.....	10
1.4 RELATION TO OTHER WPS AND DELIVERABLES	11
1.5 STRUCTURE OF THE DOCUMENT	11
2 HIGH PERFORMANCE COMPUTING SYSTEM	12
2.1 HARDWARE.....	12
2.1.1 <i>Eagle specification</i>	13
2.1.2 <i>HPC storage and tape library</i>	14
2.1.3 <i>Training cluster</i>	15
2.2 SOFTWARE	16
2.2.1 <i>Cluster Management Software</i>	17
2.2.2 <i>Available software</i>	19
2.2.3 <i>Module management</i>	20
2.2.4 <i>User Own software installation</i>	22
2.3 USER SERVICES AND SUPPORT.....	23
2.3.1 <i>Operation of the IT infrastructure</i>	23
2.3.2 <i>User services and support</i>	23
2.3.3 <i>Help Desk solution at PSNC</i>	23
2.3.4 <i>End user support</i>	23
2.3.5 <i>User support on security</i>	23
2.3.6 <i>Training centre</i>	24
3 CLOUD COMPUTING SYSTEM.....	25
3.1 CLOUD HARDWARE	25
3.2 CLOUD MANAGEMENT SYSTEM.....	26
3.3 CLOUD REQUIREMENTS.....	27
3.4 DELIVERED CLOUD RESOURCES	29
4 STORAGE AND DATA ANALYTICS SYSTEM.....	30
4.1 DATA PROCESSING INFRASTRUCTURE FRAMEWORKS	30
4.2 DATA CORPUS.....	31
5 GRANTING ACCESS.....	34
6 CONTRIBUTION TO MARVEL GOALS	36
7 CONCLUSIONS	37
BIBLIOGRAPHY.....	38

List of Tables

Table 1. EAGLE System Specification at PSNC.....	13
Table 2. EAGLE CPU nodes	14
Table 3. Eagle GPGPU nodes	14
Table 4. Eagle operating system	14
Table 5. PSNC Storage	14
Table 6. Configuration of training infrastructure.....	16
Table 7. Queue configuration on the Eagle supercomputer.....	17
Table 8. BST Cloud area specification	25
Table 9. DCW Cloud area specification	26
Table 10. Cloud infrastructure requirements for GRN3 use case	27
Table 11. Cloud infrastructure requirements for GRN4 use case	27
Table 12. Cloud infrastructure requirements for MT1 use case.....	28
Table 13. Cloud infrastructure requirements for MT3 use case.....	28
Table 14. Cloud infrastructure requirements for services common to all use cases.....	28
Table 15. Configuration of VMs delivered to the project.....	29
Table 16. Configuration of the Data Corpus.....	33

List of Figures

Figure 1. PSNC Server Room	13
Figure 2. PSNC data infrastructure and correlation with HPC	15
Figure 3. Topology of the training infrastructure.....	16
Figure 4. Topology of the Cloud Computing system.....	25
Figure 5. MARVEL multimodal Data Corpus-as-a-Service for smart cities.....	32
Figure 6. Data Corpus elements are deployed on a single VM.....	33
Figure 7. Information on MARVEL computational grant at PSNC HPC portal	35

DRAFT

List of Abbreviations

AI	Artificial Intelligence
API	Application Programming Interface
AV	Audio-Visual
AVCC	Audio Visual Crowd Counting
CI/CD	Continuous Integration / Continuous Delivery
DFB	Data Fusion Bus
DMP	Data Management Platform
DMT	Decision-Making Toolkit
DNS	Domain Name System
E2F2C	Edge-to-Fog-to-Cloud
EC	European Commission
FFMPEG	Fast Forward MPEG
GPGPU	General-Purpose computing on Graphics Processing Units
GUI	Graphical User Interface
HDFS	Hadoop Distributed Files System
HLS	HTTP Live Streaming
HPC	High Performance Computing
HTTPS	Hypertext Transfer Protocol Secure
IoT	Internet of Things
IP	Internet Protocol
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
MEMS	Micro Electro-Mechanical Systems
ML	Machine Learning
MPP	Massive Parallel Processing
MQTT	Message Queuing Telemetry Transport
MVP	Minimum Viable Product
OLAP	Online Analytical Processing
PANN	Pretrained Audio Neural Networks
PDM	Pulse Density Modulation
PFLOPS	Peta Floating-Point Operations Per Second
POE	Power Over Ethernet
RAM	Random Access Memory

REST	Representational State Transfer
RTSP	Real-Time Streaming Protocol
SED	Sound Event Detection
SNR	Signal-to-Noise Ratio
SSO	Single Sign-On
TB	Terabyte
TRL	Technology Readiness Level
UI	User Interface
UJ	User Journey
URL	Uniform Resource Locator
USB	Universal Serial Bus
UX	User experience
VCS	Version Control System
VM	Virtual Machine
WP	Work Package

DRAFT

Executive Summary

This deliverable highlights details of the infrastructure established for the MARVEL project. The document describes the hardware, software, and services offered by Supercomputing Centre of PSNC. The main focus is given on using the already available resources including the process of accessing these resources.

The HPC infrastructure requires specialised expertise to be used efficiently and thus, the deliverable provides the knowledge of mechanisms for accessing and using these resources at PSNC. Furthermore, to ensure seamless access to compute resources within the MARVEL project, and to carry out testing, integration, validation, and benchmarking procedures in an efficient manner, it is highly imperative to pay particular attention to the parameters of the offered infrastructure. The pilot and the pre-production testbeds are connected with specialised resources such as HPC and Cloud, including GPGPUs which enable taking full advantage of the AI-based algorithms for audio-visual analytics developed in WP3.

To facilitate the process of prototyping, a dedicated infrastructure was established which imitates HPC in respect of topology and job submission organisation, which contributes to reduced access time to tested resources.

Project partners conducted a thorough analysis against required cloud resources for use cases (pilots) in order to set up the infrastructure adjusted to specific demands.

Finally, the document reports on current achievements in establishing Data Corpus that is released as a service and enables smart cities to build and deploy innovative applications based on multimodal perception and intelligence.

1 Introduction

1.1 Purpose and scope of this document

The purpose of this document is to present the status of establishing the infrastructure at HPC and cloud levels for the implementation of the MARVEL framework. The infrastructure consists of computing and memory resources that are linked to specific resources such as HPC and dedicated cloud virtual machines. By enabling access to GPGPU, the capabilities of AI-based algorithms for audio-visual analytics developed in WP3 were fully exploited. In addition, the document discusses the aspects of planning and managing access to resources implemented in HPC, focusing on two access modes: batch and interactive. In addition, it explains how to use the pre-installed software using the functionality of the `module` and install user-owned applications. Then it describes what services PSNC offers to users and how to contact the support department in case of problems. Cloud computing services are presented in terms of their specification (hardware configuration) as well as determined for the project demands resources. Another goal was to present the possibility of data storage and analysis along with the status of Data Corpus implementation.

1.2 Intended readership

Deliverable D5.3 – “HPC infrastructure and resource management for audio-visual data analytics – initial version” is a public document that presents the infrastructure used for the deployment of the MARVEL framework.

The content found in this document aims to show to all stakeholders and potential users the HPC and the cloud infrastructure used for installing MARVEL services. Moreover, the resource management and rules are discussed along with their fulfilment which are required for the correct use of the offered resources.

Additionally, this document will serve as a guide to the upcoming integrated releases of the MARVEL framework, being a benchmark for expanding use cases, MARVEL-embedded components and technologies, and services offered.

1.3 Contribution to WP5 and project objectives

This deliverable has been composed within the context of ‘WP5 – Infrastructure Management and Integration’, and more specifically, it constitutes the first major output of ‘Task 5.1. HPC infrastructure’ and ‘Task 5.2. Resource management and optimised automatic usage of external computational and storage resources’. As part of the objectives of WP5, the document provides details of the provision and management of the HPC and cloud infrastructure for processing and storing data which are crucial to ensure foundation for high-throughput processing and storage resources. The configuration is casted to specific needs of MARVEL services to deliver fast complex analytics on the integration of high-level services and ML modelling.

The proposed solution offers sufficient capacity to efficiently host AI-based audio-visual analysis algorithms developed in WP3, ensuring high data processing efficiency in addition to high network and storage performance. Moreover, resources such as GPU cards required by dedicated services, the prototyping environment, along with high-end cloud-dense servers, provide the full advantage required by the R1 MARVEL integrated framework definition, released in January 2022. Services such as authentication, resource discovery, and task monitoring are offered as standard by the PSNC centre, supported by user services and administrative support. Together, these provide the environment necessary to carry out the software development process based on testing, integration, validation, and benchmarking.

1.4 Relation to other WPs and deliverables

This deliverable relies on the foundational work conducted within ‘WP1 – Setting the scene: Project setup’. More specifically, the selection of the Use Cases for demonstration draws from the detailed material on Use Case descriptions of deliverable ‘D1.2 – MARVEL’s Experimental protocol’. Additionally, deliverable ‘D1.3 – Architecture definition for MARVEL framework’ is an important source for this work, as it contains the refined architecture, which is the blueprint for this release, as well as subsequent releases. D1.3 also provides useful information that D5.3 builds upon, for example the description of available MARVEL components and their TRL, the grouping of components into building blocks that correspond to architectural layers, and the outline of integration processes that need to be applied.

The work defined within the context of ‘WP2 - MARVEL multimodal data Corpus-as-a-Service for smart cities’, ‘WP3 – AI-based distributed algorithms for multimodal perception and situational awareness’, and ‘WP4 - MARVEL E2F2C distributed ubiquitous computing framework’ stands as a reference for the definition of requirements of the infrastructure definition.

Within WP5, there has been a close collaboration with Task 5.3 and Task 5.4 with regards to the continuous integration of the framework realisation and progress against validated benchmarks, verifying the usefulness of the proposed solutions, respectively. This close dependency is expected to continue until the release of the final complete version of the MARVEL framework.

This deliverable will be a reference to the upcoming deliverable ‘D5.8 HPC infrastructure and resource management for audio-visual data analytics – final version’, where the last resource delivery and configuration will be presented.

1.5 Structure of the document

The structure of this document reflects the infrastructure components delivered to the MARVEL project: computational (HPC and Cloud) and storage (data storage and analysis) and is arranged as follows:

- Section 2 discusses the configuration of the HPC system offered to the MARVEL from three perspectives: hardware, software, and user services and support.
- Section 3 provides information about the cloud computing system where MARVEL components are deployed.
- Section 4 describes the system established for data analytics.
- Section 5 presents how granting procedures are organised at PSNC and what rules must be followed in order to get access to the offered resources.
- Section 6 links the results presented in previous sections to the overall MARVEL goals and objectives.
- Section 7 summarises and concludes this document.

2 High Performance Computing system

This section describes the infrastructure available at the Poznan Supercomputing and Networking Center (PSNC). This HPC centre is equipped with thousands of Central Processing Units (CPUs), a numerous of graphics cards for processing (GPGPU), storage disks required to provide Massive Parallel Processing (MPP) and, with this the required performance to the applications.

The goal of this elaboration is to better understand capabilities and thus set up the MARVEL infrastructure and provide seamless access to supercomputing resources for the management of the High Performance Computing and High Performance Data Analysis processes.

PSNC provides access to a world-class e-Infrastructure for the scientific community, a specific research and development environment – DIGITAL SCIENCE – for “proof of concept” projects, prototyping or large-scale pilot projects. Another branch of our activity is in DIGITAL INDUSTRY, which is focused on the creation of innovations based on Information and Communication Technologies (ICT), i.e., the execution of specific implementations for various fields of science and industry. Equally important is our work in the context of SOCIAL INNOVATION to spread knowledge and awareness of contemporary technological opportunities among various social groups, as well as actively fighting against digital exclusion. PSNC is a leading HPC centre in Poland with its 7.3 Pflops of computing power, 42 PB of online storage and data management infrastructure, providing direct support for the scientific communities in Poland as well as in Europe (e.g., Nuclear Fusion, Astrophysics, Bioinformatics, Chemistry, Nanotechnology). The PIONIER network is a nationwide broadband optical network that represents a base for research and development in the areas of information technology and telecommunications, computing sciences, applications and services for the Information Society. PIONIER connects 21 Academic Network Centres, Centres of Metropolitan Area Networks (MAN) and 5 of the HPC centres (participating as third parties in PRACE) using their own fibre connections.

2.1 Hardware

PSNC boasts a dedicated data centre (1600 m²) equipped with air and liquid cooling systems, video monitoring, fire protection systems and 24/7/365 monitoring. The total peak performance of ca. 96K CPU cores is over 7.3 PFlops. PSNC systems have appeared in the Top 500 list several times. PSNC is part of the European HPC infrastructure (PRACE), European and national grid infrastructure (EGI, PL-GRID), national (PLATON e- Science Platform, National Data Storage) and European Data Infrastructure (EUDAT).

PSNC hosts two dedicated data centres (1.600 m²) equipped with air and liquid cooling systems, video monitoring, fire protection systems and 24/7/365 monitoring. These include a primary, 2x 820 m², brand new data centre in the Berdychowo district of Poznan, delivered in the first quarter of 2015 (called “BST”) and a secondary data centre, in the Poznan central district. A view of the currently installed Petascale system is presented in Figure 1.



Figure 1. PSNC Server Room

The data centres are equipped with all the required environmental systems and components including redundant power supplies, own transformer stations, UPS, power generators, fire protection (detection and active extinguishing systems), security and access control systems and finally, 24/7 monitoring.

Furthermore, PSNC operates the modern backbone network based on 100 Gbps links between the major HPC sites and 100 Gbps links to bigger academic communities. It participates in the international academic networking consortium GÉANT, which offers multiple 100 Gbps to European countries. PSNC also peers with commercial providers in Poland through links in Poznan as well as with international operators through peering points in Hamburg and Amsterdam, amongst others.

In the following subsections, the resources made available for the implementation of project tasks will be discussed. This overview has been divided into three sections: HPC specs (supercomputer called Eagle), its disk and tape storage, and deployed training infrastructure for rapid prototyping.

2.1.1 Eagle specification

Eagle is an HPC cluster manufactured by Huawei composed of 1233 nodes of Intel Haswell and Broadwell CPUs. Nodes are equipped with 64GB – 256GB of RAM memory each and interconnected by a fast Infiniband FDR network. These resources are available to the MARVEL users in a shared mode considering the current machine load and demands of other users. Table 1 contains a detailed hardware specification of the EAGLE system at PSNC.

Table 1. EAGLE System Specification at PSNC

Nodes	Cores	Memory	Interconnect	Amount of Storage	Storage Performance
1.233	32.984	120.6 TB	56 Gb/s Infiniband FDR	4.6 PB Lustre 1 PB /home	32 GB/s Lustre File system 10 GB/s /home CNFS

EAGLE CPU Nodes

Table 2 contains the specification details of CPU nodes of the EAGLE system.

Table 2. EAGLE CPU nodes

Node merchant name	CPU model	No. of nodes	No. of processors and cores	RAM per node	Computing power per node	Node class tag
HUAWEI CH121 V3	Intel Xeon E5-2697 v3	589	2x14	64 GB	1.1 TFLOPS	intel,haswell,huawei
HUAWEI CH121 V3	Intel Xeon E5-2697 v3	530	2x14	128 GB	1.1 TFLOPS	intel,haswell,huawei
HUAWEI CH121 V3	Intel Xeon E5-2697 v3	59	2x14	256 GB	1.1 TFLOPS	intel,haswell,huawei
HUAWEI CH121 V4	Intel Xeon E5-2682 v4	55	2x16	128 GB	1.1 TFLOPS	intel,haswell,huawei

EAGLE GPU Nodes

Table 3 contains the specification details of GPU nodes of the EAGLE system.

Table 3. Eagle GPGPU nodes

CPU model	No. of nodes	No. of processors and cores	RAM per node	GPU model	No. graphic cards per node	Computing power per GPU
Xeon Gold 5115	3	2x10	92 GB	NVIDIA V100	2	7.8 TFLOPS

EAGLE Operating System

Table 4 contains the information about the EAGLE operating system, its domain name, and the task management system.

Table 4. Eagle operating system

Type	Domain Name	Task management system in batch mode
GNU Linux	eagle.man.poznan.pl	SLURM

2.1.2 HPC storage and tape library

This section provides information on the storage systems linked to the HPC. Similarly, as for HPC, presented resources are shared among different users so its efficiency depends on the current load generated by other users. The type and performance of the storage system at PSNC are delivered in Table 5.

Table 5. PSNC Storage

Storage	Storage Performance

LTO 5 3.5 PB	2 GB/s
Jaguar 28 PB	4.2 GB/s

In summary, a detailed view of the schematic infrastructure and their relations at PSNC can be obtained from Figure 2 below.

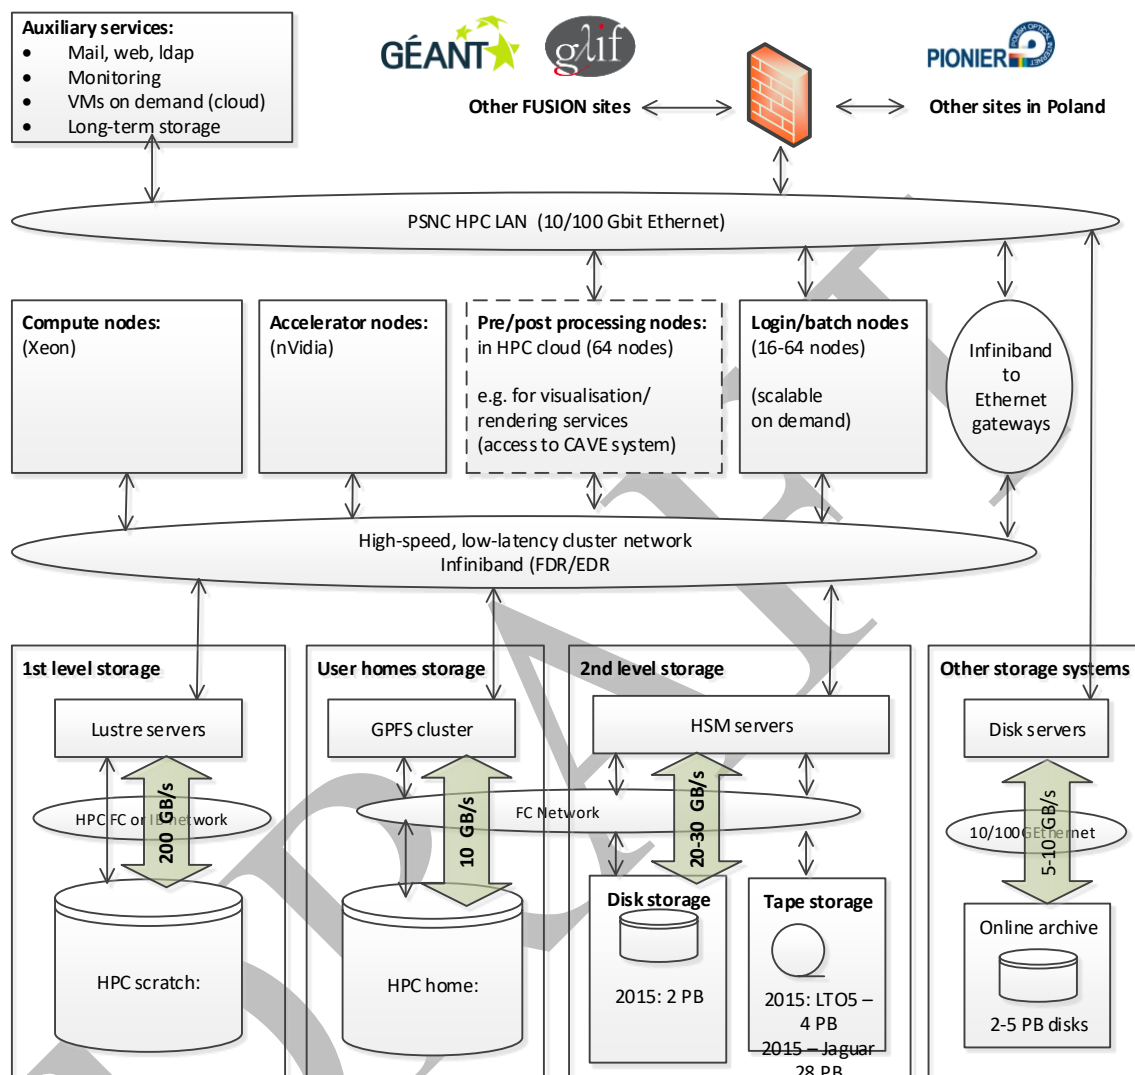


Figure 2. PSNC data infrastructure and correlation with HPC

The storage connected to the HPC is intended for storing short- and medium-term data while performing calculations. The short-term space is called SCRATCH and has a capacity of several PBs (shared by all users) used to store temporary calculation data. The medium-term space available through GPFS is called PROJECT DATA and is subject to the constraints of the available space. For MARVEL, the limit is currently set to 100 GB but can be increased if needed. The results of the calculations, which are not subject to further processing, are transferred for long-term storage as described in section 4.

2.1.3 Training cluster

MARVEL is using the Eagle cluster to run its applications, especially those which require access to GPGPUs. This requires the installation and configuration of many services that must work properly with the cluster system taking into account the requirements of the queue system.

This involves performing many tests, which are much easier to carry out in a prototyping environment.

PSNC also offers access to infrastructure for rapid software prototyping. It is important to realise that supercomputers in HPCs tend to be heavily loaded, making it compulsory to wait for the computation to run in the job queue. Having a smaller, multi-node computing cluster configured in a manner similar to the original supercomputer is essential to run jobs quickly without having to wait in a system queue. The proposed training infrastructure has a similar OS and queuing system, but only fewer cores and memory are available.

The configuration of the training cluster is detailed in the table below (Table 6).

Table 6. Configuration of training infrastructure

Name	Address	Cores	Memory	Disk space
Training0	62.3.171.192 sophora-192.man.poznan.pl	32	32 GB	40 GB HDD + 1000 GB volume
Training1	62.3.171.143 sophora-143.man.poznan.pl	32	32 GB	40 GB HDD
Training2	62.3.171.224 sophora-224.man.poznan.pl	32	32 GB	40 GB HDD
Training3	62.3.171.194 sophora-194.man.poznan.pl	32	32 GB	40 GB HDD
Training4	62.3.171.167 sophora-167.man.poznan.pl	32	32 GB	40 GB HDD

The topology of the training infrastructure setup for MARVEL prototyping is presented on Figure 3.

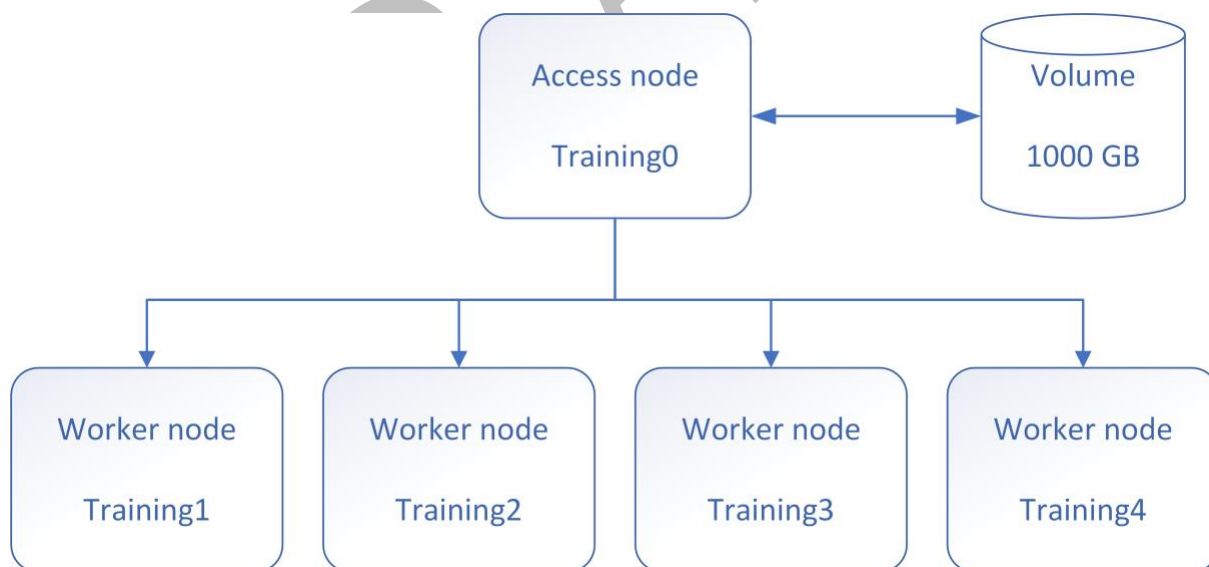


Figure 3. Topology of the training infrastructure

2.2 Software

This section describes the availability of software packages in PSNC for the MARVEL project. Software is an essential aspect in high performance computing as HPC applications differ little from standard computing applications except from the fact that HPC applications are developed

to run in parallel and require high performance, physical infrastructure with appropriate libraries and compilers to achieve the highest possible degree of efficiency and effectiveness.

Jobs can be submitted in two modes: batch mode, where application(s) are run in a series without manual interventions and interactive mode, where user interaction (presence) is mandatory in order to manipulate launch of programs, deciding on an ongoing basis to take the next processing steps. Batch mode is always preferred to run these applications in clusters, which will be discussed in detail in the following sections (2.2.1 - 2.2.4). However, this section will only discuss the available cluster management software packages and user application execution on a cluster available in PSNC.

2.2.1 Cluster Management Software

In this section, the process related to the batch system, job types, and job management is described. This knowledge is necessary to properly use HPC systems and avoid administration issues. The EAGLE server uses SLURM (SLURM workload Manager, 2022) as resource management system. A batch job can be submitted using the `sbatch` command (implemented by SLURM) along with its configuration files where the user can specify requirements for the application (e.g., memory needed, number of cores, etc.) and load necessary modules or stage files. Finally, the status of the submitted jobs can be accessed by using `sinfo`.

Queue

Running tasks through the SLURM queuing system takes place within queues, which differ in priority, limits and access rights. EAGLE has the queues presented in Table 7.

Table 7. Queue configuration on the Eagle supercomputer

Queue Name	Job Timeout	Default timeout job	Standard RAM	Comments
standard	7 days	1 day	2 GB	The queue is intended for performing tasks on one node
fast	1 hour	10 minutes	-	High priority; The maximum number of nodes 1
bigmem	7 days	1 day	2 GB	Queue for tasks requiring more than 128 GB per node. High priority but only 59 nodes belong to this partition
tesla	7 days	1 day	2 GB	Queue with GPU nodes

Submitting jobs

Prior to job submission, the appropriate command launching the application has to be embedded in the script to be correctly read by the queuing system:

```
/home/users/user/submit_script.sl
```

Example:

```
#!/bin/bash -l
#SBATCH -N 1
#SBATCH --mem 5000
#SBATCH --time=20:00:00
```

```
/path/to/binary/file.exe >/path/to/output/file.out
```

To submit the job to selected queue please use #SBATCH -p parameter:

```
#!/bin/bash -l
#SBATCH -N 1
#SBATCH --mem 5000
#SBATCH --time=20:00:00
#SBATCH -p long
/path/to/binary/file.exe >/path/to/output/file.out
```

Jobs can be submitted using sbatch command:

```
sbatch /home/users/user/submit_script.sl
```

Submitting interactive jobs

In this mode, interactive jobs can be submitted by executing the following command:

```
srun --pty /bin/bash or srun -u /bin/bash -i
```

The first command allocates a pseudo terminal that simplifies the work on a remote console. In case of any problems, please use only the second command.

Interactive jobs in graphic mode

From the user point of view, it is sufficient to log in to the cluster with the -X option

```
ssh -X and execute the following command:
srun-interactive -N 1 -n 1
```

Job submission using GPU cards

To submit a job to nodes equipped with GPU cards, it is required to use the tesla partition and add the following section to the submission script:

```
#SBATCH --gres=gpu:<no. of cards for every task>
```

An exemplary job using two cards should contain the following sections:

```
#SBATCH --gres=gpu:2
#SBATCH --partition=tesla
```

A selected number of applications are enabled to use GPUs, either by built-in functionality or using a dedicated module, usually containing "CUDA" in its name.

```

namd/2.10-ibverbs-smp-cuda    <- GPU supported version
namd/2.10-multicore(default)
namd/2.10-multicore-cuda      <- GPU supported version
namd/2.10-multicore-mic
namd/2.12-ibverbs
namd/2.12-ibverbs-smp
namd/2.12-ibverbs-smp-CUDA    <- GPU supported version

```

Checking status of the queue

To check what jobs have been submitted by a given user, please execute the command:

```
squeue -u username
```

Removing jobs

To remove the job, please use the `scancel` command and specify the corresponding job id as a parameter. Both, waiting and running jobs can be removed.

```
scancel job_id
```

2.2.2 Available software

Description of the software available to users at the HPC system. A complete documentation for installed software along with sample scripts to make use of them can be found on the website (PSNC Eagle software, 2022).

Name	Version(s)	Purpose
GCC	4.8, 4.9	Development/Compiler
Intel Parallel Studio XE Cluster Edition:	15.0.0, 15.0.3, 16.0.0	Development/Compiler
Intel C++ Compiler		Development/Compiler
Intel Fortran compiler		Acceleration Library
Intel Data Analytics		Math Library
Intel MKL (C++, Fortran) Intel TBB (C++)		Threading Library
Intel IPP (C++)		Media and Data Library
Intel Advisor (C++, Fortran) Intel Inspector (C++, Fortran)		Vectorisation/Optimisation
Intel VTune Amplifier		Debugging Tools
Intel MPI (C++, Fortran)		Performance Profiler
Intel ITAC (C++, Fortran) IMSL (Fortran)		MPI Library
		MPI Analyser and Profiler
		Numeric Library
AMD Core Math Library	5.3.1	Development/Libraries
Abaqus	6.14-2	Application/Engineering

Abinit	7.4.3	Application/Science
Amber	14	Application/Science
Bowtie	1.0	Application/Science
Codeanalyst	3.4.18	Application/Science
Cudatoolkit	7.0.28	Application/Science
Gaussian	09.D.01	Application/Science
Gromacs	5.0.4	Application/Science
Hmmer	3.1b1	Application/Science
Mapdamage	2.0	Application/Science
Matlab	R2013a	Application/Science
Mumax	3.8	Application/Science
Namd	2.10	Application/Science
Orca	3.0.3	Application/Science
Plink	1.0.9	Application/Science
Quantum Espresso	5.2.0	Application/Science
Rna-Seqc	1.1.8	Application/Science
Rsem	1.2.18	Application/Science
Siesta	3.2	Application/Science
Tabix	0.2.6	Application/Science
Tophat	2.0.13	Application/Science
Trinityrnaseq	2.0.6	Application/Science
VcfTools	0.1.14	Application/Science
Velvet	1.2.10	Application/Science
Vowal-wabbit	8.0	Application/Science
Xenome	1.0.1r	Application/Science

2.2.3 Module management

This section discusses how to manage modules at the HPC and which are the standard methods used to set up the user environment for clusters. All available compilers and libraries can be listed by typing the module available command. Every user of the systems at PSNC is allowed to compile and run their own codes. It is also possible to compile and run third-party software using either GCC or Intel compilers (Intel oneAPI DPC++/C++ Compiler, n.d.). Available compilers, applications and libraries (e.g., MKL) can be listed using the module list command. To make a module active, one has to issue the `module load <module name>` command. If the compiled application is using any library that needs to be loaded during the compilation, it is required to load these modules before the actual program is executed.

Every application that is storing intermediate data during computation or is using large datasets, needs to copy the data to the `/mnt/lustre` directory. This directory is shared among all nodes that are used for any application computation. Since this is a temporary space that can be erased at any time, output data should be saved to the user's `/home` directory afterwards.

All applications must be launched using the queue system. Applications executed on the login node, manually on the nodes or applications that use the /home directory for storing intermediate results will be terminated without warning.

Modules

Modules are the standard methods used to manage the user environment for clusters. It is enough for the user to load a module associated with a given application to set all required environment variables (`PATH`, `LD_LIBRARY_PATH` `INCLUDE`, etc.) accordingly. If the user wants to use a different version of the application, he/she has to unload the current-in-use module and load the module associated with the new version.

Usage

<code>module list</code>	To list loaded modules please use this command
<code>module avail</code>	To list modules available for a given user please use this command
<code>module help <module_name></code>	Additional info about the module can be obtained by executing this command
<code>module load <module_name></code>	Loading a module is done using this command
<code>module unload <module_name></code>	A module can be unloaded this way
<code>module show <module_name></code>	Checking for environment variables modified by given module is available via this command

Settings customisation

In the modules environment, a user can decide which modules will be loaded when logging into the system. To do this, one has to edit the file `~/.bashrc` or `~/.bash_profile` configuration and add the appropriate lines. Below is an example where a variable `ModulePath` has been added to the directory, which contains the user module and the module `open64` compiler in version 4.2.5.2.

```
[username@hostname ~]$ cat ~/.bash_profile
...
# Environment Modules
# Prepend directory to the MODULEPATH environment variable.
module use --append $HOME/.modules/my_modules
# Modules loaded at login time.
module load open64-4.2.5.2
...
```

2.2.4 User Own software installation

Users are allowed to install their own software in the \$HOME directory. This section explains how to properly use already delivered modules containing required software, like libraries, compilers, interpreters, and other supportive applications and tools. The installation procedure of the own software is presented in the subsequent section.

Install own software

It is important to note that own software compilation on the head node is forbidden. It is possible to install own software in the /home directory without asking administrators for permission. Users can also ask the administrators to compile and install any kind of software by writing a message to the standard support email. In the consecutive part of the subsection, we discuss how it should be done correctly.

Software Compilation

To compile the software on HPC system, we need to access the interactive console. To do so, a user must submit an interactive job and compile it in an interactive mode in order to utilise the resources of the remote node. Although the head node provides resources as well, these resources are limited and shared amongst all the users. Consequently, the compilation of an application on the head node impacts other users and shall be avoided in any sense.

<code>srun --pty /bin/bash</code>	Example for using SLURM
<code>srun --pty -n 12 /bin/bash</code>	To access a node with a certain number of cores a user needs to specify -n parameter
<code>srun --pty -n 12 --mem 16000 /bin/bash</code>	If a user wants to reserve a certain amount of RAM, the --mem parameter needs to be used
<code>module load gcc</code>	When the user accesses the console directly at the node, the modules for the compilation need to be loaded
<code>./configure --prefix=home_folder</code> <code>make</code> <code>make install</code>	The last step is to compile the software in the /home directory, according to the manufacturer's recommendations
<code>make -j 12</code>	In order to utilise compilation processes on several cores, the user can add the -j parameter

Interactive Job

An interactive job allows a user to obtain a shell on a computational node. Because compiling and running any programs on the head node is not allowed, interactive jobs are one way to run graphical user interface software (e.g., Matlab, Maple, etc.) or compile a particular application. The following is an example of how to run a task in interactive mode.

<code>srun --pty /bin/bash</code>	Interactive job submission
-----------------------------------	----------------------------

<code>srun --pty -n 12 /bin/bash</code>	A user has to add the <code>-n</code> parameter to access the node with the specified number of cores
<code>srun --pty -n 12 --mem 16000 /bin/bash</code>	If a user wants to reserve a certain amount of RAM, the <code>--mem</code> parameter needs to be used

2.3 User Services and Support

Efficient cooperation of MARVEL pilots with infrastructure technical support is essential to fully exploit the potential of available resources. PSNC as a hosting centre and national network operator has teams of people that can fully support all aspects of computing (especially HPC) and data infrastructures. Below are some areas in which this cooperation can be established using the synergy effect.

2.3.1 Operation of the IT infrastructure

PSNC is hosting HPC machines for more than 25 years and has an experienced team responsible for both, physical and software maintenance for various types of systems (clusters, GPU machines, and large SMP systems). Currently, the scope of the support covers all issues related to the computing infrastructures, including troubleshooting, compilation, and runtime problems. MARVEL deploys its services on novelty, demanding, and reliable machines which are managed by an experienced team of professionals following established rules.

2.3.2 User services and support

Due to the fact that MARVEL requires specific configurations for its services that fully use the current technical solutions, support and immediate response to emerging problems are expected, which is an indispensable part of contacting the support department. In addition to the infrastructure mentioned in the previous sections, PSNC offers a set of services and support activities intended to assist effective, reliable, and safe usage of the PSNC infrastructure. The services contain in particular user support for daily operations as well as support for the optimisation and efficient usage of the PSNC computing clusters.

2.3.3 Help Desk solution at PSNC

PSNC has supported users with the utility of achieving fast response times on their problems coming from distinct areas of computational infrastructure for years. The Network Operation Center (NOC) is the “Operator on Duty” on a 24/7 basis. It reacts in emergency situations and can provide tentative support nearly online. However, all MARVEL user queries and problems must be registered in the PSNC Ticketing System (TTS).

2.3.4 End user support

End User Support corresponds to the service for the user application enhancement in a broader sense. It mainly consists of improving their codes. The permanent support has been addressed to all MARVEL users having HPC accounts in PSNC.

2.3.5 User support on security

The user becomes a critical link in the security chain of every system or infrastructure. Even an optimally protected infrastructure may be endangered if its users tend to behave in an inaccurate way. Thus, it is proposed that activities be undertaken focused on supporting MARVEL users involved in the HPC fusion infrastructure in the security area. The majority of the activities will be performed during the whole project lifetime. They will be focused also on users of the

infrastructure who may not have Information Technology security skills or even lack of general domain knowledge.

2.3.6 Training centre

Since 2005, PSNC has organised training sessions for HPC users coming from different research areas. During approximately 100 trainings, most important thematic groups have been identified: sequential programming, parallel programming or queuing system usage, just to name a few. The training offer is supported by qualified HPC staff of PSNC and external specialists from the Technical University of Poznan.

As already described above, PSNC offers a wide range of services in order to support its users. Users are also guided by using the infrastructures, tools, and services provided within MARVEL while additional implementation support is given by technical partners. Also, consortium members and users of the infrastructure that MARVEL provides, can take advantage of the rich portfolio of training provided by PSNC such as parallel programming, scientific visualisation, machine and deep learning, or performance optimisation, to name a few. The full list of all HPC training is reachable on the website (PSNC training courses, 2022).

3 Cloud Computing system

PSNC Cloud computing system is located in two locations:

- BST which is composed of 170 servers, 35TB memory and 2.7 PB of storage space
- DCW which is composed of 25 servers, 10TB memory and 0.7PB of storage space

Cloud storage is designed to store the data needed for ongoing computing and is not intended to serve as a long-term service. Since the target of MARVEL is to store approximately 3.3 PB of data, this data will be stored in dedicated systems described in section 4.

These locations are interconnected with a fast backbone network with a bandwidth of 400Gb. The topology of the cloud is presented in Figure 4.




Figure 4. Topology of the Cloud Computing system

3.1 Cloud hardware

This subsection details the specification of the equipment included in the cloud system.

Table 8. BST Cloud area specification

Availability zone	Processor	Cores	Memory [GB]	Backbone network [Gbit]	No. of servers
Inula-1	Xeon(R) CPU E5-2697	56	192	Infiniband FDR (56)	48
Inula-2	Xeon(R) CPU E5-2697	56	192	Infiniband FDR (56)	48
Inula-3	Xeon(R) CPU E5-2697	56	256	Infiniband FDR (56)	48
Supermicro-2015		40	256	Ethernet 10 Gbit	14
		32	256	Ethernet 10 Gbit	4
quanta		32	96	Ethernet 10 Gbit	5
Total		8752	35328		

Table 9. DCW Cloud area specification

Availability zone	Processor	Cores	Memory [GB]	Backbone network [Gbit]	No. of servers
Klaudia	Xeon Gold 6138	80	384	Ethernet 10 Gbit	6
	AMD EPYC 7451	96	512	Ethernet 10 Gbit	14
Total		1824	9472		

3.2 Cloud management system

Cloud computing is becoming more and more popular not only for business but also in scientific applications. While computing power-intensive tasks are still mainly performed on supercomputers, analytics, data visualisation, database engines, and web page services require a runtime environment that is primarily flexible and reliable. For this type of purpose, PSNC has a cloud environment that allows you to run virtual machines with Linux (Redhat, Centos, Ubuntu, Debian, etc.) and Windows. The system, operating under the control of the OpenStack system, connects computing and application servers with data storage solutions and the network. Thanks to the use of cloud technologies, services and applications running on virtual machines become resistant to hardware failures - in the event of problems, the machines are automatically started on redundant servers. An additional advantage is the ability to adjust the amount of available resources to the changing needs of the application.

PSNC provides:

- **Flexible virtual machines running for network services under the control of the selected operating system.** If your machine is to handle network services such as e-mail, websites, etc., choosing such a machine is the best solution.
- **Computing machines in a cloud environment.** Despite the fact that the greatest tasks are counted on supercomputers, the requirements of some applications (e.g., Windows application) make it impossible to use the cluster. In such cases, you can use the fact that the servers used to run virtual machines have 10-20 times greater performance than even the fastest laptops. Thanks to the appropriate configuration of the cloud environment, the virtual machine will have a performance similar to the physical server with a given configuration - the physical resources of a given server will be dedicated to this virtual machine.
- **Private clouds.** Instead of a single machine, it is possible to create a private environment in which, within the available resources (processors, memory, disk), it is possible to independently create many virtual machines with configurations corresponding to a specific application. For this type of application, it is also possible to define network configuration between machines, firewall rules both for the entire environment and for individual machines, and the use of advanced mechanisms such as load balancers.
- **Cloud data storage and management systems** - including classic matrix systems and cluster file systems integrated with the computing and service cloud environment as well as object-oriented data storage systems (S3, Ceph and Swift).

PSNC Cloud computing system offers access to the following services:

- virtual machines
- private cloud (VPC)
- virtual disks
- object Storage

- shared storage
- API
- orchestration
- load balancers
- VPN
- firewall
- control panel
- Linux OS
- Microsoft SPLA

3.3 Cloud requirements

Project partners conducted a thorough analysis against required cloud resources for use cases in order to set up infrastructure adjusted to specific demands. The result of this analysis is given in the tables below (Table 10 - Table 14) broken down to specific pilots.

Note that some of the components listed in Table 14 run in a single instance in the cloud for all the use cases, such as DFB, DatAna Cloud, DatAna Registry, HDD, SmartViz, MARVDash, and MARVEL Data Corpus. On the other hand, the inference models (Audio Tagging, AVAD, AVCC, VCC, SED, FedL, ViAD) may require one or more instances running in the cloud for each use case, depending on the number of data sources (cameras or microphones) present in the use case.

The MARVEL M18 use cases for which the resource requirements are analysed in the following tables are:

- GRN3 - Traffic Conditions and Anomalous Events
- GRN4 - Junction Traffic Trajectory Collection
- MT1 - Monitoring of Crowded Areas
- MT3 - Monitoring of Parking Places

Table 10. Cloud infrastructure requirements for GRN3 use case

Partner	Component name	Required CPU cores	Required RAM (GB)	Required Storage (GB)	Operation	GPU Required
AU	Audio-Visual anomaly detection (AVAD)	12	48	30	Real-time	Yes
TAU	Audio Tagging	12	24	12	Real-time	Yes
UNS	FedL	2	2	1	Asynchronous	Yes
TOTAL		26	74	43		

Table 11. Cloud infrastructure requirements for GRN4 use case

Partner	Component name	Required CPU cores	Required RAM (GB)	Required Storage (GB)	Operation	GPU Required
AU	Audio-Visual crowd counting (AVCC)	8	32	10	Real-time	Yes
TAU	Sound event detection (SED)	12	24	12	Real-time	Yes
CNR	DynHP	32	64	500	Asynchronous	Yes
UNS	FedL	2	2	1	Asynchronous	No

TOTAL	54	122	523		
--------------	-----------	------------	------------	--	--

Table 12. Cloud infrastructure requirements for MT1 use case

Partner	Component name	Required CPU cores	Required RAM (GB)	Required Storage (GB)	Operation	GPU Required
AU	Visual anomaly detection (ViAD)	4	16	10	Real-time	Yes
AU	Visual crowd counting (VCC)	8	32	10	Real-time	Yes
CNR	DynHP	32	64	1000	Asynchronous	Yes
UNS	FedL	2	2	1	Asynchronous	No
TOTAL		46	114	1021		

Table 13. Cloud infrastructure requirements for MT3 use case

Partner	Component name	Required CPU cores	Required RAM (GB)	Required Storage (GB)	Operation	GPU Required
AU	Audio-Visual anomaly detection (AVAD)	8	32	20	Real-time	Yes
TAU	Sound event detection (SED)	8	16	8	Real-time	Yes
TAU	Audio Tagging	8	16	8	Real-time	Yes
UNS	FedL	2	2	1	Asynchronous	No
TOTAL		26	66	37		

Table 14. Cloud infrastructure requirements for services common to all use cases

Partner	Component name	Required CPU cores	Required RAM (GB)	Required Storage (GB)	Operation	GPU Required
ITML	Data Fusion Bus (DFB)	24	128	5000	Real-time	No
ATOS	DatAna Cloud	12	32	1000	Real-time	No
ATOS	DatAna Registry	1	2	1	Real-time	No
CNR	Hierarchical Data Distribution (HDD)	4	4	10	Asynchronous	No
ZELUS	SmartViz	4	8	60	Real-time	No
FORTH	MARVdash	16	32	1000		No
STS	MARVEL Data corpus as a service	16	32	100000	Real-time	No
TOTAL		77	238	107071		

It was finally decided and established that initially all project components will be deployed on one thick Virtual Machine composed of 32 CPUs and 64GB of RAM. Currently, an incremental approach has been used to test existing solutions and observe their resource requirements. The most resource-intensive services will be placed on dedicated machines.

3.4 Delivered cloud resources

The table below presents the configuration of the VMs delivered to meet the project demands.

Table 15. Configuration of VMs delivered to the project

No.	Name	# of VCPUs	Memory	Volume	Operating system
1.	acme-dns	2	4GB	50GiB	Ubuntu Server 20.04 LTS
2.	gitlab	4	8GB	500GiB	Ubuntu Server 20.04 LTS
3.	edgesec_vpn_supernode	2	4GB	50GiB	Ubuntu Server 20.04 LTS
4.	marvel	32	64GB	2048GiB	Ubuntu Server 20.04 LTS
5.	edge_sec_vpn_gw	2	4GB	50GiB	Ubuntu Server 20.04 LTS
6.	cloud2	16	32GB	100GiB	Ubuntu Server 20.04 LTS

The remaining resources are delivered to establish servers for the MARVEL Data Corpus as reported in section 4.2.

4 Storage and Data analytics system

Currently, the PSNC mass storage infrastructure consists of several types of systems, including mid-size disk arrays, high-performance disk arrays, software-defined disk servers, as well as specialised systems, including an efficient file server and SSD arrays. The disk system mainly consists of NetApp E5600 series storage arrays, typically with 120 x 4 TB each. They are connected to Brocade 16Gbit / s switches (96 ports/switches). HPC arrays are DDN SFA4k12x with 420 x 4 TB drives per array. Attached to them are 4 OSS Luster servers, as well as a cluster of Luster MDS servers. File servers include EMC Isilon 5x x410 with 10Gbit Ethernet interfaces. SSD arrays include EMC XtremeIO and NetApp EF500 arrays with FC 8Gbit/s and FC 16Gbit/s / IB QDR interfaces. PSNC also supports a cluster of disk servers that run OpenStack Swift with Swift and S3 interfaces, and offer RADOS and S3 as well as RBD standards. PSNC offers 42PB of online storage and data management infrastructure, providing direct support for the scientific communities in Poland as well as in Europe. Moreover, PSNC is a leader of a Polish initiative named National Data Storage - KMD (Krajowy Magazyn Danych). By the end of 2024, the infrastructure created by the KMD members will increase the capacity of the data storage infrastructure by 200 petabytes and the capacity of the tape space by 180 petabytes, and the DataLake model used in its construction will ensure its high flexibility. In addition, as a result of the project, by 2024, cost-effective services for basic data storage and management processes, as well as their processing and analytics will be provided. These resources will be partially available also to the MARVEL users in accordance with the provisions of the GA and the requirements of individual pilot applications.

4.1 Data processing infrastructure frameworks

Thanks to the growing demand for processing huge amounts of data, High Performance Data Analytics (HPDA) has become more and more popular. HPDA brings together HPC (or other dedicated infrastructure) with Big Data to efficiently analyse large-scale, operational and static data together with multiple streaming data sources. This combination turns out to be quite advantageous when volume, velocity, and variety requirements surpass the capabilities of each domain's state-of-the-art solutions.

Performing calculations would be supported by the use of frameworks that will allow for scaling the implementation both in the context of data and the infrastructure used. There are many solutions of this type on the market, adapted to the problem being solved and the preferred programming language.

In the MARVEL project, solutions from the Apache Kafka framework (Kafka platform, n.d.) are used allowing for distributed event store and stream-processing. It is important to the project goals that Kafka offers a unified, high-throughput, and low-latency platform enabling real-time data feeds as well as facilitates connecting external systems (Kafka Connect) and stream processing via a set of libraries (Kafka Streams).

Automation of the flow of data between software systems which conspicuously includes the ability to operate within clusters is crucial for MARVEL demands and that is why Apache Nifi (Apache NiFi system, n.d.) is used for this purpose. This highly configurable and secure system supports powerful and scalable directed graphs of data routing, transformation, and system mediation logic. It also offers features that are very desirable by the specificity of MARVEL applications like: loss tolerant, low latency, dynamic prioritisation, and modification of the flow at the runtime.

Elasticsearch (Elasticsearch, n.d.) is a big data framework developed in Java gaining popularity since it addresses some of the implementation shortcomings of other platforms and provides better real-time processing support. These shortcomings mainly relate to high memory requirements and the micro-batching nature of its steaming engine. Elasticsearch can be used to search various kinds of documents and perform and combine many types of searches: structured, unstructured, geo, and metric. It facilitates using distributed architecture by supporting multitenancy where each node hosts a portion of data, coordinating and delegating operations to the correct locations. This approach assures automatic rebalancing and routing for scalable, near real-time search. It is often used in conjunction with other approaches such as Logstash (the data collection and log-parsing engine) and Kibana (the analytics and visualisation platform).

Recently, PyTorch has become more and more popular amongst data analysts due to its support of the Python language (C++ enabled). It is an open-source programming framework often used to model machine learning problems due to the ease of implementation supported by a number of libraries addressing the most recent applications. Specifically, PyTorch delivers to the programmers high-level features like Tensor computing (like NumPy) with strong acceleration via graphics processing units (GPU) and deep neural networks built on a tape-based automatic differentiation system.

In order to fully utilise research efforts to accelerate Big Data frameworks, heterogeneous hardware, such as GPGPUs and FPGAs, must be considered. The process of transformation of GPGPU can be facilitated by automated procedures like HadoopCL which directly generates OpenCL kernels for GPGPU execution. Another technique is source-to-source translation into CUDA, source-to-source translation into OpenCL, and mixed-source development (CUDA via the Java Native Interface) respectively enabled by HeteroDooop, Glasswing, and HeteroSpark. Microsoft developed Dandelion, a unified compiler and runtime system for accelerating C# and F# workloads on heterogeneous resources such as GPUs and FPGAs. Mars is a MapReduce framework that is optimised for execution on GPUs written in C++. MrPhi is a MapReduce framework optimised for Intel Xeon Phi coprocessors.

4.2 Data Corpus

MARVEL's Data Corpus is aimed to prioritise strengthening of open science and open data through enriching research and innovation contributing to benchmarking of edge, fog, cloud, and AI technologies. It is released as a service and enables smart cities to build and deploy innovative applications that are based on multimodal perception and intelligence (Pillar I of the project).

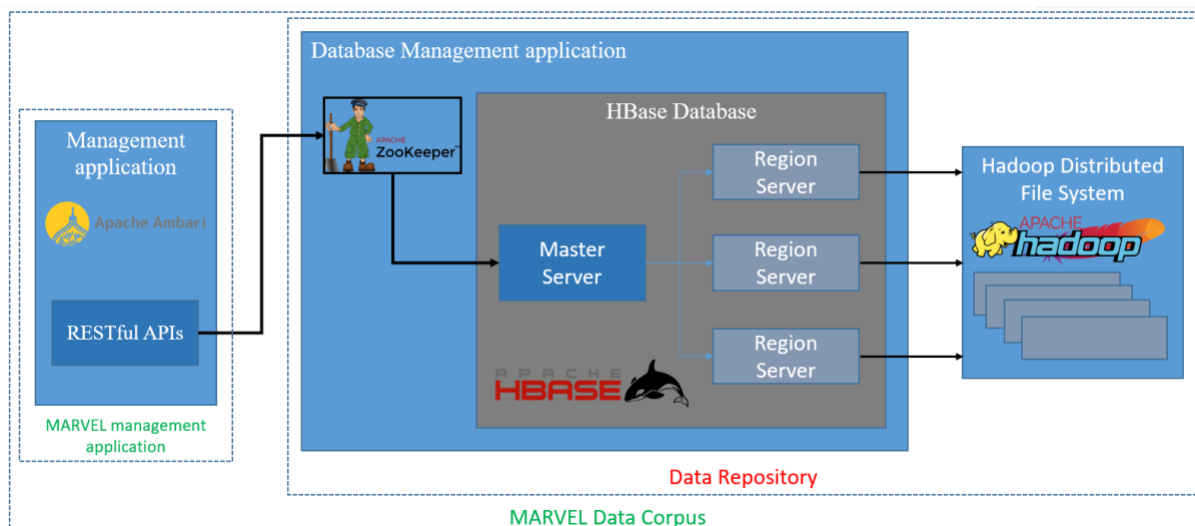


Figure 5. MARVEL multimodal Data Corpus-as-a-Service for smart cities

The following assumptions are defined for Data Corpus architecture:

- One Region Server for each pilot
- One Region Server for open datasets
- Each Region Server talks with a MasterNode on HDFS
- Each MasterNode splits the data into portions (slaves), depending on the size of the data
- Zookeeper manages the Hbase
- Queries (accessing) are made through Ambari

Based on the decisions taken under discussion among project partners (requirement analysis), resources for Data Corpus rely on virtualised infrastructure. A specific number of VMs were allocated to deploy the Hadoop HBase cluster. The storage backends for these VMs were improved in order to meet capacity and high throughput parameters by using the CEPH cluster with significant volumes.

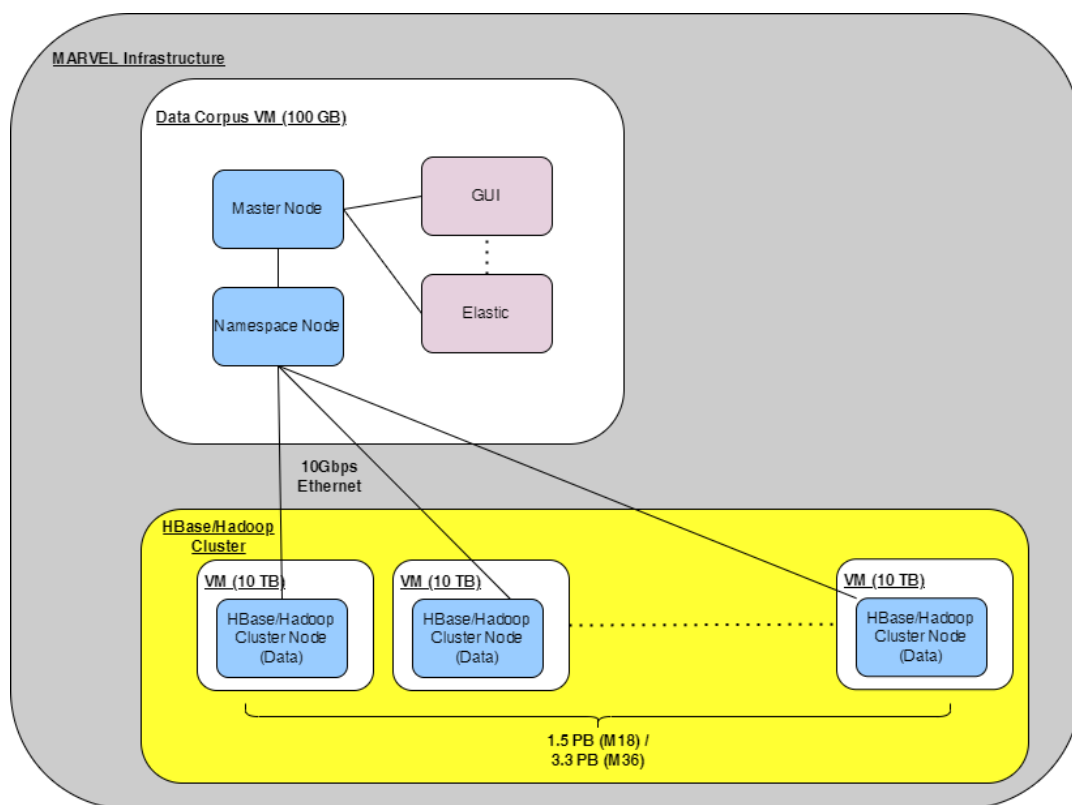


Figure 6. Data Corpus elements are deployed on a single VM

The current installation comprises of VMs presented in Table 16.

Table 16. Configuration of the Data Corpus

No.	# of VCPUs	Memory	Volume	Operating system
7.	16	32GB	1024GiB	Ubuntu Server 20.04 LTS
8.	8	16GB	100GiB	Ubuntu Server 20.04 LTS
9.	8	16GB	100GiB	Ubuntu Server 20.04 LTS
10.	8	16GB	100GiB	Ubuntu Server 20.04 LTS

It is planned for the future (around M24) to launch a new CEPH cluster which will be faster (better IO and streaming performance) and provides more scalability in terms of capacity. The cluster’s volume will be increased periodically until we reach the required repository size.

5 Granting access

A unified process for requesting HPC resources has been defined and implemented at PSNC. The domestic user represented by Principal Investigator, must submit the official proposal (computational grant) to the Council of Users. The Council of Users is responsible for the review process and takes the decision in terms of granting the access. The access rules are regulated and expressed in the “Rules of the usage of PSNC’s computational resources”. Summing up, the proposal is evaluated, and then the request for account creation is passed to the HPC Operational Department. Each year, a report on the computations and scientific results is expected from the Principal Investigator. The maximal computational grant duration is three years. In summary, the Principal Investigator needs to provide the following information on their grant:

1. Personal data with their scientific portfolio;
2. The title and description of the work foreseen to be processed;
3. Core time requirements;
4. Storage space requirements;
5. The list of the grant’s stakeholders (users).

At present, the whole grant management is processed in an electronic manner in the portal developed by PSNC: (PSNC HPC portal, 2022) (see Figure 7). The system is responsible for:

1. Gathering the data required to fill up application forms;
2. Management of the groups of users associated with the grant proposal;
3. User accounts and their accounting;
4. Assigned limits guarding;
5. Reporting process and the evaluation of its results by the Council of Users body;
6. Maintaining and presenting user jobs’ data in a smart manner;
7. Generating the yearly utilisation reports on demand with the multi-criteria options.

The portal is developed using an Open Source software, the code of which is still under development and includes continuous updates. Finally, the systems can be accessed via SSH:

```
user@host:~$ ssh pncusername@eagle.man.poznan.pl
```

As it was already mentioned, the work on EAGLE is organised around computational grants. To facilitate the project work, one of them has been established: "MARVEL grant" - Grant no. 505. Of course, it illustrates only logical organisation which follows pilots’-oriented approach in the MARVEL project. Other applications and cases, which are under investigation (e.g., Artificial Intelligence, data analytics, pre- or post-processing, and other supporting tools) will be assigned based on either technical coupling or balancing reasons.

In order to use EAGLE resources, each user must have an account registered at the HPC portal (PSNC HPC portal, 2022). The portal provides overall information about computational grants (validity, limits, machine allocation), consumed resources and publications presenting the results of computations. Each user account is assigned to the specific computational grant as well as the HPC cluster. Moreover, the portal allows to set up a password for the machine.

Grant					
Id	505				
Subject	Grant dla projektu Marvel				
Funding	The Framework Programme of the European Union				
Description	file				
Abstract	file				
Keywords describing computational techniques	AI, video stream processing				
Keywords describing issues that will be solved in the project	AI, video stream processing				
Typ grantu	pelny				
Date of grant acceptance	2021-04-08				
Data ostatniej modyfikacji grantu	2022-06-08				
Grant duration [months]	36				
Data zakończenia grantu	2024-04-08				
Priority	4				
Status	approved				
Confirmation of compliance by the secretariat	YES				
Comment					
Applying for hardware and software resources					
	Maximum CPU time usage in the project [hours]*	The total disk space usage[GB]	The name of the standard software used	User own applications	Additional requirements
klaster EAGLE	47334 / 1000000	(b.d) / 100		NIE	
Maszyny wirtualne	0 / 0	(b.d) / 0		NIE	
The required capacity of tape resources [GBs]			0 GB, What data will be archived:		
* - used/to be used					
(b.d) - no details					
show details regarding resources usage...					
Reports					
Start of reporting period		End of reporting period		Report	Details
no request regarding reports					

Figure 7. Information on MARVEL computational grant at PSNC HPC portal

Access to the Virtual Machines is granted based on RSA public keys of the user. Operations of the user account creation and importing the public keys are handled by the Virtual Machine administrator. The administrator is also responsible for server maintenance and the installation of all necessary software.

6 Contribution to MARVEL goals

The overarching goal of MARVEL is to deliver an Edge-to-Fog-to-Cloud (E2F2C) framework based on reliable infrastructure that operates in real-world environments, processing large volumes of captured AV data, enabled by multi-modal perception and intelligence. More specifically, with regards to the project objectives analysed in the DoA, the delivery of the reliable infrastructure is related to “Objective 3” which states about distributed and secure Edge-to-Fog-to-Cloud (E2F2C) computing framework for processing big data and IoT applications as well as to “Objective 5” that concentrates on creating the MARVEL Data Corpus-as-a-service contributing to BDVA standards.

The HPC, Cloud and storage infrastructure established in a manner ensuring high-level operability, directly contributes to the achievement of the above-mentioned project objectives. This contribution can be considered as a "setting milestones" towards the upcoming complete versions of the MARVEL environment in terms of infrastructure underpinning the processing of services and data volumes. This provides a stable and efficient setting for the foundation of the services developed under the other work packages of the project.

The presented hardware configurations constitute the initial proposal that meets the current requirements of design applications. As part of the project development, it should be expected that the resources made available so far will be expanded, addressing the growing demand for computing power and data space.

7 Conclusions

This document is the official output of the Work Package 5 of the MARVEL project. The work that is listed has been performed during the first eighteen months of the project life cycle, with particular attention being paid to the hardware and software infrastructure, including services that are available and can be used in a favour of the project's purposes.

The document provides a brief outline of the MARVEL project in respect of resource demands and highlights various aspects of the available infrastructure. The first part describes the hardware infrastructure of the PSNC supercomputing centre. The second section presents the available software stack, including software packages, compilers, and libraries installed on the available clusters. It also describes how to install additional software packages in computing systems. In addition, services such as training and user support, which are already present in PSNC, are described in detail. This document also covers cloud computing systems to offer physical resource support for deploying MARVEL's services. Data Corpus presented in another section, is critical to make the data available and accessible to project pilots as well as to the scientific and industrial community. It has been developed based on an architecture that provides, on the one hand, high efficiency of data processing and serving, and, on the other hand, sufficient capacity for storing large volumes of data. Finally, the report describes the regulations for accessing the systems (HPC, cloud, VMs) in the supercomputer centre from a procedural and operational point of view. Following these rules is necessary to ensure security, high-quality services, and transparency in access to resources.

The infrastructure made available for design purposes was developed on the basis of a set of requirements submitted by the owners of pilot applications and is a reference starting point for further work, including adaptation resulting from the development of the MARVEL system in the further part of the project. The final status of the established infrastructure will be reported in deliverable 'D5.8 HPC infrastructure and resource management for audio-visual data analytics – final version' at M36.

Bibliography

- SLURM workload Manager*. (2022). Retrieved from <https://slurm.schedmd.com/>
- PSNC Eagle software*. (2022). Retrieved from <https://wiki.man.poznan.pl/hpc/index.php?title=Oprogramowanie>
- PSNC training courses*. (2022). Retrieved from <https://szkolenia.man.poznan.pl/?lang=en>
- PSNC HPC portal*. (2022). Retrieved from <https://hpc.man.poznan.pl/>
- Intel oneAPI DPC++/C++ Compiler*. (n.d.). Retrieved from <https://www.intel.com/content/www/us/en/developer/tools/oneapi/dpc-compiler.html#gs.3n9rbk>
- Kafka platform*. (n.d.). Retrieved from <https://kafka.apache.org/>
- Apache NiFi system*. (n.d.). Retrieved from <https://nifi.apache.org/>
- Elasticsearch*. (n.d.). Retrieved from <https://www.elastic.co/elastic-stack/>