



D4.3 Advanced NEMO platform & laboratory testing results. Final version

Document Identifica	tion		
Status	Final	Due Date	30/05/2025
Version	1.0	Submission Date	17/06/2025
Related WP	WP4	Document Reference	D4.3
Related	D1.1, D1.2, D1.3,	Dissemination Level (*)	PU
Deliverable(s)	D2.3, D3.3, D4.1, D4.2		
Lead Participant	DITRA	Lead Author	Dimitrios Skias
Contributors	SYN, INTRA,	Reviewers	Enric Pages-Montanera
	AEGIS, SPACE,		(ATOS)
	ATOS, MAG, ENG,		
	ESOFT, SU, COMS		
			Jonathan Klimt (RWTH)

Keywords:

Integration, validation, API, SDK, Lifecycle Management, Migration Controller, automation

Disclaimer for Deliverables with dissemination level PUBLIC

This document is issued within the frame and for the purpose of the NEMO project. This project has received funding from the European Union's Horizon Europe Framework Programme under Grant Agreement No. 101070118. The opinions expressed and arguments employed herein do not necessarily reflect the official views of the European Commission.

The dissemination of this document reflects only the author's view, and the European Commission is not responsible for any use that may be made of the information it contains. This deliverable is subject to final acceptance by the European Commission.

This document and its content are the property of the NEMO Consortium. The content of all or parts of this document can be used and distributed provided that the NEMO project and the document are properly referenced.

Each NEMO Partner may use this document in conformity with the NEMO Consortium Grant Agreement provisions.

(*) Dissemination level: (PU) Public, fully open, e.g., web (Deliverables flagged as public will be automatically published in CORDIS project's page). (SEN) Sensitive, limited under the conditions of the Grant Agreement. (Classified EU-R) EU RESTRICTED under the Commission Decision No2015/444. (Classified EU-C) EU CONFIDENTIAL under the Commission Decision No2015/444. (Classified EU-S) EU SECRET under the Commission Decision No2015/444.



Document Information

List of Contributors	
Name	Partner
Rubén Ramiro	ATOS
Enric Pages-Montanera	ATOS
Antonis Gonos	ESOFT
Dimitrios Skias	INTRA
Kostas Vrioni	MAG
Astik Samal	MAG
Nikos Drosos	SPACE
Emmanouil Bakiris	SPACE
Theodore Zahariadis	SYN
Ilias Seitanidis	SYN
Victor Gabillon	TSG
Spyros Vantolas	AEGIS
Hassane Rahich	SU Y
Matija Cankar	COMS

Document His	tory		
Version	Date	Change editors	Changes
0.1	19/02/2025	Dimitrios Skias (INTRA)	ToC 1 st version
0.2	20/04/2025	Dimitrios Skias (MTRA)	Final ToC
0.3	13/05/2025	INTRA, SYN, ATOS, UC3M, TID, TSG, AEGIS, SU, MAG, COMS	Initial contributions
0.4	20/05/2025	Dimitrios Skias (INTRA)	First round of contributions
0.5	23/5/2025	J X G, MAG, TID, SYN	Second round of contributions
0.6	30/5/2025	INTRA, ATOS, UC3M, SYN, COMS	Final round of Contributions
0.7	2/6/2025	RWTH, ATOS	Peer-Review ready version
0.8	6/6/2025	INTRA, MAG, TID, SYN, ATOS	Reviewers' comments consolidation
0.9	13/6/2025	INTRA, SYN	Document finalization
1.0	17/06/2025	ATOS	Quality Control & Submission

L

Quality Control		
Role	Who (Partner short name)	Approval Date
Deliverable leader	Dimitrios Skias (INTRA)	17/06/2025
Quality manager	R. Valle Soriano (ATOS)	17/06/2025
Project Coordinator	E. Pages (ATOS)	17/06/2025

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	2 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Table of Contents

Document Information	2
Table of Contents	
List of Tables	6
List of Figures	7
List of Acronyms	9
Executive Summary	
1 Introduction	
1.1 Purpose of the document	
1.2 Relation to other project work	
1.3 Relation between D4.3 and D4.2	
1.4 Structure of the document	
2 NEMO Integration, Validation & Verification approach and tools	
2.1 NEMO CI/CD Environment & Tools	
2.1.1 Open-Source repository	
2.1.2 NEMO Automated Deployment and Configuration	
2.2 3 rd Party documentation of NEMO meta-OS	17
2.3 Cloud/Edge/IoT Integration and Validation Infrastructure	
2.4 Integration & V&V Methodology & Plan	
2.4.1 NEMO V&V documentation	
3 NEMO Integrated Platform (Final Version)	
3.1 Meta-OS functionality in NEMO yer. 1.0	
3.1.1 NEMO Indirastructure Management	
3.2 NEMO Open Call integration activities	
3.2.1 ARGO	
3.2.2 CorNOS	
323 Eros4NRG	
Genesys	
3.2. MARINEMO	
3.2.6 MetaFOX	
4 NEMO Service Management Layer updates	
4.1 Intent-based Migration Controller	
4.1.1 Overview	
4.1.2 Architecture	
4.1.3 Interaction with other NEMO components	
D4.3 Advanced NEMO platform & laboratory testing results	

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	3 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



4.1.4	Conclusion	39
4.2 Plugin	& Applications Lifecycle Manager	39
4.2.1	Architecture	40
4.2.2	Conclusion	41
4.3 Monet	zation and Consensus-based Accountability	41
4.3.1	Overview	41
4.3.2	Architecture	
4.3.3	Interaction with other NEMO components	
4.3.4	Results	
4.3.5	MOCA Business Models	47
4.4 Intent-	based SDK/API	
4.4.1	Overview	49
4.4.2	Architecture	49
4.4.3	Interaction with other NEMO components	49
4.4.4	Conclusion	
5 NEMO sce	nario-driven verification & results	52
5.1 NEMC	Cluster registration – resource provisioning	52
5.1.1	Verification scenario	53
5.1.2	Results	53
5.2 NEMC) workload registration & provisioning	58
5.2.1	Verification scenario	
5.2.2	Results	64
5.3 NEMC) workload scheduling & orchestration	68
5.3.1	Verification scenario	
5.3.2	Results	71
5.4 NEMC) workload lifecycle management	81
5.4.1	Ventication scenario	
5.4.2	Results	
5.5 NEMC	Secure Execution Environment	92
5.5.1	Verification scenario	
5.5.2	Results	
6 Conclusior	15	97
7 References		98
Annex 1-Net	work Intent	101
Annex 2-LCN	A subcomponent	103
Annex 3 – Se	rviceProviderModel	105

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	4 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Annex 4-Intent for network connectivity request

Pondino

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				esults.	Page:	5 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



List of Tables

Table 1. Differences between D4.3 and D4.2	13
Table 2. Integration testing - scenario template	27
Table 3. NEMO integration testing - Checkpoints template	27
Table 4. NEMO dev cluster (K8S)	28
Table 5. Staging 1 cluster (K8S)	29
Table 6. Staging 2 Cluster (K3S)	30
Table 7. Production Cluster (k8S)	31
Table 8. Cluster Metrics	57
Table 9. Intent for network connectivity request. L2S-M network request attributes.	61
dille	
2ett	
▼	

Document name:	D4.3 A Final ve	4.3 Advanced NEMO platform & laboratory testing results.					6 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



List of Figures

	17
Figure 1. NEMO code repository in Eclipse Research labs	-10
Figure 2. NEMO CI/CD organization	$-\frac{1}{10}$
Figure 3. NEMO Login page	$-\frac{18}{10}$
Figure 4. NEMO main management page	_ 18
Figure 5. Registered Workloads View	_ 19
Figure 6. Create Workload View	_ 20
Figure 7. Create Workload Instance View	_ 20
Figure 8. Workload Instances main View	21
Figure 9. Intent Management View	22
Figure 10. Create Intent View	22
Figure 11. Resource Provisioning View	23
Figure 12. Cluster Registration Form	23
Figure 13. Intent-Based API endpoints	24
Figure 14. Sample of the expected parameters' body of the workload POST endpoin	25
Figure 15: NEMO project phases and main meta-OS version releases	26
Figure 16. IBMC Architecture	37
Figure 17. Migration Sequence Diagram	38
Figure 18. LCM high-level architecture	<i>40</i>
Figure 19. MOCA architecture	42
Figure 20. MOCA integration diagram	43
Figure 21. Example of deploying dataset	- 46
Figure 22. Example of contract deployment	47
Figure 23. Final Architecture diagram of Intent-Based API	49
Figure 24 Cluster registration sequence diagram	- <u>5</u> 2
Figure 25 LCM home page	- <u>53</u>
Figure 26. Elemente page	$-\frac{55}{54}$
Figure 27 Cluster registration form	57 54
Figure 28 Cluster nending status	- 57 55
Figure 20. MOCA cluster registration message through RabbitMO	- 55 55
Figure 30 Meta-Orchestrator response	- 55 55
Figure 31 Pagistar cluster to blockchain	56
Figure 32 Undated cluster statut	- 50 56
Figure 32. Workload provisioning workflow in NEMO	- 50 58
Figure 34. NEMO workland conformant workflow	$-\frac{50}{50}$
Figure 35. Network workflowd deployment workflow	
Figure 35. Network workload deployment workflow	_ 65
Figure 50. Arrival of intent and first filter.	$-\frac{05}{65}$
Figure 57. Grpc connector execution	$-\frac{05}{65}$
Figure 56. L25-W unnotations to be used by the worktoad.	_ 05
Figure 59. Network resource created in S1 cluster	_ 00
Figure 40. Network resource created in S2 cluster	$-\frac{00}{67}$
Figure 41. Poa deployed in ST with L2S-M annotations	$-\frac{0}{7}$
Figure 42 Poa deployed in 52 with L25-M annotations	$-\frac{0}{7}$
Figure 48. Ping between pods in S1 and S2 clusters	$-\frac{0}{0}$
Figure 44. Fing between poas in S1 and S2 clusters	_ 08
Figure 43. Intent-Based Migration Sequence Diagram	_ 68
Figure 40. CFDRL Migration Sequence Diagram	_ 69
Figure 4/. Horizontal Scaling Sequence Diagram	_ 69
Figure 48. LCM Workload Initial Visualization	72
Figure 49. Availability Intent Creation	72
Figure 50. Ibmc-controller processing the intent	_ 73
Figure 51. Dev Cluster Ibmc-agent logs	73

Document name:	D4.3 A Final ve	14.3 Advanced NEMO platform & laboratory testing results. inal version					7 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Figure 52. Staging Cluster Ibmc-agent logs	73
Figure 53. LCM workload final visualization	74
Figure 54. LCM complete workflow view	74
Figure 55. Inputs/outputs from CFDRL point of view	78
Figure 56. JSON Schema for MO-CFDRL RabbitMQ communication	78
Figure 57. Messages used for Scaling	79
Figure 58. Log of the action taken by CFDRL and sent to the Cluster Ibmc	80
Figure 59. Message published by ibmc	80
Figure 60. Log of the Migration action being implemented by the cluster IBMC	80
Figure 61. CFDRL decides a scaling action and sends it through Rabbit MQ to MO	81
Figure 62. Meta orchestrator receives the scaling action and triggers the scaling	81
Figure 63. The number of replicas is checked and it is set to 3 as requested	81
Figure 64. Workload Lifecycle Sequence Diagram	81
Figure 65. LCM User access	83
Figure 66. Example of intent conditions	83
Figure 67. Intent Evaluator logs	84
Figure 68. Intent API receives the Intent Evaluation report	84
Figure 69. Cluster metrics from PPEF	85
Figure 70. Cluster metrics RabbitMQ payload	85
Figure 71. The CMDT component includes a SwaggerUI endpoint with up-to-date documentation and example	es.
	85
Figure 72. Detailed information about the NEMO workload pod showing its location, status, traffic, and	
response times.	86
Figure 73. Detailed information about the pod's traffic. Additionally at displays its internal traffic, as well as	
inbound and outbound traffic with foreign pods.	87
Figure 74. Tree-like representation of all deployments, excluding non-NEMO workload	
components	87
Figure 75. Manage workloads	88
Figure 76. Manage workload instances	88
Figure 77. Workload Instance lifecycle timetine	89
Figure 78. LCM workload deployment tree	89
Figure 79. Workload intents management	90
Figure 80. Resource provisioning	90
Figure 81. Cluster performance	91
Figure 82. Intent Fulfilled	91
Figure 83. Intent Violation	92
Figure 84. Unikernel creation sequence diagram	92
Figure 85. publishToSee endpoint	95
Figure 86. SEE success message	96
Figure 87. Schematic of the communication channels used by the MO API and the SEE	96

Rett

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	8 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



List of Acronyms

Abbreviation / acronym	Description
AAA	Authentication, Authorization, and Accounting
AI	Artificial Intelligence
API	Application Programming Interface
CD	Continuous Delivery
CFDRL	Cybersecure Federated Deep Reinforcement Learning
CIDR	Classless Inter-Domain Routing
CLI	Command Line Interface
CMDT	Cybersecure Microservices' Digital Twin
CI	Continuous Integration
CLI	Command-line Interface
CNCF	Cloud Native Computing Foundation
CPU	Central Processing Unit
CRD	Custom Resource Definition
DApps	Distributed Applications
DLT	Distributed Ledger Technology
DNS	Domain Name System
Dx.y	Deliverable number y belonging to WP x
E2E	End-to-End
EC	European Commission
EV	Electric Vehicles
FL	Federated Learning
GDPR	General Data Protection Regulation
GPU	Graphics Processing Onit
gRPC	Generic Remote Procedure Call
GUI	Graphical User Interface
IBA	Intent-Based API
IBMC	Intent-based Migration Controller
IdM	Mentity Management
IDS	Intrusion Detection System
IPFS	Interplanetary File System
IoT	Internet-of-Things
IT	Information Technology
K8s	Kubernetes
LAN	Local Area Network
LCM	Life-Cycle Manager
LL	Living Lab
meta-OS	Meta-Operating System
ML	Machine Learning
IDS IPES IoT IT K8s LAN LCM LL meta-OS ML	Interplanetary File System Interplanetary File System Internet-of-Things Information Technology Kubernetes Local Area Network Life-Cycle Manager Living Lab Meta-Operating System Machine Learning

Document name:	D4.3 Ao Final ve)4.3 Advanced NEMO platform & laboratory testing results. Final version					9 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Abbreviation /	Description
acronym	
MLaaS	ML-as-a-Service
mNCC	Meta Network Cluster Controller
МО	Meta-Orchestrator
MOCA	Monetization and Consensus-based Accountability
MQTT	Message Queuing Telemetry Transport
NAC	NEMO Access Control
OS	Operating System
OSS	One-stop-shop
PAYG	Pay-As-You-Go
PV	Photovoltaic
PPEF	PRESS & Policy Enforcement Framework
RaaS	Resource-as-a-Service
REST	Representational State Transfer
SDN	Software Defined Network
SEE	Secure Execution Environment
SME	Small-Medium Enterprise
SMP	Slice Manager Plugin
SOM	System-on-Module
UI	User Interface
UPS	Uninterruptible Power Supply
UX	User Experience
YAML	yet another markup language



Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	10 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Executive Summary

The present deliverable, D4.3, entitled "Advanced NEMO platform & laboratory testing results-Final version" aims to provide a complete and detailed overview of the integration process carried out within the framework of the Next Generation Meta Operating System project towards its realization. Moreover, this document sheds light on the final technical specifications of the components that were developed in the context of the Work Package 4, namely the Intent-based API, MOCA, LCM and the IBMC.

Within the scope of this document several other aspects that enhance the understanding of the integration and validation process followed in NEMO are discussed. To this end, the infrastructure related aspects, such as the hardware specifications of OneLab facilities that host the NEMO meta-OS framework are presented. The ZeroOps configuration and deployment procedure followed, as described, can be utilized as a paradigm for future deployments of the NEMO meta-OS in 3rd party infrastructure. In addition, comprehensive documentation tailored for 3rd parties that aim to utilize NEMO meta-OS is included, presenting the steps that are necessary to access, register and deploy workloads in the NEMO meta-OS.

The present document intents to provide insights on the final integrated version on the NEMO meta-OS which is the result of a Validation and Verification methodology that is established and adopted. The associated work concerns also the scenario-driven integration tests that were conducted in laboratory setting including the supporting CI/CD environment and tools.

Finally, the document is foreseen to be used as the best practices handbook for other projects in terms of integration and validation procedures. NEMO since its beginning provided state-of-the-art architectural patterns that made the process of integration and validation easier, in terms of interconnecting components and debugging issues in the End-to-End system workflows.

Rengin

Document name:	D4.3 A Final ve	D4.3 Advanced NEMO platform & laboratory testing results. Final version					11 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



1 Introduction

1.1 Purpose of the document

Deliverable D4.3 "Advanced NEMO platform & laboratory testing results-Final version" is the second and final version of the NEMO platform-laboratory results reporting. This document provides the updates and optimizations made since the delivery of deliverable D4.2 "Advanced NEMO platform & laboratory testing results-Initial version" [1] that aimed to provide the main guidelines for the testing and integrating steps to be followed towards NEMO's validation. Moreover, the updated workflows used for these components' interactions are reported in the following sections. As D4.3 is considered the final version of the integration deliverables of NEMO, an overview of the workflows used by the components of other WPs are presented towards an End-to-End workflow presentation.

1.2 Relation to other project work

Deliverable D4.3 is tightly connected to D4.2, as it is the second version of frond reports the updates of the work initially described in D4.2. In addition, D4.3 reports the issues that the component owners met during the finalization of their components. D4.3 is also strongly coupled with D4.1 [2] on which the integration methodology to be followed by the project components was defined as well as some preliminary integration activities and testing results that took place among the more mature NEMO components. In D4.3, the final version of the NEMO meta-OS was validated over as well as the updated integration plan. At Work Package (WP) level, D4.3 is firmly attached with the WP4 tasks. Tasks T4.1-T4.3 contribute to this deliverable by providing the updates carried out within the individual tools implemented by each task. Furthermore, task T4.4, responsible for the integration of the NEMO components vastly committed to the design of the individual integration pipelines as well as monitoring the procedure of integrating submodules towards the final integration.

Moreover, this deliverable has strong relation with other WPs and the technical work conducted. The components' specifications and meta-architecture as well as the functional and non-functional requirements defined within WP1 [3] and reported on the related deliverables are validated in the current deliverable and any change of them is reported. The benchmarking definition provided by task T1.3 is demonstrated through the integration procedures reported in the following sections. While WP4 contributed to WP2 and WP3 in terms of providing the best practices related to the integration procedure, the integration steps and semantics of these two WPs are being reported in this deliverable.

1.3 Relation between D4.3 and D4.2

In D4.2, the initial outcomes as well as the guidelines to be followed in the WP4 components integration were presented, in D4.3 the final updates and optimizations carried out by the components' owners are described. In addition, the final workflows of the WP4 and other WPs are presented to provide a complete picture of the End-to-End workflow of NEMO. While the two deliverables share some common chapters, they have several differences, for the ease of the reader these changes are described in Table 1. The structure of D4.3 has been changed compared to D4.2 to provide a better understanding of the NEMO integration process and the updates of the WP4 tools to the reader. New subsections were added to incorporate the integration activities carried out by the Open Call projects as well as some of the subsections of D4.2 were converted to standalone sections in D4.3 to provide more details about the progress of the Project and the integration pipelines.

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	12 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Section in D4.3	Section in D4.2	Differences
1. Introduction	1. Introduction	Several updates in this section
1.1 Purpose of the document	1.1 Purpose of the document	Minor update
1.2 Relation to other project work	1.2 Relation to other project work	The section has been updated
1.3 Relation between D4.3 and D4.2	-	New section in D4.3
1.4 Structure of the document	1.4 Structure of the document	The section has been updated
2. NEMO Integration, Validation & Verification approach and tools	2. Methodology	Several updates in this section
2.1 NEMO CI/CD Environment & Tools	2.1 NEMO CI/CD Environment & Tools	Several updates in this Section.
2.2 3rd Party documentation of NEMO meta- OS	-	New Section in D4.3
2.3 Cloud/Edge/IoT Integration and Validation Infrastructure	2.2 Cloud/Edge/IoT Integration and Validation Infrastructure	The section has been updated
2.4 Integration & V&V Methodology & Plan	2.3 Integration & V&V Methodology & Plan	Minor Updates
3 NEMO Integrated Platform (Final Version)	2.5 NEMO Integrated Platform	New Section in D4.3, Updated
3.1 Meta-OS functionality in NEMO v2	2.5.1 Meta-OS functionality in NEMO v1	The section has been updated
3.2 NEMO Open Call integration activities	-	New Section in D4.3
4. NEMO Service Management Layer	3. NEMO Service Management Layer	Major Updates
4.1 Intent-based Migration Controller	3.1 Intent-based Migration Controller	Updated
4.2 Plugin & Applications Lifecycle Manager	3.2 Plugin & Applications Lifecycle Manager	Updated
4.3 Monetization and Consensus-based Accountability	3.3 Monetization and Consensus-based Accountability	Updated
4.4 Intent-based SDK/API	3.4 Intent-based SDK/API	Updated
5. NEMO scenario-driven verification & results	4. NEMO scenario-driven verification & results	Major Updates
5.1 NEMO Cluster registration – resource provisioning	4.1 NEMO Cluster registration	Major Updates
5.2 NEMO workload registration & provisioning	4.2 NEMO workload registration, deployment & provisioning	Major Updates
5.3 NEMO workload scheduling & orchestration	4.3 NEMO workload migration	Major Updates
5.4 NEMO workload lifecycle management	4.4 NEMO workload lifecycle management	Major Updates
5.5 NEMO WP3 Integration Activities	-	New section
6. Conclusions	5. Conclusions	The section has been updated

Table 1. Differences between D4.3 and D4.2

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	13 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



1.4 Structure of the document

Renge

The remainder of this report is organized as follows:

- Section 2 contains information on the CI/CD environment of the NEMO meta-OS and on the OneLab facilities that provide for the integration activities of the first integrated version of the NEMO meta-OS. In addition, it presents the high-level architecture view of the first integrated version of the NEMO meta-OS highlighting the key integration activities for each functional layer and describes the components that are fully or partially integrated.
- Section 3 presents the high-level NEMO Integrated Platform along with the integration activities of the Open Call Projects with NEMO.
- Section 4 describes the overview, the architecture, the initial results and the interactions with other components for the modules that are comprising the Service Management Laver of the NEMO meta- OS platform, namely the intent-based Migration Controller (TBMO), the Plugin & Application Lifecycle Manager (LCM), the Monetization and Consensus based Accountability (MOCA) and the Intent-based SDK/API.
- Section 5 sheds light into the integration activities that are conducted and materialized the first integrated version of the NEMO meta-OS, following the scenario-driver V&V methodology.
- Section 6 provides conclusions and insights of the final version of the NEMO meta-OS.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	dvanced NEMO platform & laboratory testing results. Page: 14 of 111				
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



2 NEMO Integration, Validation & Verification approach and tools

2.1 NEMO CI/CD Environment & Tools

2.1.1 Open-Source repository

NEMO is making use of a GitLab CI/CD framework [4] deployed at the Eclipse Research Labs to house and keep track of the code stack developed within the project. Moreover, within NEMO a docker registry repository is being used for storing and deploying the generated containers. The pipeline adopted by NEMO is integrated with the GitLab repository and utilizes an automated build functionality that generates the docker container images [5] based on the architectural tool chains of the project. The final step of the pipeline involves the automated deployment of the latest image generated to the appropriate server. The use of a centralized versioning system such as GitLab provides several benefits:

- Automates repetitive tasks like running tests and building code.
- Integrates code frequently, allowing teams to catch issues sooner.
- Run automated tests and static analysis tools with every commit
- Automatically deploy to staging or production environments.
- Track the state of builds, tests, and deployments in real time.
- Keep a record of who deployed what, when, and why
- Link code changes directly to issues or feature requests.
- Allow for code review and discussion before merging.
- Detect vulnerabilities in dependencies or container images.

On the other hand, apart from the technical benefits of using a version control repository, there are other benefits as well. NEMO since its conception aimed to provide an opensource meta-OS framework that will enable local organizations to deploy and use a load management system in the IoT-Edge-Cloud continuum. For this reason, the NEMO's repository¹ is public and accessible along with the Open Call projects^{2,3} that aim at validating and enhancing NEMO. The repository follows a sub-project organization, Figure 1, where the components are grouped based on the implemented outcomes of the relevant tasks, while two major subprojects, Open Call 1 and Open Call 2, house the services enhancing and validating NEMO.



³ https://gitlab.eclipse.org/eclipse-research-labs/nemo-project/opencall-2

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	15 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final

¹ <u>https://gitlab.eclipse.org/eclipse-research-labs/nemo-project</u>

² <u>https://gitlab.eclipse.org/eclipse-research-labs/nemo-project/opencall-1</u>



Q Search or go	o to	📚 Eclipse Research Labs > 🍓 NEMO Project			
Proup		NEMO Project⊕			
Manage		Group ID: 3920 🖞			
코 Plan	>	NEMO aims to establish itself as the game changer of AloT-Edge-Cloud Continuum by introducing an open source, flexible, adaptable, c meta-Operating System.	ybersecure and multi-	technol	logy
I> Code ₽ Secure	>	Subgroups and projects Shared projects Archived projects Q. Search		Name	5 V
🗄 Analyze	>	> 💱 N NEMO Cybersecurity and Unified_Federated Access Control 🌐	Se :	00	සි
		> S• N NEMO Federated MLOps 🕀	8* (00	රිර
		> 3* N NEMO Infrastructure Management	00	: O 0	ධිපි
		> 💲 N NEMO Kernel 🌐	0e /	100	â
		Se N NEMO PRESS and Policy Enforcement	8° (00	86
		> 💲 N NEMO Service Management 🕀	0.0	00	සිස
		() N Nemo Docs () *)	1 ye	er a
		1 N Nemo HowTo (#)	2	2 wee	łks a



2.1.2 NEMO Automated Deployment and Configuration

Rett

In Software development one of the fundamental steps in the lifecycle of a system is the CI/CD development methodology. In this context, within NEMO a CI/CD pipeline was set up to automate the deployment and configuration of the NEMO components in various environments. In D4.2 the basic functionality of the NEMO CI/CD was described along with a short guide⁴ for the developers.

The NEMO components are organized in a cluster-oriented hierarchy within the Flux CD [6] [7] continuous delivery solution as depicted in Figure 2. Each directory contains the configuration files of the components deployed at that particular cluster. In the final version of the NEMO meta-OS, apart from the component integration, the integration of 3rd party infrastructures took place. Figure 2 illustrates both the OneLab clusters and the pilot environments. The NEMO CI/CD [6] played a pivotal role as it was possible to deploy all the NEMO core components in the Pilots' clusters without the need for extensive configuration and involving the components' owners achieving a high impact with low effort curve.



Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	16 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



🚵 NEMO / FLUX CD / Repository

% main ∽ flux-cd / clusters / Find file Edit ~ Code ~ + ~ Name Last commit Last update Donelab-dev-new nemometaos/mo-api:v2.0.10 15 hours ago Edit ingress.yaml 🗅 onelab-dev 1 month ago Conelab-k3s Merge branch 'main' into feature/deployment-controller-ibmc 2 weeks ago Conclab-production Remove MO API config from staging and producction clusters and... 1 day ago D onelab-staging Merge branch 'feature/deployment-controller-ibmc' into 'main 1 day ago smart-farming-pilot Velero chart added to staging cluster 1 day ago Smart-media-pilot Merge branch 'service/smart-media' into 'main' 16 hours ago

Figure 2. NEMO CI/CD organization

2.2 3rd Party documentation of NEMO meta-OS

Within the NEMO Meta-OS project, a framework that will enable the effortless deployment of services through Zero DevOps Tools was designed and implemented. The derived framework based on the best practices of microservices [8] architecture deployment in a Kubernetes [9] [10] [11] cluster environment provides to the 3rd party user two interfaces to interact with. A user-friendly graphical user interface based on the best practices of UI/UX and a REST API for service-to-service interaction with NEMO, allowing the 3rd party user to develop railered interactions with NEMO.

The User is initially landed on the Login page of NEMO, Figure 3, to authenticate and be categorized based on the user categories described in deliverable D1.3 [12]. The authentication mechanism used is based on the Keycloak [13] framework which is a well-known and established identity management system widely used in modern critical systems. After the user is successfully authenticated, he is redirected to the main NEMO management page, Figure 4. In this view the main action categories for using NEMO are displayed:

Workload Monitoring
 Workload Instances
 Intent Management^{5,6}
 Besource Provisioning

More details about the functionalities of each category will be provided later in this section. Finally, the option log out on the top right corner of the view concludes the list of actions for this view.

⁶<u>https://gitlab.eclipse.org/eclipse-research-labs/nemo-project/nemo-service-management/intent-based-sdk_api/intent-api/-/blob/main/network-intents.yaml?ref_type=heads</u>

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	17 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final

⁵https://gitlab.eclipse.org/eclipse-research-labs/nemo-project/nemo-service-management/intent-based-sdk_api/intent-api/-/blob/main/intent-examples.yml?ref_type=heads



	Sign in to your Username or email	LOAK account	
Nemo LCM I Services Overview	Sign In Figure 3. NEM	O Logia page	E+ Logout
III » Hosts] Workload Monitoring Workload Monitoring Manage Workloads	Workload Instances	Intent Management	Resource Provisioning
	Eigung 4 NIEMO mai		

Figure 4. NEMO main management page

The Registered Workloads View, Figure 5, provides a complete monitoring interface for the already submitted workloads. Each workload record has its unique identification number sorted in a descending order. The Workload Name and Version attributes are the unique identifiers upon the creation of each Workload record, meaning that for a given workload name there cannot be two records with the same version. In the example displayed in Figure 5, the ml-retrain workload task is displayed multiple times, however the version differs each time. Then, there is the Status attribute, the values of this attribute are Pending and Accepted. When a workload is created it gets in the pending status until it is evaluated and validated by the NEMO framework. After this procedure is finished the workload gets in the accepted status. Next there is the Ingress Support field that denotes if the workload supports or not network communication with services outside of the NEMO environment. The User field displays the user that created, and it is used for trackability of user actions within the NEMO framework. Finally, the Actions column houses the functionalities related to the management of each workload. The first item is a file

Document name:	D4.3 A Final ve	dvanced NEMO platform & laboratory testing results. Page: 18 of 111					
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



uploader, and it is used to upload the HELM [14] file of the workload. The file can be uploaded after the creation of the workload and the workload cannot transit to the accepted status before this action is performed. The second item is the instance creation button, it is used to create a workload instance based on the workload's configuration, both from the workload creation page and the uploaded HELM file. The third item in the row is the deletion action in case the user wants to delete the workload. Finally, the associated instances button is a shortcut for the Workload instances view for this specific workload. In Figure 5 within the Actions column the HELM file uploader and the deletion buttons are faded, until a workload gets to the accepted status, the user is able to upload a new HELM file and/or delete the workload. After the workload is marked as accepted, the user cannot modify the existing workload, any change required on the existing workload requires the creation of a new workload.

_											
	Services (Dverview									
	Sea	arch									
	Create	Workload									
	ID ↓	Workload Name	Version	Status	Ingress Support	ser	Balance	Actions			
	33	ml-retrain	0.1.6	ACCEPTED	•				⊘	Î	
	32	ml-retrain	0.1.5	ACCEPTED	•			*	۲	Î	
	31	ml-retrain	0.1.4	ACCEPTED	•			^	€	Î	
	30	ml-retrain	0.1.3	ACCEPTED	•			Ť	€	Î	
	29	ml-retrain	0.1.2	ACCEPTED	•				⊘	Î	
	28	ml-retrain	0.1.1	ACCEPTED	•				€	Î	
	27	ml-retrain	0.1.0	ACCEPTED	•				⊘	Î	
	26	echo-server	0.5.16	ACCEPTED	•			T	€	Î	
	25	echo-server	0.5.15	ACCEPTED	•			+	۲	Î	
	24	ml-pvc	0.1.6	ACCEPTED	•			^	⊛	Î	
						ltems per pa	ge: 10 •	1 -	10 of 28		$\langle \rangle$



In Figure 6 the workload creation form is displayed. In this view the user enters the basic information of the workload to be used as we explained earlier the combination of the name and version attributes need to be unique and these fields are mandatory. The field Supported Intents contains a dropdown list of the available intent categories that the user can select for deploying the workload in a cluster based on the needs of the service, e.g., availability, existence of dedicated GPU, energy, etc. The user can modify the intents to be used on the cluster selection by the Meta-Orchestrator also after the workload's creation through the Intent Management view. Finally, there is the Ingress Support attribute discussed before. The values assigned are True or False, with False being the default value for security purposes.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	results.	Page:	19 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



100	🖵 Reaist	tered Workloads						
Search	🖵 Cre	ate Workload						
	Name *	Version *						
Create Workload	Type Workload Name	Type Workload \						
ID 🕡 Workload Na	Workload Name is required Supported Intents			ons				
33 ml-retrain	Supported Intents*		*	9	⊙	Î	С	
32 ml-retrain				9	€	Î	С	
31 ml-retrain	Ingress Support			9	⊘	Î	С	
30 ml-retrain	False -			Ð	⊘	Î	С	
29 ml-retrain		Create	Reset	9	⊘	Î	С	
28 ml-retrain	0.1.1 AUGEPTED	аккашоунени	us.com -	œ	⊛	Î	С	
07 ml rotroin		lakka Qaunalia		-	0	-		

Figure 6. Create Workload Vie

The last step in the pipeline defined for executing a service in NEMO is the workload instance creation, Figure 7. After a workload is created and gets on the accepted status, an instance of it can be created. The user has to fill the name of the instance (Release Name) and if applicable a proposed cluster to be executed. The attributes: Values Override, Include CRDs, Is Upgrade and NoHooks are out of scope of NEMO as these are Kubernetes oriented parameters. The workload instance can be created without any further action, however, by clicking the Create Intent Template button a submenu opens. In this submenu the user is able to select one or more requirements that will enhance the execution of his service. In the sample displayed in Figure 7 shows that the desired execution environment for the service is a cluster able to provide an availability higher than 90%. Moreover, the user is able to set a specific time range that wants this migration to take place, e.g., high intense processing periods, etc.

	🖵 Create Wo	rkload Instace	
Release Name *			
Type Release Name			
Release Name Is required Cluster Name			
nemo-smart-farming	÷		
Values Override		Include C RD 3	
0	11.	False	~
Is Upgrade		No Hooks	
False	÷	False	*
Intent Type *			
Select Intent Type	÷	Remove Intent	
Start Date-lime		End Date-time	
Choose start datetime		Choose end datetime	-
Target Name	Target Condition	Target Value	
availability	IS_GREATER_THAN +	90	
Add Target			
		Create In	tent Template Render Re

Figure 7. Create Workload Instance View

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	20 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Another high-level category in the NEMO platform is Workload Instance View, Figure 8. In this view the user can see the status of the workload instances previously deployed and manage them. The Release name refers to the name provided in the previous form while the Workload ID refers to the parent workload created at the first step. The status denotes whether the instance is Rendered, Deployed or deleted. The Cluster field shows the cluster that the current instance is deployed based on the Meta-Orchestrator's decision-making algorithm and the requirements (Intents) that the user set up. The instance ID is the unique identifier of each instance. The combination of Release name and Instance ID is unique identifier of each deployed workload instance across all clusters connected to NEMO. Moreover, a set of useful interactions in terms of status monitoring and management are provided in the action column. First, a historic representation of the various statuses of the workload instance e.g., created, deployed, suspended, deleted along with the timestamp of each status update. Then there is the manifest viewer, which can be used for examining the configuration parameters in case an error occurs and finally the deletion button can be used to permanently terminate the service instance.

Nemo LCM	III Se	ervice	s Overview						
					🖵 Workloa	d Instances			
		Se	arch						
		ID ,	Release Name	Workload ID	Status	Cluster	Instance ID	Actions	
		89	smart-xr-retrain-2	33	DELETED	dev-onelab	218b064	~	
		88	smart-xr-retrain-1	33	DEPLOYED	dev-onelab	4b0d188	~	
		87	smart-xr-retrain	33	DEPLOYED	dev-onelab	7c3bf46	~	Ĩ
		86	ml-retrain4	29	DELETED	dev-onelab	661ba90	~	Ĩ
		85	ml-retrain3	28	DELETED	dev-onelab	9cc3430	~	Ĩ
		84	ml-retrain2	27	DELETED	dev-onelab	54e69b4	~	Ĩ
		83	ml-retrain1	28	DELETED	dev-onelab	2aaaff1	~	
			•	Digura 8	Workload In	nstances mai	in View		

NEMO was designed to nightly workloads based on specific requirements (Intents), for this reason the migrations take place based on tailored and sticked requests. NEMO provides to the user the possibility to modify on the fly these requirements for any given running workload instance. In this context, the Intents management View was created to house this functionality, Figure 9. As with the previously described tiews) the status field denotes the current condition of a specific Intent type on a specific workload instance. The status can be *Fulfilled* or *not fulfilled*, *compliant*, *degraded* or *fulfilmentfailed*. In the Intent Type column, the specific requirement (Intent) is displayed while in the next column the specific Workload instance is given. It is worth noting that one workload instance can have multiple intent types, even of different types. In the actions column there is the Intent Report option which provides helpful insights for understanding why an Intent failed by displaying detailed information of the error occurred. In addition, the user can modify the existing Intent by altering the desired target value. Finally, a simple control panel offers to the user the possibility to start and suspend the execution of an Intent with the related buttons as well as to delete it.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	21 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



LCM	III Services Overv	iew				
			Intents Management	ent		
	Search.					
	Create Inter	t				
	ID ↓	Status	Intent Type	Workload Instance ID	Actions	
	132	NOT_FULFILLED COMPLIANT	cloud_continuum	spain-n	X → = =	
	131	NOT_FULFILLED COMPLIANT	cloud_continuum	spain-n	X > = =	
	130	NOT_FULFILLED COMPLIANT	cloud_continuum	spain-n) X > = =	
	129	NOT_FULFILLED COMPLIANT	cloud_continuum	spain-n) 📜 🔸 🔳 🖷	
	128	NOT_FULFILLED DEGRADED	DeliverComputingWorkload	f0407bb	X > = =	
	127	NOT_FULFILLED FULFILMENTFAILED	MachineLearning	f0407bb) X > = =	
	126	NOT_FULFILLED COMPLIANT	cloud_continuum	spain-n) 🗶 🔸 🔳 🖷	
	125	NOT_FULFILLED COMPLIANT	cloud_continuum	spain-n	X > = =	
	124	NOT_FULFILLED COMPLIANT	cloud_continuum	spain-n) (X) (X) (X) (X) (X) (X) (X) (X) (X) (X	
	123	NOT_FULFILLED COMPLIANT	cloud_continuum	spain-n	X > = =	
				items per pa	ge: 10 👻 1 - 10 of 132 <)	, I

Figure 9. Intent Management View

From the Intents Management main View, the user is able to see the existing Intents but also to create new Intents, in Figure 10 the form of adding new Intents to an existing workload is displayed. The advantage of this view is that the user can add intents after the creation of a workload instance as well as for advanced users to create an intent using YAML files. The use of YAML files is used within NEMO for the connection and orchestration of complex components such as the mNCC for network creation for pods across the clusters.

		Intents	Management			
Search		🖵 Crea	te Intents			
Create Intent	Use yaml file					
ID ↓ Sta	Workload Instance ID *		Start Date-time			
132	Type Workload Instance ID		Choose start datetime	Ē		1
131						
130	Intent Type *		End Date-time			Ē
129	Select Intent Type	•	Choose end datetime			Î
128	Target Name	Target Condition	Target Value			Î
127	Type Target Name	Select C 🔻	Type Target Value			T
126						
125	Add Field					Î
124				Create	Reset	ī
123		ciouo_continuum	shau.u		<u>, , , , , , , , , , , , , , , , , , , </u>	

Figure 10. Create Intent View

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	22 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



The last menu provided by the NEMO Management User Interface is the Resource Provisioning View, Figure 11. In this View the user can see the available clusters connected to the NEMO platform along with some useful information such as their processing capabilities and real-time performance metrics.

Search						
reate Cluster						
ID ↓	Name	Endpoint	CPUs	Storage	Memory	Actions
933189e	dev-onelab	https:// api.main.nemo.onelab.eu:644	68	1620	172	50 B
8d63cc0	test-cluster	https:// api.main.nemo.onelab.eu:644	10	300	120	:
67c21fc	pro-onelab	https:// api.prod.nemo.onelab.eu:644	32	1350	62	:: i
60a2ee7	k3s-onelab	https:// api.s2.nemo.onelab.eu:6443	16	1080	32	
179592b	staging-onelab	https:// api.staging1.nemo.onelab.eu	32	1080	62	51 T
					Items per page: 10	▼ 1 - 5 of 5 <

Clusters Table

Figure 11. Resource Provisioning Xiew

Moreover, through this menu option, the users providing their own cluster infrastructure can register them through this view, Figure 12. The infrastructure providers need to fill the information related to the processing, environmental and financial aspects of their machines. The cluster registration process is divided into two steps, a) provide the basic information used by the NEMO components for the decision-making process of pods' deployment and b) the underlying inter-cluster connectivity. The registration form presented is responsible for tacking the first part of the registration, while the second part of the registration is achieved by aploading the kubeConfig file of the Kubernetes cluster to be registered. The Graphical User Interface previously described is one of the ways to interact with the NEMO platform. The GUI is indented to be used from non-technical personnel that need to effortlessly deploy a workload.

			🖵 Regi	ister Cluster		
	Search Create Cluster	Cluster Config File Upload Config File	e			
	ID ↓	Cluster Name		Managed API		tions
	933189e	Type Cluster Name		Type Cluster Name	e	
	8d63cc0	CPUs #	Memory (MB)	Storage (GB)	V-RAM (MB)	
\mathbf{A})	67c21fc	Type # of CPUs	Type Memory Size	Type Storage Size	Type V-RAM Size	s =
	60a2ee7	Availability %	Green Energy %		Cost	s .
	179592b	Select Availability	- Select Green	Energy % 👻	Select Cost	s •
		CPU Base Rate (in ms)		Memory Base Rate (in M	IBs)	- 5 of 5 < >
		Type CPU Base Rate		Type Memory Bas	e Rate	
					Submit	

Figure 12. Cluster Registration Form

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	23 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Another way of interacting with NEMO is through a REST API [15], Figure 13. Intent-Based API endpoints. Through this solution a power user can make a set of API calls in order to create tailored pipelines in terms of deploying, configuring and monitoring applications (workloads). A documentation page⁷ was created to enable 3rd party interactions by providing both a complete set of instructions and the data types that are being used as well as a sandbox environment where the user can interact without coding or an external GUI for REST calls.

intent	^
GET /intent/ List/Creates Intents.	intent_list 🗸 🗎
POST /intent/ List/Creates Intents.	intent_create 🗸 🔒
POST /intent/template/	intent_template_create 🗸 🇎
GET /intent/types/	intent_types_list 🗸 🇎
PUT /intent/{id}/action/	intent_action_update 🗸 🔒
PUT /intent/{id}/target/ Updates an intent target with a new value.	intent_target_update ∨ 🔒
workload	^
GET /workload/	workload_list 🗸 🔒
POST /workload/	workload_create 🗸 🍵
GET /workload/instance/	workload_instance_list 🗸 🍵
PUT /workload/instance/{instance_id}/delete/	workload_instance_delete_update 🗸 🍵
GET /workload/instance/{instance_id}/manifests/	workload_instance_manifests_list 🗸 🍵
POST /workload/upload/ Upload a workload document underlying helm chart.	workload_upload_create 🗸 🍵
GET /workload/{id}/ Retrieve, update and delete a workload document	workload_read 🗸 🔒
PUT /workload/{id}/ Retrieve. update and delete a workload document	workload_update 🗸 🔒
PATCH /workload/{id}/ Retrieve, update and delete a workload document	workload_partial_update 🗸 🔒
DELETE /workload/{id}/ Retrieve, update and delete a workload document	workload_delete 🗸 👜 💱
POST /workload/{id}/template/ Renders kubernetes manifests for given workload document.	workload_template_create 🗸 🔒

ighte 12. Intent-Based API endpoints

Figure 14, presents a sample of the expected parameters' body of the workload POST endpoint.

⁷ <u>https://intent-api.nemo.onelab.eu/api/v1/swagger/</u>

2er

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	24 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



POST /w	orkload/		workload_create 🔨
st or Create a	new workload document(s)	
arameters			Try it out
me	Description		
ta * required	Example Value Model		
ody)	WorkloadDocumen	tCreate v (string title: Name moderapth: 100	
	version*	Interview Cool document name string title: Version mode(negth: 32 min(negth: arts) string title: Version	
	schema*	Schema v { description: The document schema	
	type	} string title: Type The workload document type Enum:	
	intents	<pre>> Array [1] → [f supported intents</pre>	
		Intents v string title: Intents maxiength: 100 minlength: 1 Intent userlabel	
	ingress_support	boolean title: Ingress support Whether the workload document can be exposed via NEMO	
	}		

2.3 Cloud/Edge/IoT Integration and Validation Infrastructure

The complete integration of different systems in the Cloud-Edge-IoT continuum into a single framework is a very complex task and requires a lot of effort and precision. The individual systems in each layer of the continuum Cloud, Edge and IoT are designed to provide max throughput, low latency and increased processing capabilities to meet the requirements of a modern system. While these systems can flawlessly work individually, several issues may arise when they need to be interconnected. NEMO is a complex framework that aims to provide severe automations at a low level across the continuum in terms of deployment and device management. In a typical Cloud-Edge-IoT continuum, the cloud computing layer provides a centralized environment where the main processing power and storage space exist. On the contrary, Edge serves, as the frontiar of processing as they are located close to the IoT devices. These are responsible for processing the incoming data and taking decisions in real-time.

Within the NEMO framework an operational management of the Cloud-Edge-IoT continuum takes place. The infrastructure where the NEMO meta-OS platform is being deployed varies based on the equipment present on the Pilots' premises. To guarantee a complete and flawless integration of the platform, a set of minimum requirements was distributed among the partners. This guaranteed that the underlying technologies, such as Kubernetes can work without issues. Typically, the main issue within the cloud-Edge IoT continuum interactions is the connectivity due to the communications taking place over different technologies with varying capabilities, e.g., wireless channels, low bitrate channels, etc. In NEMO, this was considered since the design phase where lightweight communication protocols and messaging mechanisms, such as RabbitMQ [16] that are resilient to packet loses and require a small set of resources. These channels within the proposed meta-OS are used to exchange real-time information about the resource availability on the connected clusters which is crucial for the NEMO users as it can be used for taking strategical decisions on where a workload should be deployed. In addition, the channels are able to handle high volume data images for migrating the workloads along with the related data volumes.

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	25 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Integration & V&V Methodology & Plan 2.4

Since the NEMO conceptualization phase, a detailed plan for the integration and validation phases was created. This plan was initially presented in deliverables D1.3 and D4.1 while an updated version reporting the alignment with that plan as well as the updates of the tools was presented in deliverable D4.2. The plan was aiming to create a universal cross component and pilot template for monitoring the various phases of development to keep track of the updates and to identify any risks arose. The plan adopted is based on the best modern practices of software development, NEMO followed an agile and incremental approach of iteration cycles, grouped in 3 Phases, as depicted in Figure 15.

Phase 1: Baseline (M1-M18). Provides the initial NEMO Proof of Concept. Phase 1 starts with system, specification of the meta-OS Architecture and decomposition (WP1), design analysis prototyping (WP2-WP4), integration, testing and validation of all key meta-OS components (WP4), The outcome will be NEMO Ver. A and initial Living Labs validation and the selection of the new consortium members and new components from Open Call #1 to be implemented with Phase

Phase 2: Advance (M19-M30). All NEMO components are further developed (WP2-WP4), while NEMO is expanded with new functionality added from the new consortium members accepted via Open Call #1. Stronger integration with 5G networks and MANO systems will be realized and validated in Living Labs). The outcome will be NEMO Ver. B and Living Labs validation, along with new AIoT applications and services from Open Call #2.

Phase 3: Mature (M30-M36). Focus on validation and optimization, and more realistic field conditions testing and verification, not only from NEMO consortium buy also from 3rd parties selected via Open Call #2, increasing system TRL and preparing NEMO Ver 1.0, validated in Living Labs. This phase also strengthens activities related to engagement of open-source communities and relevant initiatives, ensuring accessibility, sustainability and availability in open-source platforms.





At the competition of each phase a retrospective analysis along with a risk management assessment took place. This guaranteer that the project would proceed based on the plan conducted at the beginning of the project and any deviation or risk appeared during the development phase would be immediately ideratified and resolved without jeopardizing the progress of the project. The template presented in the following subsection was used to document and monitor any issues occurred.

2.4.1 NEMO V&V documentation

The NEMO End-to-End integration pipeline is a complex procedure; in order to provide high-quality tests, the End-to-End pipeline was divided in smaller and simpler test pipelines. Associated technical details and stemming results are provided in Section 5 of this document. As mentioned earlier, a common template applicable across the different functional scenarios present in NEMO was defined. This played a pivotal role as it allowed to properly document all the possible steps and interaction took place in an organized and well-written manner. At each iteration phase presented before, the following template was evaluated, and new features were introduced in case a feature was missing or modified in case it

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	26 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



was not adding value to the project. To better understand what needs to be done as well as to verify the goals' success, three main attributes were defined:

- Scenario
- Outcome
- Checkpoints

In Table 2 the template for the integration tests explaining the scenarios and what needs to be tested along with the functional and non-functional requirements for each scenario are documented. Moreover, an overview of the test plan is present to assist the engineers to perform the tests. After the completion of a test procedure, the involved partners were filling in the integration scenario checkpoint template, Table 3, which was later used in the retrospective analysis to provide information about the open ispres.

Table 2. Integration testing - scenario template

Test 1:	
Objective	
Components	
Requirements alignment	
Features to be tested	
Test setup	
Steps	1. 2. 3.

Table 3. NEW integration testing - Checkpoints template

Che	Checklist for Test1						
		Yes	No	Comments			
1	Is a service created?	\checkmark					
2	Is the device registration completed successfully?	 					
3	Is the device sending its data successfully?	~					
4	the data stored in Database / Registry?	\checkmark					

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				esults.	Page:	27 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



3 NEMO Integrated Platform (Final Version)

The final version of the NEMO meta-OS was established following a comprehensive integration and validation procedure. The conducted integration tests concerned both the NEMO developed technical solutions and the 3rd party ones introduced by the 2 Open Calls.

3.1 Meta-OS functionality in NEMO ver. 1.0

3.1.1 NEMO Infrastructure Management

The infrastructure management across the various Kubernetes clusters has been designed for scalebility, flexibility, and high availability, utilizing a multi-cluster approach. Several clusters have been deployed, each tailored for specific roles such as control planes, worker nodes, and specialized hardware, including GPU-enabled nodes.

• Nemo main cluster: A control-plane node (nemo-dev-master) and live worker nodes (nemodev-worker1 to nemo-dev-worker5) have been configured within this development cluster, running Kubernetes version 1.29.13. This setup provides a robust environment for development and testing purposes. Table 4 provides a detailed view of the technical specifications of the NEMO main cluster.

Table 4. NEMO dev cluster (K8S

Node name	Node Type	Specifications
k8smaster.onelab.eu	Master	 CPU: 8 CPU Cores RAM: 16GB Storage: 140GB Ephemeral OS-Image: Ubuntu 22.04.4 LTS Kernel Version: 5.15.0-116- generic Container Runtime: containerd://1.6.28
k8sworker1.onelabeu	Worker and Storage	 CPU: 16 CPU Cores RAM: 32GB Storage: 120GB Ephemeral + 150GB Ceph OS-Image: Ubuntu 22.04.4 LTS Kernel Version: 5.15.0-116-generic CONTAINER-RUNTIME: containerd://1.6.28
k8sworker2.onelab.eu	Worker and Storage	 CPU: 16 CPU Cores RAM: 32GB Storage: 120GB Ephemeral + 150GB Ceph OS-Image: Ubuntu 22.04.4 LTS Kernel Version: 5.15.0-116-generic CONTAINER-RUNTIME: containerd://1.6.28

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	28 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Node name	Node Type	Specifications
k8sworker3.onelab.eu	Worker and Storage	 CPU: 16 CPU Cores RAM: 32GB Storage: 120GB Ephemeral + 150GB Ceph OS-Image: Ubuntu 22.04.4 LTS Kernel Version: 5.15.0-116-generic Container Runtime: containerd://1.6.28
k8sworker4.onelab.eu	Worker and Storage	 CPU: 16 CPU Cores RAM: 32GB Storage: 120GB Ephemeral + 150GB Ceph OS-image: Ubuntu 22.044 LTS Kernel Version: 5.15.0-116-generic Contained Runtime: contained: 71.6.28
k8sworker5.onelab.eu	Worker and Storage	 CP0-16 CPU Cores RAM: 32GB RAM Storage: 120GB Ephemeral + 150GB Ceph OS-image: Ubuntu 22.04.4 LTS Kernel Version: 5.15.0-116-generic Container Runtime: containerd://1.6.28

• Nemo Staging 1 Cluster: Focused on scaling operations, this cluster comprises a control-plane node (nemo-s1-master) and several worker nodes (nemo-s1-worker1 to nemo-s1-worker3), running Kubernetes versions 1 207 and 1.31.3, thus ensuring reliability for both development and production workleads. Table 9 presents the technical details of the NEMO Staging 1 cluster.

	· · · · · · · · · · · · · · · · · · ·	Table 5.	Staging 1 cluster (K8S)
	Node Name	Node Type	Specifications
R	neme d-master	Master	 CPU: 8 CPU Cores RAM: 16GB Storage: 120GB Ephemeral + 150GB Ceph OS Image: Ubuntu 22.04.3 LTS Kernel Version: 5.15.0-78-generic Container Runtime: containerd://1.7.12
,	nemo-s1-worker1	Worker and Storage	 CPU: 16 CPU Cores RAM: 16GB Storage: 120GB Ephemeral + 150GB Ceph OS Image: Ubuntu 22.04.3 LTS Kernel Version: 5.15.0-78-generic Container Runtime: containerd://1.7.12
	nemo-s1-worker2	Worker and Storage	 CPU: 16 CPU Cores RAM: 16GB

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Page: 29 of 111					29 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Node Name	Node Type	Specifications
		 Storage: 120GB Ephemeral + 150GB Ceph OS Image: Ubuntu 22.04.3 LTS Kernel Version: 5.15.0-78-generic Container Runtime: containerd://1.7.12
nemo-s1-worker3	Worker and Storage	 CPU: 16 CPU Cores RAM: 16GB Storage: 120GB Ephemeral + 150GB Ceph OS Image: Ubuntu 22.04.3 LTS Kernel Version: 5.15.0-78-generic Container Runtime: containerd://1.7.12

• Nemo Staging 2 Cluster: The lightweight k3s-based cluster consists of a control-plane node (nemo-k3s-master) and three worker nodes (nemo-k3s-node-1 to nemo-k3s-node-3), operating on version 1.30.6+k3s1. A specialized AMD-based worker node (nemo-k3s-node-3) was added to support the integration of the NEMO component, ensuring optimal performance for resource-intensive tasks. More details about the Staging 2 cluster are presented on Table 6.

	Node name	Node Type	Specifications
	nemo-k3s-master	Master	• CPU: 4 CPU Cores
			• RAM: 8 GB
			Storage: 64 GB External SSD
			Os Image: Ubuntu 24.10
			• Kernel Version: 6.11.0-1004-raspi
			Container Runtime: containerd://1.7.22-k3s1
	nemo-k3s-node-1	Worker	CPU: 4 CPU Cores
		and	• BAM: 8 GB
		Storage	Storage: 1TB External SSD
			• OS Image: Ubuntu 24.10
			Kernel Version: 6.11.0-1004-raspi
	• •		• Container Runtime: containerd://1.7.22-k3s1
	nemo-k3s-node-2	Worker	• CPU: 4 CPU Cores
		and	• RAM: 8 GB
		Storage	• Storage: 64 GB External SSD
			• OS Image: Ubuntu 24.10
			• Kernel Version: 6.11.0-1004-raspi
			• Container Runtime: containerd://1.7.22-k3s1
()	nemo-k3s-node-3		• CPU: 16 CPU Cores
X			• RAM: 16 GB
			• Storage: 120GB Ephemeral + 150GB Ceph
			• OS Image: Ubuntu 24.10
			• Kernel Version: 6.11.0-19-generic
			• Container Runtime: containerd://1.7.12

• **Nemo Production Cluster**: Dedicated to production workloads, this cluster includes a GPUenabled worker node (nemo-prod-gpu-worker) and several worker nodes (nemo-prod-worker1 to nemo-prod-worker3), along with a control-plane node (nemo-prod-master), all operating on Kubernetes version 1.30.7. This configuration ensures high-performance processing for tasks such as machine learning. The NEMO Production cluster details are demonstrated on Table 7.

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						30 of 111
Reference:	D4.3 Dissemination: PU Version: 1.0				Status:	Final	

Table 6. Staging 2 Cluster (K3S)



Node Name	Node Type	Specifications
nemo-prod-master	Master	 CPU: 4 CPU Cores RAM: 8 GB Storage: 80 GB Ephemeral OS Image: Ubuntu 22.04.3 LTS Kernel Version: 5.15.0-78-generic Container Runtime: containerd://1.7.12
nemo-prod-worker1	Worker and Storage	 CPU: 8 CPU Cores RAM: 16 GB Storage: 250GB Ephemeral + 150GB Ceph OS Image: Ubuntu 22.04.3 LTS Kernel Version: 5.15.0-78-generic Container Runtime: containerd://1.1.12
nemo-prod-worker2	Worker and Storage	 CPU: 8 CPU Cores RAM: 16 GB Storage: 120GB EphemerAl + 150GB Ceph OS Image: Ubuntu 22,04.4 DFS Kernel Version: 5.160-78 generic Container Runtime: containerd://1.7.12
nemo-prod-worker3	Worker and Storage	 CPU: 8 CPU Gores RAM: 16 GB Storage: 120GB Ephemeral + 150GB Ceph US Image! Ubuntu 22.04.2 LTS Kernel Version: 5.15.0-78-generic Container Runtime: containerd://1.7.12
nemo-prod-gpu- worker	Worker and Storage	 CPU: 4 CPU Cores RAM: 8 GB Storage: 120GB Ephemeral + 150GB Ceph OS Image: Ubuntu 22.04.4 LTS Kernel Version: 5.15.0-78-generic Container Runtime: containerd://1.7.12

Table 7. Production Cluster (k8S)

To mitigate the impact of power outages, dedicated UPS systems have been implemented for each host running virtual machines that support the cluster nodes. This solution ensures uninterrupted operation during power failures, effectively reducing downtime. These infrastructure improvements, along with the addition of specialized hardware for NEMO integration, have significantly enhanced the resilience, performance, and reliability of the entire system.

3.2 NEMO Open Call integration activities

In order to provide homogenised, flexible orchestration of varied workloads across heterogeneous and scattered devices, the concept of the meta-Operating System (meta-OS) refers to the efficient integration of highly diversified hardware and software resources. Because of NEMO's expandable architecture, additional features can be added as NEMO plugins. Since the NEMO Kernel serves as the core system, plugins are designed to add capabilities to the core, such as flexibility, extensibility, and isolation of bespoke meta-OS logic or applications. Horizontal or domain-independent services that seek to offer some fundamental and standard features that expand the NEMO capabilities are referred to as plugins.

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						31 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



NEMO has already realized extensions through additions developed by the 1st Open Call projects. NEMO 1st Open Call invited SMEs active as edge computing, edge and/or native cloud software development, operating systems, IoT/5G networks and IoT manufacturing entities to join the NEMO ecosystem by offering: a) NEMO meta-architecture extensions b) software components/plugins not covered by current NEMO implementation plan c) new network or service/resources metering/automated control components or d) porting NEMO on new, highly heterogeneous IoT devices.

Six winning proposals have been identified, specifically MetaFOX, by Vodéna, Serbia working on Machine Learning and AI technologies; CorMOS, by Business & IoT Integrated Solutions, Cyprus focusing on cross-domain orchestration of Meta-OS edge resources; ARGO, by Intellia ICT, Greece, investigating porting AR/XR technologies in the meta-OS; Eros4NRG, by Martel Innovate BV, Netherlands offering ZeroTrust IoT Analytics with focus on Smart Energy applications; GENESYS, by SWHARD srl, Italy providing an edge gateway for the NEMO Meta-OS; and MARINE110, by BEAM Innovation, Romania working on an efficient resource utilization and maritime network slicing plugin for the NEMO Meta-OS.

In addition, NEMO introduced Open Call #2 in the framework of which the final xed NEMO integrated framework will be further validated, as the NEMO Open Call #2 projects with utilize the NEMO meta-OS functionality supporting the various use cases realized by the 3rd parties.

The rest of the chapter provides a brief description of the NEMO open call #1 projects' technical concept and describes the integration details of each technical solution developed as a NEMO meta-OS architecture extension.

3.2.1 ARGO

The ARGO project, standing for "Augmented Reality next Generation Operational Systems in NEMO", is a pioneering initiative designed to align seamlessly with the objectives and scope of the Next Generation Meta Operating System (NEMO) program. The motivation behind ARGO's proposition lies in its integration with the NEMO architecture, enhancing the existing framework with innovative AR solutions. At the core of ARGO is the integration of NEMO's flexible meta-architecture with specialized AR devices, such as the Vuzix M4000 and Vuzix Shield, and the backend support system of the AR devices. These lightweight, monocular, and binocular AR glasses are designed for prolonged industrial use, offering immersive experience without compromising user comfort during full-shift operations. ARGO's target is to enhance industrial efficiency, safety, and training using advanced AR technology. The ARGO system has been already deployed on the NEMO provided development environment in OneLab and preliminary tests of the system on the premises of the Foundation of the Hellenic World (FHW) in the context of Smart Media/ City & XR Living Lab were conducted.

3.2.2 CorMQ6

The CorMOS Orchestration Engine is a specialized software component developed to facilitate the deployment of applications to appropriate cluster nodes based on user-defined requirements, goals, and intents. The functionality of this software component relies on data stored in a JSON file that captures the application's or users' goals/intents in a platform-independent way, along with a YAML file that specifies the internal components of the application, their interactions, and additional resource constraints. The CorMOS Orchestration Engine further analyses the data in the JSON file and updates the YAML file with the appropriate instructions (e.g., node selectors and/or affinity rules) to ensure that the components of the application are deployed to the correct nodes of the cluster, thus achieving user goals such as performance optimization and efficient resource utilization. Moreover, the CorMOS developed Telemetry Data Management system focuses on a few parameters of particular interest that directly affect the behavior of the CorMOS Orchestration Engine.

The CorMOS orchestration engine will be integrated with the NEMO architecture and its logical placement would be within or alongside the Meta-Orchestrator (MO). The Meta-Orchestrator plays a

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						32 of 111
Reference:	D4.3 Dissemination: PU Version: 1.0					Status:	Final



critical role in coordinating workload execution, deployment, and migration across the IoT-Edge-Cloud continuum. Given its pivotal function, the CorMOS Orchestration Engine could interact with key NEMO components to ensure seamless integration and to leverage the platform's existing capabilities effectively. From an architectural standpoint, the CorMOS Orchestration Engine should be positioned as a high-level component within the orchestration layer. This placement would allow it to either extend the Meta-Orchestrator (MO) functionality while serving as the decision-making hub for workload execution in context-specific infrastructures and environments. To achieve optimal orchestration, the CorMOS Orchestration Engine should retrieve information from various monitoring, networking, security, and policy enforcement components, ensuring that workload placement and execution are conducted in an intelligent, secure, and resource-efficient manner.

3.2.3 Eros4NRG

Eros4NRG is a platform designed to monitor and analyse energy production and construction, enabling more efficient decision-making through predictive analytics. It acts as a one-stop-shop (OSS) for energy stakeholders by integrating data from electric vehicles (EVs), EV charging stations, photovoltaic (PV) plants, and smart headquarters. The platform enhances trust and transparency in energy data, while addressing key challenges such as: (i) Unreliable data from IoT sources impacting machine learning operations; (ii) poor organization of static energy data leading to information loss; and (iii) the need for user-friendly AI/ML insights for non-technical stakeholders.

Eros4NRG's final integration is guided by a set of clear functional and non-functional requirements, ensuring that the system operates efficiently across modern cloud-native infrastructures (such as MinIO, PostgreSQL, and Docker). Designed with a service-oriented architecture, the platform enables streamlined data collection, processing, visualization, and security. The Eros4NRG project relies on extracting raw data from EMOTION and ASM ferni systems via API calls. This continuous stream of data enables the platform to monitor and analyse energy consumption and production, as well as EV battery performance in (near) real time. The module starts by sending API requests to the EMOTION and ASM Terni systems. Building on the final integration, Eros4NRG will undergo several enhancements and additions to maximize nepotential and improve user experience.

3.2.4 Genesys

GENESYS project aims to extend NEMO's scope and technology by porting the meta-operating system on a new IoT Edge device, fully customisable and industry-ready, developed by SWHARD. Targeting the Industry 5.0 and G-IoT paradigms, aligned with the NEMO *Smart Manufacturing & Industry 4.0 Use Cases & Living Leb*, the project's primary goal is to install NEMO on an industrial-grade Systemon-Module (SOM), delivering a functional product while providing extensive documentation for seamless installation on Edge platforms.

FLEX is categorised as a micro-edge system. Conventionally, the purpose of this class of edges is to collect data from field sensors, do some light pre-processing and transfer aggregated information to the cloud. Despite the low resources available (by definition), the expected average workload for the CPU is usually very low in a common edge-cloud design.

In NEMO's perspective the micro-edge is a cluster node which still serves the two purposes above. Being part of the cluster, it will also provide processing power not only for the elaboration of its "own" sensors, but for other tasks. The Meta-Orchestrator is responsible to decide, deploy and monitor the tasks to be run in every node, including the micro-edges. This concept leads to a greater optimization of all the available resources in the cluster.

Following the MVP implementation of the Genesys framework, the successful porting of the Software components required to let the FLEX enter the NEMO META-OS ecosystem, is the keystone of GENESYS. Although the work has not yet demonstrated with a real application, the FLEX nodes have

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						33 of 111
Reference:	D4.3 Dissemination: PU Version: 1.0					Status:	Final



been successfully joined the test cluster, ran a simple application and a set of benchmarks. A side activity that has been done was the test of the integration of Coral NPU into the cluster.

3.2.5 MARINEMO

The MARINEMO Slice Manager Plugin (SMP) is designed as an add-on for 5G-enabled cellular networks, its main purpose being to efficiently optimize the communication resources of the targeted systems by applying AI/ML-powered intents for slicing and user profile reconfigurations.

The API Engine of the Slice Manager Plugin facilitates external interactions, providing a crucial interface for third-party systems to influence and monitor the network's behaviour. It offers APIs for defining testbed-specific operational parameters, alerting mechanisms, and performance analysis. This component not only enhances the plugin's adaptability to different environments but also ensures that it can serve a broad range of use cases from operational adjustments to detailed performance analysis.

Regarding the Slice Manager Plugin integration points, the plugin needs to be connected to testbedspecific agents for network API exposure and performance monitoring. Specifically, the MARINEMO Slice Manager Plugin will be integrated within META-OS ecosystem through the Telefonica Private 5G testbed and its Cumucore 5G Core Network APIs. Furthermore, the SMP while instantiated on top of a container orchestration platform, more specifically, the OneLab Kuberneter cluster.

The MARINEMO SMP will be deployed on the META-OS Kubernetes based virtualization platform by following the next steps.

- Access to the Kubernetes cluster for testing the pugin for testing and validating the SMP interactions with the META-OS APIs the plugin docker-based container will be directly deployed on the NEMO Kubernetes cluster Access to the META-OS Docker Hub repository for testing and validating the deployment capabilities of the NEMO META-OS platform, access to the META-OS Docker Hub repository will be used for uploading the SMP docker image and instantiating the container on the Rubernetes cluster by running the META-OS Gitlab CD scripts.
- Access to the Gitlab META-OS repository for complete CI/CD for testing and validating the complete CI/CD capabilities of the META-OS platform, access to the project-related Gitlab repository will be used for uploading all the SMP source code and running the CI scripts for creating the plugin docket image and the CD scripts for instantiating the container on the Kubernetes cluster.

3.2.6 MetaFOX

The MetaFOX project⁸ is an advanced automated machine learning (AutoML) component, which significantly simplifies the initial model creation within NEMO META-OS by automating the process of model selection, feature engineering, and hyperparameter tuning. AutoML is a transformative approach that simplifies complex machine learning processes, enabling both experts and non-specialists to design and deploy models efficiently. MetaFOX uses the power of automation in machine learning to streamline model development, enhance the quality of models, and democratize AI accessibility. It significantly simplifies the initial model creation for federated learning (FL) and transfer learning (TL) by automating the process of model selection, feature engineering, and hyperparameter tuning.

The MetaFOX solution, when deployed on Kubernetes, consists of a collection of components running in separate Pods, each governed by Deployments that handle the creation and scaling of these Pods. Services provide stable network endpoints for communication within the cluster, while Persistent Volumes (PVs) and Persistent Volume Claims (PVCs) ensure the durability of data for stateful components. In the context of the deployment, Minikube was used as the Kubernetes environment on a

⁸ <u>https://metafox.readthedocs.io/en/latest/</u>

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						34 of 111
Reference:	D4.3 Dissemination: PU Version: 1.0					Status:	Final



Jor J

dedicated Linux server, providing a local cluster suitable for both development and production-like evaluations.

Preliminary integration tests were conducted achieving a successful deployment providing for early validation of the component's compatibility with OneLab's infrastructure. The successful deployment carried out has provided sufficient evidence of the MetaFOX component's compatibility with the OneLab environment. These results support the feasibility of full-scale integration, and they lay the groundwork for further testing and optimization. Future work will focus on consolidating these findings into a robust deployment procedure tailored for OneLab.

 Document name:
 D4.3 Advanced NEMO platform & laboratory testing results. Final version
 Page:
 35 of 111

 Reference:
 D4.3
 Dissemination:
 PU
 Version:
 1.0
 Status:
 Final



4 NEMO Service Management Layer updates

This section reports the latest specifications and design options for the NEMO components of the Service Management layer in the NEMO architecture. These components provide middleware between core NEMO functionality and workloads but also end users. They support ZeroOps principles and expose interfaces to external entities (services or users). Moreover, the supported services include Lifecycle Management and DLT-based accountability of workload or infrastructure usage and collectively contribute to NEMO openness and adoption by third parties, referring to appreading or infrastructure owners, as well as developing entities.

4.1 Intent-based Migration Controller

4.1.1 Overview



The Intent-based Migration Controller (IBMC) is responsible for orchestrating the seamless transfer of workloads across the IoT, Edge, and Cloud Continuum. Rather than relying an manual configurations, it uses intent-based networking to automatically interpret high-level objectives and translate them into concrete migration actions. This allows the system to adapt to changing conditions, ensure efficient use of resources, scale effectively, and uphold service reliability. The IBMC's tole is essential in managing the highly dynamic and distributed nature of the meta-OS environment.

4.1.2 Architecture

The IBMC architecture is represented in Figure 16, which follows the hierarchical structure of the Open Cluster Management (OCM) HUB-managed cluster⁹. The key sub-components that make up this architecture are described below:

- **Ibmc-controller**: This component serves as the entry point for initiating the workload migration process. Upon receiving a new intent from the Intent-Based API, the ibmc-controller checks the current workload status and determiner whether it complies with the specified intent parameters. If any of these parameters are not satisfied, the controller selects a suitable target cluster from the list of available managed clusters and sends a migration trigger message to the ibmc-agent deployed in the selected cluster.
- **Ibmc-agent**: The ibme-agent functions as a black-box component responsible for executing the actual migration process. Each cluster runs its own ibmc-agent, which listens to a dedicated RabbitMQ queue for incoming migration messages. When a message arrives at the source cluster, the corresponding ibmc-agent initiates a backup of the specified workload using Velero¹⁰ Upon successful completion of the backup, a message is sent to the target cluster, where the local ibmc-agent receives the message and proceeds to restore the workload actordingly.

Velero: Velero is a key technology in the workload migration process, as it is deployed across all clusters and ensures that workloads can be efficiently backed up and restored, enabling eamless transitions between clusters. Its integration into the migration workflow allows for consistent data protection and minimizes downtime during migrations.

• **Rook Ceph:** Rook Ceph¹¹ is used as an S3-compatible object store for Velero, providing reliable and scalable storage for backup data. By integrating with Velero, Rook Ceph ensures that

¹¹ https://rook.io/docs/rook/latest/Storage-Configuration/Object-Storage-RGW/object-storage/

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						36 of 111
Reference:	D4.3 Dissemination: PU Version: 1.0					Status:	Final

⁹ <u>https://open-cluster-management.io/docs/concepts/architecture/</u>

¹⁰ <u>https://velero.io/</u>




workload snapshots and metadata are securely stored and readily available for restoration during migration processes.

Figure 17, depicts the sequence diagram that concerns the NEMO workload migration process. The corresponding integration and validation test are presented in section 5 of the document.



Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	37 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final





4.1.3 Interaction with other NEMO components

- 1. The Intent-API publishes in the notify' Rabbitmq queue an Intent with one or more requirements (Availability, GreenEnergy, CPU, Memory etc).
- 2. The ibmc-controller retrieves vorkload status from the Intent-API.
- 3. The ibmc-controller retrieves from the MO the information related to each of the managed clusters.
- 4. If the workload is already deployed in any cluster and the intent requirements are not met, then a migration action is triggered.
- 5. A new target cluster that complies with the intent requirements is selected from the available Managed Clusters list.
- 6. The ibne-controller sends a message via RabbitMQ queue to the ibme-agent containing the workload D and the selected target cluster.

backup is created with all of the workload associated resources and it is uploaded to the Rook Ceph S3 Bucket located in the HUB cluster.

- 8. When the backup is completed, a message is sent to the target cluster's ibmc-agent to continue with the migration process.
- 9. The backup is restored in the target cluster, creating all the workload resources.
- 10. A message is sent to the MO notifying about the workload migration completion.
- 11. The MO updates the corresponding Manifestwork so that it matches the new workload deployment status.
- 12. A similar message is sent to the Intent-API updating the workload status, specifying the cluster where it has been deployed.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	38 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



4.1.3.1 Meta-Orchestrator (MO)

The IBMC integrates with the MO primarily to gather cluster-related data and coordinate migration execution. The ibmc-controller queries the MO to retrieve up-to-date information about all managed clusters, which it uses to evaluate compliance with intent requirements. After a migration is completed, the ibmc-agent notifies the MO, which then updates the ManifestWork to reflect the workload's new deployment state.

4.1.3.2 Intent-Based API

The Intent-Based API maintains all information related to deployed workloads and is responsible for creating and managing the intents that initiate workload migrations. When a migration is triggered, the intent's requirements are used to identify the new target cluster, which must meet all specified criteria. Upon completion of the migration, the ibmc-agent sends a notification to the Intent-API, updating the workload's status, including the new cluster where it has been deployed. This ensures that the Intent-API remains synchronized with the latest workload deployment state.

4.1.3.3 MNCC

The notify queue to which the ibmc-controller listens, is also used for the L2S-M network creation process which also relies on Intents. The creation of these networks, like workload migrations, is triggered by specific intent requirements. For more details, refer to section 2.2.

4.1.3.4 CFDRL

The CFDRL uses machine learning and relies on trained models to detect the sub-optimal workload operation, leading to the triggering of the workload migration. When this occurs, the CFDRL sends the required information to the IBMC, in order perform the migration.

4.1.4 Conclusion

The IBMC, has completed its functionality implementation and has been successfully integrated with the other components, including MO, Intern Based API, mNCC, and CFDRL. This integration enables the IBMC to support seamless workload deployment and migration across the IoT-to-Edge-to-Cloud continuum, maintaining operational efficiency within the dynamic meta-OS environment.

The functionality of the Intent Based Engration Controller (IBMC) has been significantly extended to support intelligent redeployment decisions driven by dynamic, user-defined intents. This enhancement enables the IBMC to continuously monitor the operational environment and autonomously evaluate incoming intents, which encapsulate high-level objectives such as performance optimization, energy efficiency or latency reduction requirements.

The addition of this intent-driven redeployment logic marks a shift toward a more proactive and adaptive control mechanism within the meta-OS framework, ensuring that service delivery remains aligned with evolving user needs and operational constraints.

The IBNC is deployed across the NEMO pilots' infrastructures, where it will be validated through pilotspecific use cases to demonstrate the IBMC's potential to enhance service continuity and resource efficiency in complex, distributed computing environments.

4.2 Plugin & Applications Lifecycle Manager

The Plugin & Applications Lifecycle Manager (LCM) is a flexible and unified solution for managing plugins and applications across the NEMO ecosystem. Acting as the interface between NEMO users and the system, the LCM enables on-demand deployment of workloads (services, applications, and plugins) within the NEMO environment.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	39 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



While workloads operate on the NEMO meta-OS, an event-driven system continuously monitors key performance metrics. Simultaneously, a dedicated security controller tracks security-related events, alerts users to anomalies, and takes proactive steps to mitigate potential cyber threats.

The LCM's user interface is designed to integrate seamlessly with other core NEMO components - such as the Intent-based API, PPEF, MOCA, and CMDT - delivering unified and intuitive user experience. It provides access to features including user profile management, workload deployment and monitoring, security oversight, and historical data analysis, all tailored to specific user roles of NEMO ecosystem.

4.2.1 Architecture

The LCM comprises of a set of subcomponents namely the LCM CD, LCM Controller, Security Controller, Event-based Response, LCM Repository and LCM Dashboard.

The final LCM high-level architecture of NEMO meta-OS is depicted in the development view diagram in Figure 18. In brief, the software components that make up the LCM module are the following.

LCM CD is based on ARGO CD framework [17] [18] [19] to manage NEMO workbads provided by NEMO partners or NEMO Open Call participants and deploys workloads in 53 bucket container.

LCM Controller is a control mechanism that facilitates communication between LCM submodules and the NEMO ecosystem, offering endpoints for sending and receiving information.

Security Controller handles runtime security monitoring of NEMO workloads, notifying both users and relevant NEMO components of detected events.

Event-based Response module is designed to implement automated actions in response to events initiated by user input or detected by other NEMO components.

LCM Repository is used to store data related to workload lifecycle, security incidents, detected events and other workload related information to provide historical analysis and runtime statuses.

LCM Dashboard serves as the gateway between end-users and the NEMO meta-OS ecosystem, granting privileged users access to manage their workloads and monitor both performance and security.



Figure 18. LCM high-level architecture

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	40 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



4.2.1.1 LCM Repository

LCM Repository uses Elasticsearch¹² for storing, searching, and analysing data provided by various NEMO components like Intent-based API, PPEF, MOCA and CMDT. Elasticsearch provides fast search responses and comes with extensive REST APIs for storing and searching the data. Stored data include the status and lifecycle of the workloads, security events detected, workload performance and resource usage.

4.2.1.2 LCM Visualization

LCM Visualization serves as the primary interface of the NEMO ecosystem, offering tailored tools for each user role to manage workloads and resources within the NEMO meta-OS and monitor activities. Its goal is to deliver an intuitive and seamless experience across the NEMO ecosystem, catering to both advanced users and those with less technical expertise. The interface presents relevant information about the workload lifecycle, usage, and security of workloads and resources in a clear, consist format, with multiple levels of detail (e.g., workload, resource, user, and system views). Detailed explanation of the LCM views were provided in section 2.2, as part of the 3rd party documentation and are also described extensively in the context of section 5.

4.2.2 Conclusion

The LCM component has been successfully integrated into the NEMO ecosystem, deployed in OneLab environment, providing features that include:

- Workload lifecycle management and monitoring
- Resources provisioning and usage monitoring
- Security/vulnerability scanning and morntoring
- Visual interface for experienced and non-experienced users

The final result offers seamless workload deployment and monitoring through an intuitive interface.

4.3 Monetization and Consensus-based Accountability

4.3.1 Overview

The Monetization and Consensus-based Accountability (MOCA) component lies at the heart of NEMO's pre-commercial ecosystem, providing a secure, transparent, and equitable framework that enables both resource providers (offering compute, storage, network, or data) and service consumers to seamlessly earn and spend, 'credits' across the AIoT–Edge–Cloud continuum. As the backbone of NEMO's Resource-ana Service (RaaS) marketplace, MOCA guarantees timely, accurate reward settlement for providers and precise billing for consumers, all underpinned by a tamper-proof audit trail maintained on a distributed ledger.

Since the previous deliverable, D4.2, significant progress has been made to expand MOCA's capabilities and refine its functionalities. Key advancements include:

- **Mart Contract Deployment Tool**: Implementation of a dedicated tool to streamline the deployment and management of smart contracts utilized by MOCA.
- ML-as-a-Service (MLaaS) Smart Contract: Development and deployment of a specific smart contract tailored for the accounting and monetization of Machine Learning services offered and consumed via the NEMO platform.
- Enhanced Reward Mechanisms: Introduction of more sophisticated smart contract functionalities to accurately reward service providers based on various contribution factors.

¹² https://www.elastic.co/elasticsearch

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	41 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



• Network Resource and Data Monetization: Extension of the MOCA framework to encompass the accounting and potential monetization of network resources and data assets within the NEMO ecosystem.

These enhancements aim to increase MOCA's flexibility, broaden the scope of monetizable resources and services, and support more complex business models emerging within the NEMO continuum, further solidifying its role in facilitating a trustworthy Resource-as-a-Service (RaaS) marketplace.

4.3.2 Architecture

MOCA's architecture has been described in detail in D4.2. In the current document, there is one notable addition to the component's architecture, namely the Smart Contract Deployment Tool. In Figure 19 the updated architecture diagram for MOCA is depicted. The Smart Contract Deployment Tool is responsible for allowing the users to upload their own contracts to the Smart Contracts component and allow them to enhance the tool with new accounting logic.

Regarding the workflow presented in the deliverable D4.2, it remains the same without any major changes. The component can be accessed through the Event Server's RESTARL to comply with the original component architecture, where the Event Server acts as the main communication interface with the rest of the MOCA sub-components.



Figure 19. MOCA architecture

1.23 Interaction with other NEMO components

The interactions of the MOCA component with other NEMO components have remained the same as presented in deliverable D4.2. However, for consistency, in Figure 20 an updated version of the integration diagram is provided including the Smart Contract Deployment Tool subcomponent.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	42 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final





Figure 20. MOCA integration diagram

4.3.4 Results

In deliverable D4.2 the initial set of MOCA's features and capabilities were described in this document the vastly expanded accounting capabilities of the component are presented Initially, MOCA utilized a set of smart contracts to perform the necessary calculations of the rewards and costs in the NEMO continuum, namely (a) NemoTokenEstimation, (b) NemoFunds, (c) InfrastructureOwnerModel and (d) ServiceProviderModel. In addition to these smart contracts, MOCA was enriched with supplementary calculation methods that enhanced its versatility, such as:

- Calculate the costs of network usage of a workload
- Calculate the costs of the Machine Learning workloads
- Allow the users to upload a dataset and reward them
- Apply cost deductions for work oads that run on energy efficient clusters

4.3.4.1 Network Resource Monetization

To calculate the network usage costs, a new smart contract, the NetworkUsageModel, Listing 1, has been created to make the necessary cost calculations for a workload. It is worth noting that only the workloads that are exposed through MEMO, with the usage of the ingress attribute, are subject to this cost deduction. As an incentive, when the workloads with the ingress attribute enabled are registered, MOCA rewards them with 5 additional tokens, apart from the initial 5 tokens for the registration. Every workload which is deployed as an external service is given 5GBs of free network usage before they start getting charged. Above the threshold, the workload is charged per GB consumed. The network usage of a workload is calculated periodically by MOCA and it is defined as the sum of the transmitted and received data of the workload.

NetworkUsageModel sol
pragma solidity ^0.8.0;
import "./NemoFunds.sol";
contract NetworkUsageModel {
NemoFunds public nemoFunds;
uint256 constant MAX_FREE_NETWORK_USAGE = 500000000;
uint256 constant CHARGE PER GB = 10000000;
struct NetworkMetrics {
string serviceld;
string clusterId;
uint256 networkUsage; // in GB
}
event NetworkComputeTokens(
string serviceId,
string clusterId,
uint256 networkUsage,

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	results.	Page:	43 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



```
uint256 tokens
  );
  constructor(address nemoFundsAddress) {
    nemoFunds = NemoFunds(_nemoFundsAddress);
  function computeCredits(NetworkMetrics memory _metrics) public {
    require(
       nemoFunds.isCustomerRegistered(_metrics.serviceId),
       "The service is not registered!"
    );
    uint256 _networkUsage = _metrics.networkUsage;
    uint256 _tokens = 0;
    if (_networkUsage > MAX_FREE_NETWORK_USAGE) {
       _tokens =
         (( networkUsage - MAX FREE NETWORK USAGE) * CHARGE PER GB)/
         10 ** 8;
    nemoFunds.makeTransaction(
       _metrics.serviceId,
       metrics.clusterId,
       tokens
    ).
    emit NetworkComputeTokens(
       _metrics.serviceId,
       metrics.clusterId,
       networkUsage,
       tokens
    );
 }
}
```

Listing 1. The NetworkUsag Model smort contract details

4.3.4.2 ML-As-A-Service (MLaaS)

NEMO supports the deployment of Machine Learning models, which usually require GPU power to perform efficiently processing intensive processes such as the training of a large model. In this context, it was vital to offer a smart contract dedicated to the use of GPU resources, MLModel, Listing 2. The NEMO continuum offers GPUs with time slicing enabled, which allows for multiple jobs to be executed in parallel simultaneously. Therefore to accurately calculate the usage cost for a Machine Learning workload, it was important to consider, except for the GPU percentage that has utilized, also the time it had reserved that particular resource. Another crucial aspect that was used for the MLaaS was the deployment cost of such infrastructure based on socioeconomic factors.

MachineLearningModel.sol
pragma solidity ^0.8.0;
import "./NemeTokenEstimationSetupContract.sol";
import "/NemoFunds.sol";
contract MachineLearningModel {
NemoTokenEstimationSetupContract public nemoTokenEstimationSetup;
Nemo unds public nemoFunds;
address public owner;
uint256 public gpuRate;
// Number of seconds per hour
uint256 constant SECONDS_PER_HOUR = 360000000000;
struct MachineLearningMetrics {
string serviceld;
string clusterId;

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	44 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



```
string region;
  uint256 gpuUtil;
  uint256 allocatedTime;
  uint256 usageFraction;
}
event MachineLearningComputeTokens(
  string serviceId,
  string clusterId,
  string region,
  uint256 gpuUtil,
                                                                                 to the
  uint256 tokens
);
constructor(
  address _nemoTokenEstimationSetupContractAddress,
  address _nemoFundsAddress
){
  nemoTokenEstimationSetup = NemoTokenEstimationSetupContract(
     _nemoTokenEstimationSetupContractAddress
  );
  nemoFunds = NemoFunds(_nemoFundsAddress);
}
modifier checkRegionData(string memory region) {
  require(
     nemoTokenEstimationSetup.isRegionSet(region),
     "Data for region must be set before calling this fun
                                                     ion."
  );
}
function computeCredits(MachineLearningMetrics memory _metrics) public {
  (uint256 _regionalGpuCosts) = nemoTokenEstimationSetup.getRegionGpuInfo(
     _metrics.region
  );
  string memory serviceId =
                             metrics.serviceId;
  string memory _clusterId
                             metrics.clusterId;
  uint256 _tokens
                  = 0
  _tokens =
     ((_metrics.allocatedTime * _metrics.usageFraction) /
       SECONDS_PER_HOUR) +
      regionalGpuCosts;
     poFunds.makeTransaction(_serviceId, _clusterId, _tokens);
  emi MachineLearningComputeTokens(
     _metrics.serviceId,
     metrics.clusterId,
     _metrics.region,
     _metrics.gpuUtil,
     tokens
  );
}
```

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	45 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Listing 2. The MachineLearningModel smart contract details

4.3.4.3 Data Monetization

One of the rapidly increasing aspects of the AI/ML domain is dataset sharing. In this context, MOCA provides a monetization mechanism of data for NEMO users. MOCA's Event Server allows the upload of datasets for 3rd party users to utilize them on their own applications. In return, the dataset provider is rewarded with 10 tokens. The datasets are stored in an internal S3 bucket that is accessible through MOCA's API. In Figure 21 the endpoint related to the dataset's uploading is presented along with the payload fields required, while on Listing 3 the details of the related smart contract are shown

POST /da	ataset/upload				dataset_upload_create 🔨
Parameters					Cancel
ame	Description				
уре	The type of the dataset				
(tormData)	poisoned-tomatoes-dataset				
axLength: 46 inLength: 1					
lescription	A description for the dataset				
tring formData)	Images fo poisoned and healthy tomato				
axLength: 80					
netadata	The metadata of the dataset				
tring formData)	natoes, poisoned, classification, images				
axLength: 55					
inLength: 1					
iataset ile	The uploaded dataset				
formData)	Browse poisoned-tomdataset.tgz				
	I	Execu Figure 21. Example	te of deploying dat	taset	
taMode	1.50	Execu Figure 21. Example	te of deploying dat	aset	
taMode	1.sol	Execu Figure 21. Example	te of deploying dat	taset	
itaMode agma so	1.sol	Execu Figure 21. Example	te of deploying dat	taset	
agma so ntract Da	I . sol lidity ^0.8.0; ataModel {	Execu Figure 21. Example	te of deploying dat	taset	
taMode agma so ntract Da struct Da	I . sol lidity ^0.8.0; ataModel { ataInfo {	Execu Figure 21. Example	of deploying dat	aset	
ataMode agma so ntract Da struct Da string	I . sol lidity ^0.8.0; ataModel { ataInfo { dataId; dataId;	Execu Figure 21. Example	of deploying dat	aset	
agma so agma so ntract Da struct Da string string string	I . sol lidity ^0.8.0; ataModel { ataInfo { dataId; dataType; description;	Execut	of deploying dat	aset	
ataMode agma so struct Da struct Da string string string string	I . sol lidity ^0.8.0; ataModel { ataInfo { dataId; dataType; description; metadata;	Execut	of deploying dat	aset	
taMode agma so struct Da struct Da string string string string string	I.sol lidity ^0.8.0; ataModel { ataInfo { dataId; dataType; description; metadata; endpoint;	Execution	of deploying dat	aset	
taMode agma so struct Da string string string string string string	I . sol lidity ^0.8.0; ataModel { ataInfo { dataId; dataType; description; metadata; endpoint;	Execution Figure 21. Example	of deploying dat	taset	
taMode agma so struct Da string string string string string string event Da	I lidity ^0.8.0; ataModel { ataInfo { dataId; dataId; dataType; description; metadata; endpoint; ataComputeTokens(string dat	Execu Figure 21. Example	of deploying dat		
taMode agma so atract Da struct Da string string string string string string string string function uint25	I 1.sol lidity ^0.8.0; ataModel { ataInfo { dataId; dataType; description; metadata; endpoint; ataComputeTokens(string dat computeCredits(DataInfo me i6_tokens = 0;	Executive State of the second	of deploying dat	taset	
taMode agma so atract Da struct Da string	I . sol lidity ^0.8.0; ataModel { ataInfo { dataId; dataType; description; metadata; endpoint; ataComputeTokens(string dat computeCredits(DataInfo me 66_tokens = 0; ns = 1000000000;	Executive State of the second	of deploying dat	taset	
ataMode agma so struct Da string	I 1.sol lidity ^0.8.0; ataModel { ataInfo { dataId; dataType; description; metadata; endpoint; ataComputeTokens(string dat computeCredits(DataInfo me 66_tokens = 0; ns = 1000000000; DataComputeTokens(_info.da	Executive Statement of Statemen	of deploying dat		
ataMode agma so struct Da string stri	I 1.sol lidity ^0.8.0; ataModel { ataInfo { dataId; dataType; description; metadata; endpoint; ataComputeTokens(string dat computeCredits(DataInfo me 66_tokens = 0; ns = 1000000000; DataComputeTokens(_info.da	Executive Strain	of deploying dat		

Listing 3. The DataModel smart contract details

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	46 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final

}



4.3.4.4 Green Energy rewards

In this version of MOCA, we have also incorporated a rewarding mechanism for the workloads that run on "green" clusters. A workload that runs on a green cluster will receive a reduction to their resource usage costs, depending on the percentage of green energy supported by the infrastructure. This method aims to give incentive to the workload owners to prefer green deployment clusters, so that:

- 1. They reduce their workload costs and
- 2. Indirectly, they help with making the NEMO continuum more energy efficient.

The *ServiceProviderModel* smart contract was modified appropriately to consider the green energy percentage of the deployment platform. Annex 3 – ServiceProviderModel provides the details of the contract.

4.3.4.5 Smart Contracts Deployment Tool

Another major addition to the MOCA component is the **Smart Contracts Deployment Fool**. This subcomponent allows the users to upload their own smart contracts in the private Quorun blockchain used for the contract deployment. In this way, the users can interact easily with the blockchain, customize and test MOCA with their own smart contracts, implementing different approaches in the accounting logic. Figure 22 demonstrates a simple example of the body of the request to upload the contract. The user only needs to provide their smart contract.

POST /contract/upload/ This will call the deployer class and deploy the contract just uploaded.	<pre>post_contract_upload_</pre>
It will return the deployed contract address	
Parameters	Cancel
No parameters	
Request body	multipart/form-data v
file string(\$binary) Browse Handler.sol Send empty value Execute	
Solution W oure 22 Example of contract deployment	

Figure 22. Example of contract deployme

4.3.5 MOCA Business Models

The MOCA component, as mentioned earlier in this document, is not a business model by itself but rather a foundational enabler for a variety of business models within the NEMO ecosystem. It establishes a transparent, automated, and fair marketplace for resources and services across the AIoT-Edge-Cloud continuum. The core of MOCA's design facilitates a "credits" or "tokens" based economy, where contributions are rewarded and consumption is billed, all managed via smart contracts on a distributed ledger.

The primary business models MOCA enables can be viewed from the perspective of different stakeholders:

4.3.5.1 Resource-as-a-Service (RaaS) for Providers:

• **Computing Resource Monetization**: Infrastructure owners (as per *InfrastructureOwnerModel*) can offer their computing resources (CPU, RAM) and get rewarded based on usage by service providers (as per *ServiceProviderModel*). MOCA's smart contracts ensure accurate accounting and reward distribution.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	47 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



- Network Resource Monetization: Providers offering network access can monetize this, especially for workloads requiring external exposure (as per *NetworkUsageModel*). The model incentivizes initial free usage, then charges per GB, creating a clear value proposition.
- **Specialized Hardware Monetization (e.g., GPUs for MLaaS):** Owners of specialized hardware like GPUs can offer these for Machine Learning (ML) workloads (as per *MachineLearningModel*). MOCA allows for nuanced billing based on GPU utilization, allocated time, and even regional cost differences, enabling a viable ML-as-a-Service (MLaaS) model.
- **Data Monetization:** Entities possessing valuable datasets can make them available on the NEMO platform (as per *DataModel*). They are rewarded with tokens when their data is accessed or utilized, creating a marketplace for data itself.
- **Incentivizing Green Energy**: Resource providers utilizing green energy sources can ofter their infrastructure at potentially premium rates or benefit from NEMO's reward mechanisms (as seen in the *ServiceProviderModel* modifications for green energy rewards), encouraging sustainable practices and creating a niche market.

4.3.5.2 Pay-As-You-Go (PAYG) & On-Demand Services for Consumers

- Flexible Resource Consumption: Service developers and end-users can consume various resources (compute, network, ML processing, data) on-demand without significant upfront investment. They are billed transparently through MOCA based on their actual usage.
- Access to Specialized Services: Consumers can readily access and pay for specialized services like MLaaS or specific datasets, fostering innovation and reducing barriers to entry for developing complex applications.
- **Cost Optimization:** Incentives for using green clusters, along with tiered pricing models (e.g., free initial network usage), enable consumers to reduce and optimize their operational costs.

4.3.5.3 Custom Business Logic via Spart Contract Deployment:

- Extensible Monetization Schemes The Smart Contract Deployment Tool (Section 4.3.4.5) is a crucial enabler for novel business models. It allows NEMO users (providers or even third-party integrators) to deploy their own smart contracts with custom accounting and monetization logic. This opens the door for:
 - Subscription based models.
 - Revenue-sharing agreements.
 - Usage tier based pricing is not natively covered by default MOCA contracts.
 - Loyalty programs or bundled service offerings.
- Fostering a Brorder Ecosystem: This flexibility allows the NEMO marketplace to adapt and evolve, accommodating new types of resources or service delivery models as they emerge, driver by the community or specific enterprise needs.

In essence, MOCA aims to create a vibrant, self-sustaining economic ecosystem. It empowers resource providers to generate value from their assets and service consumers to access these assets efficiently and transparently. The extensibility offered by custom smart contract deployment ensures that the business models operating within NEMO can evolve and diversify over time. The distributed ledger technology underpinnings provide the trust and auditability necessary for such an ecosystem to thrive.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	48 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



4.4 Intent-based SDK/API

4.4.1 Overview

Intent-based SDK/API (IBA) is the interface of NEMO to third party interactions, with the finetuning of this component being continuous to incorporate new features as the project progresses. To this end, several activities towards the finalization of the Intent-based SDK/API were carried out since the delivery of deliverable D4.2. These activities can be divided into two categories, a) the integration related activities, and b) core optimization activities that enhance the existing functionalities of the tool.

4.4.2 Architecture

The Intent Based API was developed in a microservices oriented architecture aiming to a highavailability and highly scalable system. The current architecture, Figure 23, demonstrates the IBA subcomponents as well as their interaction with other NEMO components. This allows a fast and easy integration of new components without major modifications in the entire IBA system, just only in the specific subcomponent that interacts with the newly introduced NEMO component. An example of this architecture will be presented in the subsequent subsections with the integration of mNCC to IBA.



Figure 23. Final Architecture diagram of Intent-Based API

4.4.3 Interaction with other NEMO components

The activities related to the Intent-based SDK/API in the 2nd integration phase of the project are mainly focused on the integration between IBA and other NEMO components. As described in the Sub-section 4.4.2, Intent-based SDK/API is located at the front-frontier of the NEMO framework, meaning that it is the main input of the information used by the NEMO components. In this context, IBA is participating

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	49 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



in many integration pipelines which are described in Section 5. The related work carried out involved the integration with Meta-Orchestrator's components, such as the IBMC and the deployment/undeployment components. In this integration pipeline, new message exchanging queues were created, while the existing ones were updated to house the needs of the new features presented by the Meta-Orchestrator as well as to provide the maximum optimization for the deployment of workloads. In addition, several integration actions were performed to incorporate IBA with NEMO Access Control component. This was of paramount importance as it allows the use of IBA from external parties that are located outside of the NEMO ecosystem, e.g., a service in an external cluster. In terms of workload management, several modifications were made regarding the IBA-CMDT interactions. More specifically, the work carried out was focused on the integration of Linkerd¹³ into the features of BA to be able to handle all the states of a workload, such as creation, deployment, running status and termination by adding the necessary annotations to the generated manifest file, Listing 4 Finally, is the integration of IBA with the networking components (mNCC & L2S-M) of NEMO. For this integration pipeline, severe modifications were required at the IBA. A new endpoint was created to handle the new Intent request along with the related Intent, Annex 1, and its validator. As for the communication pipelines for inter component communication, several RabbitMQ channels were preated to allow the communication and state management of the network slices from the IBA-Listing 5.



nemo.rabbitmq INFO 2025-03-27 14:53:59,245 rabbitmq [x] Sent to mncc, routing_key=mncc.ibs, message: {"workload": "61epb162.1d52-4e4d-8b55-275187a780ba", "pod": {"labels": {"l2sm": true, "l2sm/app": "[parameter]"}, "annotations": {"l2sm/networks": "spain-network"}, "env": {"DNS_NAME": "[parameter].spain-network.inter.l2sm"}}}

isting 5. RabbitMQ communication channel with mNCC and sample payload

On the other hand, apart from the integration related development of IBA, several developments for enhancing both the User and the intra-service interfaces were done. Firstly, a lifecycle event monitoring subcomponent was created to better handle the state of the workloads, Annex 2. This was a crucial feature as it allows the end-user to have knowledge of each state of a workload that is passed through the Graphic User Interface. Then, the user can understand if his workload has been initiated, running or terminated. In case an error occurs, the user can retrieve the information related to that failure in a user-friendly way. Another enhancement that took place in the 2nd integration phase was related to the MOCA component. As the service's needs changed some of the MOCA's endpoints got deprecated and later

¹³https://linkerd.io/

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	50 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



removed from the IBA-MOCA integration pipeline as this was reported in deliverable D4.2. In addition, as the project progresses and based on the iterative Quality Assurance process employed in NEMO as well as the integration process a minor number of bugs along with service improvements were identified. To tackle these newly identified needs, IBA enhanced its existing toolset by presenting new endpoints for housing the needs of a modern microservice. Moreover, the need for specific information in Intents needed for the evaluation and deployment of a workload, lead to the creation of use case specific Intents, such as the Machine Learning Intent, Listing 6, that forces the workload to be deployed in a GPU compatible cluster node.



Listing 6. MachineLearning Intent sample

4.4.4 Conclusion

The Intent Based API was designed to be the interface between the Graphical User Interface or the 3rd party User and the NEMO components. As mentioned in the previous subsections, IBA plays a pivotal role within the NEMO architecture as it orchestrates the communication channels and is the initial point for many crucial NEMO workload pipelines. Since the delivery of D4.2, IBA was enhanced with functionalities introduced by other components as well as internal optimizations that guarantee the efficient operation of NEMO. The workflow pipelines where the IBA participated in are described in Section 5 of this document.



Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	51 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



5 NEMO scenario-driven verification & results

In this section the activities and tests performed towards a complete and integrated NEMO infrastructure are presented. Apart from the WP4 tools' integration activities, as presented in D4.2, for completeness in this document also the integration pipelines related to WP2 and WP3 tools are presented. This section aims to provide in a complete and coherent way the information regarding the tools' interactions within NEMO. As the NEMO framework consists of many tools one of the first critical issues was the testing procedure. To solve this issue, a categorization of the tools took place based on sub-system functionalities was selected. Through this system breakdown dedicated pipelines created focusing on small and dedicated functions of the System.

Each scenario, as outlined in the subsequent subsections, focuses on a distinct operational aspect of the NEMO meta-OS. The results presented herein provide empirical evidence of the platforms capabilities and its readiness to support the diverse use cases envisioned within the NEMO project, highlighting both successful integrations and areas for further refinement. The scenarios cover critical functionalities including cluster registration, workload registration and provisioning, workload scheduling and orchestration, workload lifecycle management, and integration activities stempting from WP3, such as the Secure Execution Environment.

These dedicated integration pipelines provided several benefits in the integration procedure as the integration, testing and optimization of the components based on specific functionalities made the entire process easier and more robust, as well as aligning with the best practices of Software development principles.

5.1 NEMO Cluster registration – resource provisioning

This section details the NEMO cluster registration scenario related activities undertaken. The integration results as reflected in the cluster registration sequence diagram described in Figure 24 are presented below.



Figure 24. Cluster registration sequence diagram

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	52 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



5.1.1 Verification scenario

Test 1: NEMO Clu	uster registration
Objective	To verify the cluster registration process in NEMO that facilitates the resource provisioning triggered by the NEMO partner (resource owner)
Components	• LCM
	Intent-based API
	• MO
	• MOCA
	• RabbitMQ
Features to be tested	The feature that this scenario aims to test are the cluster registration process which is initiated by the NEMO partner (cluster provider) through the LCM UI & intent based API. Then, the newly registered cluster is added into the NEMO meta-OS ecosystem by the MO. The results (status) of this process are then visualized to the user.
Test setup	All the participating components were deployed in the NEMO meta-OS cloud/edge infrastructure at OneLab (dev cluster 1).
Steps	 Cluster registration through the LCM UI Cluster registration message communication to MO Cluster addition process by MO Cluster status provisioning to RabbitMQ Cluster status update visualization in LCM UI

5.1.2 Results

This section documents the process that is described in the scenario above step by step.

5.1.2.1 Cluster registration through the LCM UI

To register a cluster through the LCM UI, first, we go through the LCM UI home page to the "Resource Provisioning" tab, Figure 25

⊞ » Hosts |



Figure 25. LCM home page

In the "Resource Provisioning" page, we have an overview of some basic details for the clusters, like their ID, name, resources (CPU, RAM, disk, GPU), their deployment status and the endpoint they are available at, as it is shown in Figure 26.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	53 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



			🖵 Clusters Table	e					
Search									
Create Cluster									
ID 🕂	Name	Status	Endpoint	CPUs	Storage	Memory	Vram	Actions	
933189e	dev-onelab	ок	https://api.main.nemo.onelab.eu:6443	68	1620	172	0		
8d63cc0	test-cluster	ОК	https://api.main.nemo.onelab.eu:6443	10	300	120	0		
67c21fc	pro-onelab	ОК	https://api.prod.nemo.onelab.eu:6443	32	1350	62	14	11 I	
60a2ee7	k3s-onelab	ОК	https://api.s2.nemo.onelab.eu:6443	16	1080	32	0	S 🕯	
1e58cae	nemo-smart-farming	ОК	https://83.235.169.221:35443	6	1000	8	0		
179592b	staging-onelab	ОК	https://api.staging1.nemo.onelab.eu:6443	32	1080	62	0		
142-224	energy	ОК	https://comsensus.eu:13387	5	5	5	0	1	

Figure 26. Clusters overview page

When a user registers a cluster, they are prompted to filling the form shown in Figure 27. The form requires of the user to upload the configuration file of their cluster and fill in information like the name of the cluster, its endpoint, the resources of the cluster, the availability percentage, its green energy percentage and possible costs they want to apply to the workload that will be deployed on their premises.

	└─ Re	egister Cluster	
Cluster Config File			
🛛 kube-config			
Cluster Name		Managed API	
test-cluster-2		https://api.mai	n.nemo.onelab.eu:6442
CPUs #	Memory (GB)	Storage (GB)	V-RAM (GB)
12	200	30	0
Availability %	Green Energy	%	Cost
99%	- 60%	•	low_cost
CPU Base Rate (in millitoken	is)	Memory Base Rate	e (in millitokens)
12		10	

Figure 27. Cluster registration form

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						54 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



5.1.2.2 Cluster registration message communication to MO

Once the submission form is completed correctly, LCM sends to NEMO the registration request (the request is proxied through the Intent API). MOCA receives the request informs the Meta Orchestrator of request, in order to join the cluster with the NEMO continuum. Until the join process is completed, the cluster is saved with state "Pending", Figure 28.

ervices Overview										
	Clusters Table									
Search										
Create Cluster	Name	Status	Endpoint	CPUs	Storage	Memory	Vram	Actions		
933189e	dev-onelab	ОК	https://api.main.nemo.onelab.eu:6443	68	1620	172	0	s 💼		
8d63cc0	test-cluster	ОК	https://api.main.nemo.onelab.eu:6443	10	300	120	0			
8909fd2	test-cluster-2	PENDING	https://api.main.nemo.onelab.eu:6442	12	30	200	0	si 💼		
67c21fc	pro-onelab	ОК	https://api.prod.nemo.onelab.eu:6443	32	1350	62	14			
60a2ee7	k3s-onelab	ОК	https://api.s2.nemo.onelab.eu:6443	16	1080	32	0			
1e58cae	nemo-smart-farming	ОК	https://83.235.169.221:35443	6	1000	8	0			
179592b	staging-onelab	ОК	https://api.staging1.nemo.onelab.eu:6443	32	1080	62	0			
142e334	energy	ОК	https://comsensus.eu:13387	5	5	5	0	ii 1		
						Items per	page: 10	▼ 1 - 8 of 8 <		

Figure 28. Cluster pending status

5.1.2.3 Cluster addition process by M

Figure 29 shows the message from MOCA the Meta Orchestrator received. The message includes all the information the user filled in in the registration form.



ter status provisioning to RabbitMQ

Once the Meta Orchestrator has joined the cluster, it sends through RabbitMQ a message with the status of the action. Since the process was completed successfully, the status returned is "ok", Figure 30.

→ Rubecti Logs Roca-cluster-parser-sobodrbd4/-KXT2] - n nemo-svc
[12/May/2025 18:57:38] - [nemo.rabbitmq:124] - [INFO] [*] Waiting for messages. To exit press CTRL+C
[20/May/2025 14:38:49] - [nemo.rabbitmq:63] - [INFO] [x] Received {'action': 'join', 'managed_api': 'https://api.main.nemo.onelab.eu:6443', 'cluster_name': 'test-cluster-2', 'operation_id': '2603d351-d2cc-46a0-9f
44-5b4147d88598', 'timestamp': '2025-05-20T16:38:49+02:00', 'status': 'ok'}
[20/May/2025 14:38:51] - [app.utils.utils:282] - [INFO] http://127.0.0.1:8080/ipfs/QmTtpAaYdtWEsSUweDtRyCffW6uTdaMvuaHYBSNrPZKZc
[20/May/2025 14:38:56] - [nemo.rabbitmg:81] - [INFO] This is the CID: QmTtpAaYdtlwEs5UweDtRyCffW6UTdaMvuaHYBSNrPZKZC
[20/Nay/2025 14:39:02] - [nemo.rabbitmg:94] - [INFO] The cluster was registered with transaction hash 0x9c52620b258fbbe0a6130d60666b363bdccbc64ce7981053757f531eb1e5474
[20/May/2025 14:38:56] - [nemo.rabbitmq:81] - [INF0] This is the CID: QmTupAaYditMESSUPCTFWouTdaMwaAMBSMr72X2C [20/May/2025 14:39:02] - [nemo.rabbitmq:94] - [INF0] The cluster was registered with transaction hash 0x9c52020b250fbbe0a6130d6d6d6ib363bdccbc64ce798105377f531eb1e5474

Figure 30. Meta-Orchestrator response

After receiving this message, MOCA continues to store the provided cluster configuration file and register the cluster to the blockchain. Figure 30 provides some logs that show that both actions were completed successfully. The file was stored in IPFS and can be retrieved, if necessary, with the produced

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						55 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



CID, a unique identifier of that file. IPFS here is used as an extra storage layer which guarantees that any tampering attempts to the cluster configuration file will be easily detectable, as any changes to the content will produce a new CID. Since the configuration file holds sensitive information, it is important to ensure the integrity of the data. After this is completed successfully, the cluster is registered in the blockchain. Figure 31 shows the details of the transaction.

[26/May/223 34:39:02] [app.event_receivers:11] []NF0] ['args: [['customerid': '2603d351-dzcc-46a8-9f44-5b4147d88598', 'customeritype': 'infrastructure', 'balance': 10868080800, 'neno8alanceActionId': 2]], 'event', 'customerige: 'infrastructure', 'balance': 'infrastructure', 'balace': 'infrastructure', 'balance': 'infrastructure',

Figure 31. Register cluster to blockchain

5.1.2.5 Cluster status update visualization in LCM UI

At the end of the process, MOCA updates the status of the cluster to "OK". This is visible through the LCM UI, as shown in Figure 32.

ervices Overview										
Clusters Table										
Orangh	Saarah									
Search										
Create Cluster										
ID 🕁	Name	Status	Endpoint	CPUs	Storage	Memory	Vram	Actions		
933189e	dev-onelab	ОК	https://api.main.nemo.onelab.eu:6443	68	1620	172	0			
8d63cc0	test-cluster	ОК	https://api.main.nemo.onelab.eu:6443	10	300	120	0			
67c21fc	pro-onelab	ОК	https://api.prod.nemo.onelab.eu:6443	32	1350	62	14			
60a2ee7	k3s-onelab	ОК	https://api.s2.nemo.onelab.eu:6443	16	1080	32	0			
2603d35	test-cluster-2	ОК	https://api.main.nemo.onelab.eu:6443	12	200	30	0			
1e58cae	nemo-smart-farming	ОК	https://83.235.169.221:35443	6	1000	8	0		_	
179592b	staging-onelab	ОК	https://api.staging1.nemo.onelab.eu:6443	32	1080	62	0			
142e334	energy	ОК	https://comsensus.eu:13387	5	5	5	0			
						Items per	r page: 10	₩ 1 - 8 of 8	< >	

Figure 32. Updated cluster status

5.1.2.6 Results from mela-Orchestrator

In step 3, "Cluster addition process by MO," MO receives the cluster joining request from the RabbitMQ queue. After several communications between MO's subcomponents, MO API registers the cluster if the request is well formed and there is no error. Listing 7 presents an example of a joining request payload that the MO is expecting to receive. If there is no error, the MO will send an ok message, Listing 8. Otherwise, if an issue has occurred, either syntactic error or connectivity related, with the request, the MO returns an error, Listing 9. The error presented in Listing 9, error 409 is related to a joining request for an existing cluster.

{	
	"availability": "99.9%",
	"cluster_name": "test",
	"cost": "low_cost",
	"cpu_base_rate": 10,
	"cpus": 10,

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						56 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



"green_energy": "20%", "id": "28aa341f-246d-4890-886b-79529d8e8b7e", "managed_api": "https://api.s2.nemo.onelab.eu:6443", "memory": 200, "memory_base_rate": 10, "storage": 300, "timestamp": "2025-02-04T10:37:55.640728Z", "vram": 1024 }





"action": "join",
"cluster_name": "test_cluster",
"error": "request failed: http request failed with status 409",
"id": "4981421-test-join",
"status": "error"

Listing 9. Error joining request

Finally, after the joining request succeeds, each cluster's metrics persist in the MO database. The metrics used for the Workload Placement were described in deliverable D3.3 [20], Table 8.

Tabl	e 8.	Cluster	Metrics
1 401	U U.	Cluster	methos

Field	Туре	Title	Description			
claster_hame	string	Cluster name	The name of the Cluster that will be deployed.			
cpus 🗡	number	CPUs	The number of CPUs of the Cluster.			
memory	number	Memory	The RAM of the Cluster in MB.			
storage	number	Storage	The disk storage of the Cluster in MB.			
availability	string	Availability	The percentage of time that the cluster is up (99.9%, 99%, 90%).			

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						57 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Field	Туре	Title	Description
green_energy	string	Green energy	The percentage of RES powering the cluster (0%, 20%, 40%, 60%, 80%, 100%).
cost	string	Cost	The cost type of a cluster (low cost, high performance). Enum
cpu_base_rate	number	CPU base rate	The CPU cost of the cluster by the CPU capacity of the cluster (in milliseconds).
memory_base_rate	number	Memory base rate	The memory cost of the cluster by the memory capacity of the cluster (in MBs).
vram	number	Video Random Access Memory	The cluster VRAM available in MB

5.2 NEMO workload registration & provisioning

Once the cluster registration has been performed within the NEMO infrastructure, a NEMO user can proceed to register, deploy and provision the corresponding workflows in the infrastructure. Figure 33, Figure 34 and Figure 35 depict the whole process necessary for the proper deployment of a workload within NEMO.

Since both the workload provisioning, Figure 33, and deployment, Figure 34, were already described in deliverable D4.2. This section of this document focuses on the interactions that occur during the deployment process of a NEMO workload, especially detailing how NEMO workloads take advantage of the mNCC flexibility to support isolated communication between its functionalities. Supplementary to the workload provisioning, the workload deployment process showed in the Figure 34 is also a key part for the completion of this experiment.

Following the initial stages of the worklead deployment process illustrated in Figure 33 and Figure 34, the workflow continues with the creation of the network elements necessary to establish isolated communication channels between NEMO Workloads deployed over the project's infrastructure. Figure 35 shows the creation process of a virtual network by the mNCC NEMO component.



Figure 33. Workload provisioning workflow in NEMO

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	58 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final





NEMO Workload Deployment

Figure 34. NEMO workload deployment workflow

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	59 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final







The virtual networks created by the mNCC, enable isolated communications between two (or more) pods located in one, or multiple, Kubernetes (K8s) clusters in NEMO. This step must be performed before the deployment of the pods themselves within the clusters used for the deployments, since each virtual network is defined (and created) as a K8s resources that each pod must consume at its instantiation stage.

In this regard, when a new virtual network is created through the mNCC component, the pods that consume this K8s resource (i.e., pods attachto the virtual network) will be able to communicate between each other as if they were located in the same Local Area Network (LAN), regardless of their physical location. More details about these virtual network resources can be seen in Deliverables D2.2 and D2.3.

It is important to address that, although the original workflow included the creation of multiple network slices where various NEMO administrative domains could be split based on the needs of each platform at network level, dynamically modifying the network topology needed for each use case, the workflow presented in this integration document will feature a single network slice. The purpose of this decision is two-fold: on the one hand, due to the relatively low complexity of the current NEMO infrastructure, splitting the network into multiple slices will not provide significant improvements for the isolation of NEMO deployments (since it will be done using the virtual networks created in the default slice. On the other hand, it simplifies the NEMO workload deployment process, as the user will not need to decide in which slide should the NEMO deployment be located. Nevertheless, this slice creation and modification functionality is enabled within the mNCC to enhance the scalability of NEMO once more clusters are aggregated to the NEMO project infrastructure.

In order to start the deployment of each virtual network in the default slice present in NEMO, the Intentbased API component must create the associated network intents that represent the network interconnecting the pods from a NEMO workload.

In Table 9 an overview of how to construct an intent to request a network connection between different clusters is presented. Later in this section, each value will be explained. In Table 9, a complete example of a L2S-M Intent is described.

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	60 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Key	Description	Value
Intent.id	Identifier for the intent	<string></string>
Intent.userLabel	Label to identify the intent among others posted by the intentAPI	cloud_continuum
Intent.intentExpectations	List of expectations for the network	<intentexpectations></intentexpectations>
intentExpectation.expectationId	Identifier for the expectation	<string,number< th=""></string,number<>
intentExpectation.expectationVerb	Verb defining what type of action to apply in the object.	DELIVER***, ENSURE
intentExpectation.expectationObject	Type of object to be delivered by the network	expectationObject>
expectationObject.objectType	Defines type of network inter	K8S_L2_NETWORK*, K8S_CLUSTER_CONFIG**
expectationObject.objectContexts	Set of attributes required to define a specific object	<objectcontexts></objectcontexts>
objectContext.contextAttribute objectContext.contextCondition objectContext.contextValueRange	String defining the attribute. The value or range is set on contextValueRange. And the condition restricting the value or range in contextCondition	name***, providerName*, domain*, pod_cidr* bearer_token**
intentExpectation.expectationTarget	Objective to fulfil by the expectation object	isolatedConnectivity*
intentExpectaion.expectationContext	Scope of the related expectation.	namespace**
Intent.intentContext	Scope of the entire intent. In the NEMO platform the main context is the related workload	NEMO_WORKLOAD

Table 9. Intent for network connectivity request. L2S-M network request attributes.

The network intentis read from the exchange/queue "nemo.api.workload/network-intent" and processed only by the tBS if the *userLabel* matches a network intent. The is classified and translated as described with more detail in D2.2 and D2.3 deliverables. Additionally, is doing some simple network management in order to fit L2S-M requirements and abstract simple networking behaviours from the higher levels in the architecture.

Once the intent has been processed in the Intent Based System (IBS), this component will generate a gRPC [21] [22] [23] service request towards L2S-M (particularly, the L2S-M MD component). This request must include the following fields:

- Network Network Class-less Inter Domain Routing (CIRD), which will be used to automatically assign IP addresses to pods using the virtual network.
- A Provider endpoint. This provider is a combination of a DNS Server and as SDN Controller which will act as a central point which manages the network.

Document name:	D4.3 A Final ve	D4.3 Advanced NEMO platform & laboratory testing results. Final version					61 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



- List of clusters where the network is going to be created. This list includes:
 - Kubernetes endpoint, including the API Key and Bearer Token of a virtual network user.
 - Namespace to be used inside the cluster.
 - The address pool, contained inside the CIDR for the virtual network. For instance, if the chosen CIDR in a new virtual network is 10.1.0.0/16, each cluster would use a subnet like 10.1.n.0/24. This subdivision is processed and calculated by the IBS.

Once this request is received in L2S-M, it will create a virtual network resource in each specified K8s cluster. This resource will be created at the same namespace as the NEMO workload deployment, since virtual networks are *namespaced* resources that must be present in the same namespace of the pole that will use them. Furthermore, once this resource is created in each cluster, the inter-doman L2S-M component will register this creation and enable the inter-domain connectivity between them as well as generating the DNS domain record that must be added in the pol deployments to enable their connectivity within the virtual network (avoiding the need to manually configure the IP addressing of each pod).

To enable pods to utilize the virtual network resource, L2S-M includes the required configuration fields in the gRPC response sent to the IBS. Specifically, any NEMO deployment intending to use the newly provisioned network must specify the following fields in their deployment descriptors:

- L2S-M annotation label, to clarify that the networking resource is managed by this component.
- L2S-M name of the application label (defined by the deployment owner), which will be used to identify the application with DNS.
- The name of the L2S-M virtual network to which the pods will be attached
- DNS name for each pod (L2S-M provides the DNS domain; the deployment owner is responsible for assigning the specific name to the application)

L2S-M populates these values in the gRPC response to the IBS, effectively providing all the necessary information required for the deployment to leverage the virtual network within the Kubernetes cluster.

The Intent-based API has a RabbitMQ consumer listening to *exchange=mncc*, *(routing_key,queue)=(mncc.ibs, mnce.ibs)* for the values mentioned above. A sample message payload expected by the IBA is presented in Listing 10.



Listing 10. Sample payload received by IBA from the RabbitMQ queue

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	62 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



The Intent-based API then updates the workload manifests with the above additions (pod labels, annotations and environment) while also performing substitution on the <workload-name> matching the name of the workload with the given NEMO Workload ID.

When the manifests are updated, the Intent-based API informs the MO by publishing a RabbitMQ message on the *exchange=nemo.api.workload*, *(routing_key, queue)=(run,run)* with the message payload containing "status" field set to "updating", signaling to the MO to update the NEMO Workload residing in the corresponding cluster.

Test 2: NEMO wo	rkload registration and provisioning
Objective	Verify the NEMO workload registration, deployment and provisioning process
Components	 NEMO Workload Registration LCM Intent-based API RabbitMQ NEMO Workload deployment LCM Intent-based API CMDT RabbitMQ Meta-Orchestrator NEMO Access Control NEMO Workload privisioning Intent-based API Access Control NEMO Network Deployment Intent-based API Access Control NEMO Network Deployment Intent-based API Access Control NEMO Network Deployment Intent-based API Access Control
Features to be tested	 To verify the workload registration, deployment and provisioning process in NEMO, the following features will be tested: Workload Registration Workload Deployment Workload Provisioning Network Deployment The feature that this scenario aims to test are the workload registration process which is initiated by the NEMO consumer (workload provider) through the LCM GUI & Intent-based API. Then, the newly registered workload is requested to be deployed into the NEMO meta-OS ecosystem by the MO. Following the workload provisioning process is triggered which is facilitated by the Intent-based API and the Access Control components. The results including the workload status are visualized to the user. The Intent-based API initiates a network creation intent, which is propagated through RBMQ to mNCC and L2S-M, triggering gRPC service setup. Once the network is ready, updated deployment fields are returned through the same path to the Intent-based API.

5.2.1 Verification scenario

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	63 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



, and the second s	F
The updated work appropriate network	oad is then published and deployed to the managed clusters with annotations by MO.
Test setup The associated comp	ponents are deployed in OneLab facilities
StepsThe steps identified1.NEMO wo a.a.W b.b.Ex c.c.NEMO wo 	in the associated sequence diagrams are listed below: rkload registration orkload registration by the NEMO user through the LCM UI secution of workload validation process in Intent-based API otification of the LCM UI about the status of the workload registration rkload deployment orkload deployment by the NEMO user through the LCM GU secution of workload validation in Intent-based API orkload deployment by the NEMO user through the LCM GU secution of workload validation in Intent-based API orkload ployment by the NEMO user through the LCM GU secution of the deployment request to the LCM GU secution of the deployment request to the MO seployment operation process triggered by MO orkload placement using cluster metrics of MO eployment operation process executer 50 MO ommunication and update of the deployment operation status to the tent-API spose the services using NEMO access control isualization of the updred status to the LCM UI rkload provisioning EMO workload provisioning is triggered by the Intent-based API EMO Access Control Wolkload setup EMO Access Control Wolkload setup EMO Access Control Keycloak plugin functionality rformancuresilience of Kong ¹⁴ Plugin twork Deployment encreatest for network creation by Intent-Base API tent request for network creation by Intent-Base API tent request for metwork creation by Intent-Base API tent request for metwork creation by Intent-Base API tent request for mother keeived mNCC EPC service created by L2SM twork ready notification + deployment fields update from mNCC to entbased API. bush updated workload into RabbitMQ teeive updated workload bt MO odated workload by MO- O deploys pod with network annotations and updates workload into anaged Clusters selected by the HUB

The rest of the subsections present the corresponding results step by step as indicated by the verification scenario.

5.2.2.1 NEMO workload deployment

This section of the present document will be centred around the interactions between the NEMO components and the mNCC component to create network intents that workloads will use in order to communicate through the infrastructure.

¹⁴ [46] [27]

5.22 Results

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	64 of 111
Reference:	D4.3	04.3 Dissemination: PU Version: 1.0					Final



When a new intent arrives with the intent-notify key to the Rabbit-MQ is then read by the IBS. The first step is to filter and detect the type of intent by the label as show in Figure 36. In case the label is not referring to cloud continuum connectivity, the intent is discarded.



Figure 36. Arrival of intent and first filter.

Then, the intent is classified and processed in the IBS as explained in deliverable D2 3. Then, the intent triggers the execution to a call of L2S-M grpc server. This is shown in the last lines of Figure 37, where the grpc executioner is called and starts the exchange of the network request.



Once the intent has been received from the IBS component, L2S-M proceeds to create the virtual network in the relevant clusters, as shown in Figure 39 and Figure 40 and each network has been properly created in the K8s clusters as a custom resource. In consequence, its details can be obtained through the command line interface of the cluster, showcasing how the combination of the IBS and the L2S-M components enable the creation of virtual networks in the NEMO infrastructure (i.e., the mNCC NEMO component).

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	65 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



NAME	AVAILABILITY CONNECTED_PODS AGE
spain-network	1 AVallable 1 150 S kubectl describe l2network snain-network-1 -n fcc6f34d-f33a-4hfe-a66h-764d9bea
d113contex	nemo-s1
Name:	spain-network-1
Namespace:	fcc6f34d-f33a-4bfe-a66b-764d9bedd113
Ladels: Appotations: 1	
API Version:	l2sm.l2sm.k8s.local/v1
Kind:	L2Network
Metadata:	
Finalizers:	lestamp: 2025-04-30113:30:092
l2sm.oper	tor.io/finalizer
Generation:	1
Resource Ve	sion: 63712928
Spec:	4/88826/-1009-4/8/-8102-94/114890034
Config:	10.1.0.0/16
Network CID	: 10.1.0.0/16
Pod Address	Range: 10.1.2.0/24
Dns Groc	ort: 30818
Dns Port:	30053
Domain:	132.227.122.23
Name: Of Post:	default-slice
Sdn Port:	30808
Type:	vnet
Status:	
10.1.2.1:	
Connected P	d Count: 1
Internal Co	nectivity: Available
Last Assign	d IP: 10.1.2.1
Events:	
	Figure 39. Network resource created in S1 cluster
alex@alex-pc	skubectl get l2networks -n 6512a712-ae0f-4839-9703-fcec4fb13828context nemo-
s2	
NAME	AVAILABILITY CONNECTED_PODS AGE
alex@alex-pc	-\$ kubectl describe l2networks -n 6512a712-ae0f-4839-9703-fcec4fb13828 spain-net
work-1con	ext nemo-s2
Name:	spain-network-1
Namespace: Labels:	<pre>chone></pre>
Annotations:	<none></none>
API Version:	l2sm.l2sm.k8s.local/v1
Kind: Metadata:	LZNETWORK
Creation T	mestamp: 2025-04-30T13:28:06Z
Finalizers	
l2sm.oper	ator.io/finalizer
Resource V	rsion: 54706992
UID:	35a6a3ab-8ea3-4727-9a10-84cbc92e687a
Spec:	
Config: Network CT	10.1.0.0/16 R: 10.1.0.0/16
Pod Addres	Range: 10.1.1.0/24
Provider:	
Dns Grpc	Port: 30818
Domain:	132.227.122.23
bond cirr.	default-slice
Name:	6633
Name: Of Port:	
Name: Of Port: Sdn Port:	30808
Name: Of Port: Sdn Port: Type: Status:	30808 vnet
Name: Of Port: Sdn Port: Type: Status: Assigned I	30808 vnet Ps:
Name: Of Port: Sdn Port: Type: Status: Assigned I 10.1.1.1	30808 vnet Ps:
Name: Of Port: Sdn Port: Type: Status: Assigned I 10.1.1.1 10.1.1.4	30808 vnet Ps:
Name: Of Port: Sdn Port Type: Status: Assigned I 10.1.1.4 Connected f Internal Co	30808 vnet Ps: od Count: 2 nnectivity: Available
Name: Of Port: Sdn Port Type: Status: Assigned I 10.1.1.4 Connected f Internal Co Last Assign	30808 vnet Ps: od Count: 2 nnectivity: Available ed IP: 10.1.1.4
Name: Of Port: Sdn Port Type: Status: Assigned I 10.1.1.1 10.1.1.4 Connected f Internal Co Last Assign Provider Co	30808 vnet Ps: od Count: 2 nnectivity: Available ed IP: 10.1.1.4 nnectivity: Available

Figure 40. Network resource created in S2 cluster

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	66 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Since the L2S-M annotations reported by the mNCC were added to each pod during the deployment process, each pod consequently deployed in each cluster attached to its corresponding network. These details can be found both in the description of the virtual network (i.e., Figure 39 and Figure 40) as well as the annotations present in each pod within the K8s cluster, as depicted in Figure 41 and Figure 42.

<pre>alex@alex-pc:-\$ kubectl get pods -n fcc6f34d- NAME echo-server-integration-500-7ccbd97fc-2ctrws echo-server-integration-500-test-connection alex@alex-pc:-\$ kubectl describe pod echo-ser f33a-4bfe-a66b-764d9bedd113context nemo-s1 l2sm=true l2sm/app=echo-server-integr l2sm/networks: spain-networ DNS NAME: echo-server-integration-500.</pre>	f33a-4bfe READY 2/2 1/2 ver-integ grep 1 ation-500 k-1 spain-net	e-a66b-764d9 STATUS Running NotReady gration-500- .2sm) :work-1.inte	vbedd113c RESTARTS 0 0 7ccbd97fc-2 er.l2sm	ontext nemo-s1 AGE 15d 15d ctrws -n fcc6f34d-
Figure 41. Pod deployed	in S1 with	n L2S-M anno	otations	
<pre>alex@alex-pc:~\$ kubectl get pods -n 6512a712-</pre>	ae0f-4839	9-9703-fcec	4fb13828o	context nemo-s2
NAME	READY	STATUS	RESTARTS	AGE
echo-server-integration-501-89b7c85b6-ggsz9	2/2	Running	Θ	15d
echo-server-integration-501-89b7c85b6-s6xck	2/2	Running	Θ	15d
echo-server-integration-501-test-connection	1/2	NotReady	Θ	15d
alex@alex-pc:~\$ kubectl describe pod echo-ser e0f-4839-9703-fcec4fb13828context nemo-s2 l2sm=true	ver-integ grep l:	gration-501 2sm	-89b7c85b6- <u>c</u>	ggsz9 -n 6512a712-a
<pre>l2sm/app=echo-server-integr l2sm/networks: spain-networ DNS_NAME:</pre>	ation-50: k-1	1 twork_1 int	ar 12rm	
alox@alox_act_\$ kubact] describe and ache sos	vor inte	ration 501	90b7c95b6	Exck p 65122712 2
a c a		2 cm	-050700500-3	
	l dieb r	2.519		
	ation 50	1		
12cm/potworks: spain potwork		L		
DNS NAME: echo-server-interration 501	spain-poi	twork 1 int	ar 12cm	
DRS_NAMEecho-server-threed attom-sol.	spach-ne			

Figure 42 Bod deployed in S2 with L2S-M annotations

It is noteworthy that these details provide the DNS name that pods can use to communicate between each other in the newty created virtual networks. Hence, each pod does not need to know the IP address of the other component, but only its DNS entry, simplifying the connectivity setup between each other with a fully compatible solution with the standard CoreDNS present in standard K8s clusters. With this name, each pod ear communicate with one another using its entry, as demonstrated in Figure 43 and Figure 44, wherea ping command is executed between the pod present in one of the clusters.

alex@alex-pc:-\$ kubectl exec -it -n fcc6f34d-f33a-4bfe-a66b-764d9bedd113 --context nemo-s1 ec ho-server-integration-500-7ccbd97fc-2ctrw -c echo-server -- ping echo-server-integration-501. spain-network-1.inter.l2sm -c 3 PING echo-server-integration-501.spain-network-1.inter.l2sm (10.1.1.1): 56 data bytes 64 bytes from 10.1.1.1: seq=0 ttl=64 time=30.564 ms 64 bytes from 10.1.1.1: seq=1 ttl=64 time=1.234 ms 64 bytes from 10.1.1.1: seq=2 ttl=64 time=1.225 ms ---- echo-server-integration-501.spain-network-1.inter.l2sm ping statistics ----3 packets transmitted, 3 packets received, 0% packet loss round-trip min/avg/max = 1.225/11.007/30.564 ms

Figure 43. Ping between pods in S1 and S2 clusters

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	67 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



alex@alex-pc:~\$ kubectl exec -it -n fcc6f34d-f33a-4bfe-a66b-764d9bedd113context nemo-s1 ec
ho-server-integration-500-7ccbd97fc-2ctrw -c echo-server ping echo-server-integration-501.
spain-network-1.inter.l2sm -c 3
PING echo-server-integration-501.spain-network-1.inter.l2sm (10.1.1.4): 56 data bytes
64 bytes from 10.1.1.4: seq=0 ttl=64 time=7.069 ms
64 bytes from 10.1.1.4: seq=1 ttl=64 time=1.216 ms
64 bytes from 10.1.1.4: seq=2 ttl=64 time=1.311 ms
echo-server-integration-501.spain-network-1.inter.l2sm ping statistics
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 1.216/3.198/7.069 ms

Figure 44. Ping between pods in S1 and S2 clusters

Since the execution of the ping was successful, these results showcase that NEMO workloads can communicate with specialised networking tools thanks to the flexibility of the mNCC

5.3 NEMO workload scheduling & orchestration

This integration scenario aims to illustrate the workload migration process and orchestration. This process is represented by three different tasks:

- Intent-Based workload migration
- CFDRL workload migration
- Workload horizontal scaling

The first two tasks concern the migration of the workload migration from one cluster to another. The main difference concerns the component that triggers the said migration. This process is depicted in Figure 45.



Figure 45. Intent-Based Migration Sequence Diagram

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	68 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



For the Intent-Based migration, the Intent-Based API is responsible for triggering the migration process by publishing an intent in RabbitMQ. This intent contains several requirements that the cluster where the workload is currently deployed must comply with. If one of the requirements is not met, the IBMC selects a new target cluster and performs the migration.

In the case of the CFDRL migration, machine learning algorithms are used to detect whether the cluster where the workload is deployed is not suitable anymore, using the information from workload metrics and trained models. A new target is then selected by the CFDRL and a migration message is sent to the IBMC to trigger migration as shown in Figure 46.



Finally, the workload horizontal scaling scenario also relies on the CFDRL component to detect the need of upscaling or downscaling the number of replicas of a deployed workload. When this happens, a message is sent to the MO, which performs the scaling action, Figure 47.



Figure 47. Horizontal Scaling Sequence Diagram

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	69 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Objective	The objective of this task is to validate the NEMO workload migration process.
Components	IBMC Intent-Based API MO CFDRL CMDT PPEF mNCC
Features to be tested	 Workload migration: By the creation of an intent, containing one or more requirements which the deployment cluster must met. If any of these requirements is not fulfilled, the ibmc-controller triggers a migration to a more suitable cluster By the CFDRL component once its inference states that it is preferable for the workload's optimal operation to be moved from cluster A to cluster B. Once the request is communicated to the Meta-Orchestrator component then the workload migration is executed. Horizontal scaling of a deployed workload
Test setup	All the participating components were deployed in the NEMO meta-OS cloud/edge infrastructure at OneLab dev cluster Land staging cluster 1).
l est setup	All the participating components were deployed in the NEMO meta-OS cloud/edge infrastructure at OneLab (dev cluster 1 and staging cluster 1).

5.3.1 Verification scenario





Test 3: NEMO w	orkload migration
Steps	The steps identified in the associated sequence diagrams are listed below:
	a. Intent-Based Migration
	1. Intent-based API publishes an Intent with one or more requirements (Availability, GreenEnergy, CPU, Memory etc).
	2. Ibmc-controller retrieves workload status from the Intent-Based API.
	3. Ibmc-controller retrieves the managed clusters information from the MQ
	4. If the workload is already deployed and the intent requirements are not met, then a migration action is triggered.
	5. A new target cluster compliant with the intent requirements is selected from the available managed clusters list.
	6. Ibmc-controller sends a message to the source cluster ibmc-agent.
	7. A backup of the workload is created and uploaded to Rook-Ceph 3 Bucket located in the HUB cluster.
	8. Upon backup completion, a message is sent to the target cluster ibmc-agent to continue with the migration process.
	9. The backup is restored in the target cluster
	10. A message is sent to the MO notifying the workload migration completion.
	11. MO updates the corresponding Manfestwork to match the new workload deployment status.
	12. A message is sent to the intent-Based API updating the workload status, specifying the cluster where it has been deployed.
	b. CFDRL Migration
	1. CFDRL gathers workload information from CMDT, PPEF and mNCC.
	2. CFDRL algorithm calculates the most suitable cluster for workload migration. 3. CFDRL sends a message to the source cluster ibmc-agent
	4. A backup of the workload is created and uploaded to Rook Ceph S3 Bucket located
	in the HUB cluster.
	5. Upon backup completion, a message is sent to the target cluster ibmc-agent to
	6. AThe backup is restored in the target cluster.
	Amessage is sent to the MO notifying the workload migration completion.
	8. MO updates the corresponding Manifestwork to match the new workload
	deployment status. 9 A message is sent to the Intent-Based API undating the workload status specifying
	the cluster where it has been deployed.
	c. Horizontal Scaling
	1. CFDRL gathers workload information from CMDT, PPEF and mNCC.
	 CFDRL algorithm calculates the required scaling action. CFDRL sends a horizontal scaling message to MO
	4. MO scales up/down the workload replicas.
	5. MO sends completion a message back to CFDRL

5.3.2 Results

The rest of the subsections present the corresponding results step by step as indicated by the verification scenario.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	71 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



5.3.2.1 Intent-Based Migration

Step 1

The initial conditions for the migration scenario require a workload already deployed in the source cluster from where it will be migrated, as shown in Figure 48. LCM Workload Initial Visualization. In this figure, the LCM IU is used to see that the workload with ID "3b004d45-3f08-48b8-b1b3-2abe06509d66" is currently deployed in the dev cluster.

ID Release I	lame	Workload ID	Status	Cluster	Instance ID	Actions		
113 test500		1	DEPLOYED	dev-onelab	3b004d4	~		4
				Items per page:	10 👻	1 – 1 of 1	<	>
		Figure	48. LCM Work	load Initial Visual	zation			
rigger the r	nigration	process, a	n availability i	intent is created a	is shown in	Figure 4	9. Ava	ailal
t Creation	which re	quires that	t the workload	is deployed in a	cluster wit	h an avai	lahilita	a or
nt Creation, 95%.	, which re	equires that	t the workload	is deployed in a	cluster wit	h an avai	lability	y gr
nt Creation, 95%.	, which re	equires that	t the workload	l is deployed in a	cluster wit	h an avai	lability	y gr
nt Creation, 95%.	, which re	equires that	t the workload	l is deployed in a ate Intents	cluster wit	h an avai	lability	y gr
ngger the f nt Creation, 95%.	Use yaml file	equires that	t the workload	l is deployed in a	cluster wit	h an avai	lability	y gr
ngger the f nt Creation, 95%.	Use yaml file	equires that	t the workload	is deployed in a ate Intents	cluster wit	h an avai	lability	y gr
ngger the f nt Creation, 95%.	Use yaml file Workload Instar 3b004d45-	ce ID * 3f08-48b8-b1b3-2	t the workload Cre Cre	is deployed in a ate Intents Start Date-time 01/05/2025, 00:00:00	cluster wit	h an avai	lability	y gr
nt Creation, 95%.	Use yaml file Workload Instar 3b004d45- Intent Type *	ce ID * 3f08-48b8-b1b3-2	t the workload	is deployed in a ate Intents	cluster wit	h an avai	lability	y gr
nt Creation, 95%.	Use yaml file Workload Instan 3b004d45- Intent Type *	ce ID * 3f08-48b8-b1b3-2	t the workload Cre Cre	is deployed in a ate Intents Start Date-time 01/05/2025, 00:00:00 End Date-time 22/06/2025, 00:00:00	cluster wit	h an avai	lability	y gr
nt Creation, 95%.	Use yaml file Workload Instar 3b004d45- Intent Type * Availability Target Name	ce ID * 3f08-48b8-b1b3-2	t the workload Cre Cre Cabe06509d Target Condition	is deployed in a ate Intents Start Date-time 01/05/2025, 00:00:00 End Date-time 22/06/2025, 00:00:00 Target Value	cluster wit	h an avai	lability	y gr

Figure 49. Availability Intent Creation

Steps 2-6

After the intent is created, it is published by the Intent-Based API. In Figure 50, the ibmc-controller deployed in the dev cluster (HUB cluster) receives the message and proceeds to check if the current deployment cluster meets the availability requirement.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	72 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final


Since the dev cluster has an availability value of only 90%, the intent requirements are not fulfilled, so a new target cluster is selected for the redeployment of the workload. The list of managed clusters shows that the staging cluster has an availability of 99.9% thus it is selected as a target cluster for migration.

Finally, a message is sent to the ibmc-agent deployed in the dev-cluster to trigger the migration process.

ignacio@DESKTOP-8888H6I:-\$ kubectl logs -f ibmc-controller-7bf6f6cd8f-4nc67
2025/05/22 13:50:44 Connected to RabbitMQ
2825/85/22 13:50:44 Awaiting messages
2825/85/22 14:45:89 Message received
2825/85/22 14:46:89 {Instance ID: 36864d45-3688-4868-b1b3-2abe06589d66, Morkload status: deployed, Deployment cluster: dev-onelab
2025/05/22 14:46:09 Managed clusters info: map[dev-onelab:map[Availability:90% Cost:low_cost CpuBaseRate:10 GreenEnergy:0% MemoryBaseRate:18 St
atus:True Storage:1620 Vram:0] k3s-onelab:map[Availability:95% Cost:loa_cost CpuBaseRate:10 GreenEnergy:48% MeeoryBaseRate:10 Status:True Storag
e:1888 Vram:8] nemo-smart-farming:map[Availability:3% Cost:low_cost CpuBaseRate:8 GreenEnergy:8% MemoryBaseRate:8 Status:Unknown Storage:8 Vram:
B] pro-onelab:map[Availability:95% Cost:low_cost CpuBaseRate:10 GreenEnergy:28% NemoryBaseRate:10 Status:True Storage:1358 Vram:14] staging-onel
ab:map[Availability:99.9% Cost:low_cost CpuBaseRate:10 GreenEnergy:60% MemoryBaseRate:10 Status:True Storage:1080 Vram:0]]
2825/85/22 14:46:89 Intent target values: [map[targetCondition:IS_GREATER_THAN targetName:availability targetValueRange:95.8]]
2025/05/22 14:46:09 Workload deployed in cluster dev-onelab, proceeding to check if migration is required
2025/05/22 14:46:09 Cluster dev-onelab does not meet intent availability requirement
2025/05/22 10:06:09 Current deployment cluster devionetab does not meet all intent requirements, proceeding to select a new cluster
2025/05/22 10:06:09 Cluster k3s-onelab does not meet intent availability requirement
2025/05/22 14:46:09 Skipping cluster nemo-smart-farming due to invalid or missing Status: Unknown
2025/05/22 14:46:09 Cluster pro-onelab does not meet intent availability requirement
2025/05/22 14:46:09 Cluster staging-onelab meets all intent requirements
2025/05/22 14:40:09 Ressage published to exchange 'mo' with routing key 'dev-onelab'
2025/05/22 14:46:09 Migration triggerend from dev onelab to staging onelab

Figure 50. Ibmc-controller processing the intent

Steps 7 & 8

Figure 51 shows the migration message arriving to the dev cluster ibme agent, where the backup of the workload resources is performed. Once the backup is finished, a message is sent to the staging cluster ibme agent to continue with the migration.

ignacio8DESMTOP-9R08H61: \$ kubectl logs -f ibmc-agent-699978cbf6-jpzn5 2025/95/22 13:58:44 Connected to RabbitMQ 2025/95/22 13:58:44 Amaiting wessages... 2025/95/22 13:58:45 Message received: 2025/95/22 14:46:49 {WorkloadID: 3b884045-3f88-48b8-b1b3-Zabe86589d66, Action: backup, SourceCluster: dev-onelab, TargetCluster: staging-onelab] 2025/95/22 14:46:49 {WorkloadID: 3b884045-3f88-48b8-b1b3-Zabe86589d66, Action: backup, SourceCluster: dev-onelab, TargetCluster: staging-onelab] 2025/95/22 14:46:49 {WorkloadID: 3b884045-3f88-48b8-b1b3-Zabe86589d66, Action: backup, SourceCluster: dev-onelab, TargetCluster: staging-onelab] 2025/95/22 14:46:49 {WorkloadID: 3b884045-3f88-48b8-b1b3-Zabe86589d66, Action: backup, SourceCluster: dev-onelab, TargetCluster: staging-onelab] 2025/95/22 14:46:29 Ressage published to exchange 'wo' mith routing key 'staging-onelab 2025/95/22 14:46:29 Teackuptane: dev-onelab exchange 'wo' mith routing key 'staging-onelab 2025/95/22 14:46:29 Teackuptane: dev-onelab exchange 'wo' mith routing key 'staging-onelab 2025/95/22 14:46:29 Teackuptane: dev-onelab exchange 'wo' mith routing key 'staging-onelab 2025/95/22 14:46:29 Teackuptane: dev-onelab exchange 'wo' zith routing key 'staging-onelab 2025/95/22 14:46:29 Teackuptane: dev-onelab exchange '20250527100000' kernicoadID: 3b00000' 3f88-48b8-b1b3-2abe6569d66, Action: restore, SourceCluster: dev-onelab.

Figure 51. Dev Cluster Ibmc-agent logs

Steps 9-12

In Figure 52, the restore message is received by the staging cluster. There, the ibmc-agent waits for the syncing of the backup uploaded to Rock-Ceph and proceeds to restore the workload resources. When the restore is completed, a message is sent to both MO and Intent-Based API to notify them that the migration has been successfully completed.



Figure 52. Staging Cluster Ibmc-agent logs

The workload status can be checked again in LCM, where the deployment cluster has been updated from the dev cluster to the staging one, as shown in Figure 53.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	73 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



D Release Name	Workload ID	Status	Cluster	Instance ID	Actions	
113 test500	1	DEPLOYED	staging-onelab	3b004d4	~	•
			Items per page:	10 💌	1 – 1 of 1	<

. .



Figure 54. LCM complete workflow view

Migration and Scaling tasks

5.3

Part of the CFDRL has been tailored towards managing and orchestrating workloads in a Kubernetes environment, with a focus on leveraging reinforcement learning (RL) techniques to optimize resource usage and performance of workloads. Below is an explanation of the project's purpose and how it is structured.

The primary goal of this project is to manage workloads efficiently in a cloud-native environment. It uses reinforcement learning (referred to as CFDRL in the project) to make intelligent decisions about workload orchestration. This involves optimizing resource allocation by scaling or migrating workloads dynamically to improve the overall performance of the system.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	74 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Scaling and migrating tasks

In the context of workloads in a Kubernetes environment, **scaling** and **migration** are two fundamental resource management tasks:

- Scaling refers to adjusting the number of replicas (pods) running for a given workload. This is typically done to handle changes in demand: increasing replicas to improve performance or availability when load increases and decreasing them to save resources when demand drops. In Kubernetes, scaling is a common operation and can be performed automatically (e.g., with the Horizontal Pod Autoscaler) or manually. In the CFDRL system, the scaling task involves learning a policy that decides, based on observed metrics (like CPU, RAM, and latence), how many replicas each workload should have at any given time to optimize resource usage and performance.
- **Migration** involves moving a workload (or its replicas) from one cluster or node to another. This can be necessary for load balancing, fault tolerance, cost optimization, or compliance reasons. Migration is more complex than scaling because it requires coordination between clusters or nodes, and may involve data transfer, downtime, or changes in network latency. In the CFDRL system, the migration task is formulated as learning a policy that decides, based on the current state of the system, to which cluster or node a workload should be moved to achieve optimal performance or resilience.

For simplicity, the CFDRL program tackles them by training and praintaining **two separate** reinforcement learning models: one specialized for scaling dectrions and another for migration decisions. Each model learns from its own set of state-action-reward experiences, allowing the system to optimize both tasks independently and effectively in a dynamic Kubernetes environment. The difference between the two models is minor, for this reason the scaling architecture and then report the changes in the migration case are described.

Architectural Overview

The project is built around a modular architecture that integrates several key components:

<u>Containerization with Docker:</u> The project uses Docker to containerize its applications. Multiple Dockerfiles are provided, each tailered for specific purposes. This containerized approach ensures portability and consistency across different environments.

<u>Kubernetes</u> for Orchestration: Kubernetes is used as the orchestration platform, with YAML configuration files defining the deployment and storage requirements. Persistent Volumes (PV) and Persistent Volume Chains (PVC) are configured to manage data storage of the models learned by CFDRL.

<u>Reinforcement Learning Core</u>: The core functionality of the project revolves around reinforcement learning. The main script (main_cfdrl.py) implements the logic for managing workloads and interacting with external APIs at likely uses reinforcement learning algorithms to analyze workload states and make decisions that ortimize system performance. Metrics are logged and stored for further analysis, and old datais periodically cleaned up to maintain efficiency.

<u>Automation and Deployment:</u> Shell scripts are provided to automate the process of building Docker images, pushing them to a container registry, and restarting Kubernetes deployments. This streamlines the deployment process and reduces the potential for human error.

Reinforcement Learning Algorithm in the Project

The CFDRL implements a Deep Q-Network (DQN) algorithm for reinforcement learning (RL). DQN is used to optimize workload orchestration in a Kubernetes-based environment by dynamically deciding the number of replicas for workloads based on system metrics such as CPU usage, RAM usage, and

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	75 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



latency. Below is a detailed explanation of the RL algorithm and its implementation in the project, based on the provided main_cfdrl.py and dqn.py files available in the git eclipse of the NEMO project.

Purpose of the RL Algorithm

The RL algorithm aims to:

- 1. Dynamically adjust the number of replicas for workloads to optimize resource usage and performance.
- 2. Minimize latency and resource over-provisioning while maintaining workload efficiency.
- 3. Learn an optimal policy for workload orchestration through interaction with the environment.

Key Components of the RL Algorithm

1. State Representation: The state is a vector that represents the current environment, including

- Number of replicas for a workload.
- CPU usage and target.
- RAM usage and target.
- Latency.

The function build_state_for_workload(workload_id) in main_cfdrl.py constructs this state vector for each workload by fetching metrics from shared data structures.

<u>2. Action Space and frequency:</u> The action space is discrete and represents possible decisions the RL agent can make:

- Increase the number of replicas.
- Decrease the number of replicas.
- Maintain the current number of replicas

The Discrete class in dqn.py defines this action space, ensuring that actions are valid integers within a specified range.

The frequency at which the CFDRL agent takes actions is set to every 1 to 3 minutes, depending on the specific task being performed. This timing is chosen because some of the critical information provided by other system components, such as workload metrics and cluster states, is only updated once per minute. Additionally, certain actions, especially migrating workloads between clusters—can take several minutes to complete in practice. By spacing out decisions in this way, the system avoids overloading the infrastructure with frequent changes and ensures that each action is based on the most recent and reliable data available from the environment. This approach balances responsiveness with system stability and operational efficiency.

<u>3. Reward Function</u>. The reward function incentivizes efficient resource usage and penalizes inefficiencies. It is defined as:

$$reward = number_{of_{renlicas}} - latency + \max(ram_{usage} - ram_{target})$$

It is helpful to carefully balance all the metrics (e.g., CPU usage, RAM usage, number of replicas).

Negative rewards are assigned for higher latency and more replicas. Positive rewards are given for achieving CPU and RAM targets.

4. Deep Q-Network (DQN)

The DQN is implemented in the DQN class in dqn.py. It consists of a Backbone Network: A multi-layer perceptron (MLP) that processes the state and outputs Q-values for each action and a Target Network: A copy of the backbone network used to stabilize training by providing fixed Q-value targets during updates.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	76 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



The DQNTrainer class manages the training process, including computing the loss using the Bellman equation and updating the target network periodically to improve stability.

5. <u>Replay Buffer:</u> The ReplayBuffer class in dqn.py stores transitions (state, action, reward, next state) for experience replay. Random sampling from it allows the agent to learn from past experiences and break the correlation between consecutive samples. Storage is available with permanent mounted volumes in the Kubernetes manifest so that the replay buffer is available to any run of the RL algorithm.

<u>6. Exploration Strategy:</u> The project uses an epsilon-greedy strategy for exploration, implemented in the EpsilonGreedy class. The agent selects random actions with a probability epsilon, which decays over time to encourage exploitation of learned policies.

<u>Neural Network Architecture:</u> The DQN uses a multi-layer perceptron (MLP) as its backbone network. The architecture is defined in the MLP class in dqn.py.

- Input: State vector (as defined above).
- Hidden Layers: Configurable number of layers and neurons, with ReLU activation.
- Output: Q-values for each action.

Algorithm Workflow:

- 1. Initialization
 - The agent initializes its neural network, replay buffer and exploration strategy.
- 2. Main Loop
 - For each workload:
 - The agent observes the current state.
 - It selects an action (e.g., adjust replicas).
 - The action is executed, and the environment returns the next state and reward.
 - The transition is stored in the replay buffer
 - Learning: The agent samples a batch of transitions and updates the Q-value function using the Bellman equation and DQN criterion.
- 3. Periodic Updates
 - The target network is updated periodically.
 - Model weights are saved to disk at regular intervals.

<u>Defining a reward function for latency</u>. Defining a reward function for latency in this system is inherently challenging due to the nature of the workload distribution and the way latency is measured.

Latency is Linked to Node-to-Node Communication. Latency is typically a measure of the delay in communication between nodes in a distributed system. However, in this case, most replicas are deployed on the same node. As a result, there is little to no measurable latency because intra-node communication is significantly laster and often negligible compared to inter-node communication. When replicas are collocated on the same node, the latency metric does not provide meaningful feedback about the system's performance. This makes it difficult to use latency as a reliable component of the reward function.

Components CFDRL takes Information From

The CFDRL system interacts with multiple components to gather information and make decisions for workload orchestration Figure 55 and Figure 56. These interactions are facilitated through RabbitMQ (via the pika library) and HTTP API requests (via the requests library). Below is an explanation of the components CFDRL communicates with and how these communications are implemented.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	77 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final





Meta-Orchestrator (MO): The Meta-Orchestrator is responsible for managing workload orchestration decision. It receives messages from CFDRL about the number of replicas for workloads and provides feedback.Communication : RabbitMQ is used for communication. Messages are sent to the cfdrl_mo_hs queue using RabbitMQ's basic_publish method.

CMDT (Cluster Monitoring and Decision Tracking): CMDT provides information about the current state of workloads, such as the number of replicas and the node where the workload is running.Communication: RabbitMQ is used for communication. Messages are received from the CMDT queue.

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	78 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



NMCC (Network Metrics Cost Calculator): NMCC provides network-related metrics, such as latency between nodes or clusters. Communication: RabbitMQ is used for communication. Messages are received from the mncc.metrics.cfdrl queue.

PPEF through Intent API: The Intent API provides workload-specific metrics, such as CPU and RAM usage targets, as well as achieved values. It also provides information about workload clusters. Communication: HTTP GET requests are sent to the Intent API using the requests library. The API response is parsed, and workload metrics are extracted and stored in shared data structures.

Shared Data Structures and Multi-threaded implementations

The program is designed to be multithreaded, enabling concurrent execution of tasks while sharing data between threads. This approach is crucial for handling multiple sources of information (e.g., habbilWIQ queues, APIs) and performing reinforcement learning computations simultaneously. Below is an explanation of how multithreading is implemented and how shared data is managed.

The program uses Python's threading module to create and manage multiple threads. Each thread is responsible for a specific task, such as consuming messages from RabbitMO queues calling APIs, or running the reinforcement learning loop. The threads start concurrently and tun in parallel.

The information gathered from these components is stored in shared data structures, such as:

- SharedWorkloadData: Stores workload-specific metrics (e.g., replicas, latency, CPU/RAM usage).
- SharedReplayBuffer: Stores transitions for reinforcement learning.
- SharedMNCCBuffer: Stores network metrics from MCC.

These shared data structures are protected by locks (Lock) to insure thread-safe access.

Running the CFDRL for the Scaling task:

META ORCHESTRATOR:
MO REPLY received new message b'{"cluster name":"k3s-onelab"."workload id":"caedfa4f-daa6-4049-a89b-d722d14317b9"."number replicas"
ss 'bytes'>
after processing: received new message {'cluster name': 'k3s-onelab', 'workload id': 'caedfa4f-daa6-4049-a89b-d722d14317b9', 'numbe
: ''} of type <class 'dict'=""></class>
CMDT:
CMDT: Time 2025-05-23 07:09:06.063786 cmdt received new message {'cluster': 'dev-onelab', 'region': 'eu-west-1', 'node': 'nemo-dev-
s-kube-state-metrics', 'job': 'kubernetes-service-endpoints', 'pod': 'echo-server-intent-18-7659fd576c-6cgk7', 'uid': '8d478b62-a86
069aef0', 'current_state': 'Running', 'status_changes_1h': 0, 'traffic_stats': {'req_bytes': None, 'req_rate': None, 'res_bytes': N
<pre>}}, 'replicas': 1} of type<class 'dict'=""></class></pre>
replicas {'caedfa4f-daa6-4049-a89b-d722d14317b9': 1, 'dae8689c-d810-4a3b-872b-c6cbfc52a3e4': 1, 'e3b30d15-3f72-45b6-99ff-4a3c888535
dd-6d910d35b27d': 1, '985c2858-58f5-4dcf-9d92-606f3eec27ce': 1, '9fc817e5-0578-4186-8d9a-09a127720f25': 1, 'aff12104-0054-4afa-a15c
f6c-45b6-af30-a3ba0c232612': 1, 'da738fb0-d84a-48b2-b5e1-e80a91443e4e': 1, 'f8554c37-c8b6-4058-b7db-7e882edccd33': 1, '092f27a0-a9a
'0fd78817-a411-432f-81f0-89addbc2da8b': 1, '123e093e-56a7-405d-9f26-c38aa8d28ab6': 1, '1955d604-245e-467d-8b65-d06abbc0a605': 1, '4
40c5b': 1, '61edb162-1d52-4e4d-8b55-275187a780ba': 1, '6ca37a04-6b61-4631-953a-6ba7f4215175': 1, '7d7f4430-9e71-4b80-a9de-1ddd684cf
7b0-3a1dac4cd6a5': 1, '970ef7b2-158f-4c88-9f6a-5d1ecb40f369': 1, '137bd579-85da-47b7-b1d3-bfbbf4353142': 1, '32e53f55-3a76-4401-8a7
28e1-4f73-b45c-21cc02991bc5': 1, '781052a5-4270-4113-847a-8730cdf55ba7': 1, '9d95bbc7-698c-41e7-9f2e-dd3de3d32d25': 1, 'a9f47e17-59
'e860d4cc-f915-4314-941d-5946c571daf1': 1, 'fcc6f34d-f33a-4bfe-a66b-764d9bedd113': 1, '2b9c9ade-0296-4cf6-b0b9-6b83a069aef0': 1, '
0a957f': 1, '5b065d03-f617-463b-8425-7de0f4594c28': 1, '0c8ee5f5-bc41-42ed-8a24-71e974dc219b': 1, '3b004d45-3f08-48b8-b1b3-2abe0650
sending request to intent api
PPEF RESPONSE
PPEF: 11me 2025-05-23 07:09:09.555693
response «Response [200]> <class requests.models.response=""> 200</class>
worktoad_td edi14289-0615-41e3-9050-5a0050080227 cpu_achteved None ram_achteved None cpu_usage_target 500 ram_usage_target 100
work toad_td_dtocs4f4_cgr3_4225_4265_abe523f4d1_cpu_achieved_V35_ram_achieved_V60_v6_cpu_usade_target_S00_ram_usade_target_100
work toad_to 0/964914-c8/3-4231-8634-80500321091 Cpu_achteved wone fam_achteved wone Cpu_usage_target 300 fam_usage_target 100
worktoad_td udesosserusionasionasionario et al. and an entered service and an entered of an entered of a construction and a service and a serv
work toau_to caediaa adadada a a a a a a a a a a a a a a

Figure 57. Messages used for Scaling

The CFDRL program was deployed and executed continuously on the OneLab cluster of the NEMO platform, where it interacted in real time with other system components such as the Meta-Orchestrator, CMDT, NMCC, and the Intent API. Over a period of 3 days, with the agent collecting one state-action-reward tuple per minute, the system accumulated a total of 4,320 transitions (3 days \times 24 hours \times 60 minutes). After this training period, the reinforcement learning agent achieved an average reward of - 5.84, indicating improved performance compared to a random policy, which yielded an average reward of -6.54. One way to interpret this is that compared to the random policy, CFDRL almost save on average

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	79 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



time 2025-05-27 13:02:38

one replica to obtain the same cpu usage. This demonstrates that the RL-based approach was able to learn and apply more effective scaling decisions than random actions in this dynamic environment.

Running the CFDRL for the Migration task:

u<mark>ben@debian-ruben:~\$ k log</mark>s cfdrl-rabbit-deployment-migration-555b8f6d6c-m2ndw -n nemoent message to CLuste<u>r</u>s: {'WorkloadID': 'caedfa4f-daa6-4049-a89b-d722d14317b9', 'Action

First Note that the model used for migration was very similar to the one in scaling but adding also in the state space information about the current location of all the workloads. The action space was then the destination cluster chosen.

After learning a model specific to the migration task, we run the CFDRL with and without the migration actions and measured over a period of one day an increase of our reward from -11.2 to -10. At hanks to the migration actions hinting at an improvement due to our model.

However, we also noticed that the variability of the reward over long period of time makes the comparisons between methods complexes.

Viewing the CFDRL and MO and Cluster IBMC interactions:

Here we report screenshots that demonstrate how an action is decided by the FDRL and then implemented by either the Meta orchestrator or the Cluster IBMC.

1. Migration:

Figure 58. Log of the action taken by CFDRL and soft to the Cheter Ibmc

ruben@debian-ruben:_\$ k logs ibmc-agent-76c9d499d7-mnhsn -n nemo-kernel tail -n 32 head -n 15
2025/05/27 12:58:58 Awaiting messages
2025/05/27 13:02:58 Message received:
2025/05/27 13:02:58 {WorkloadID: caedfa4f-daa6-4049-a89b-d722d14317b9, Action: restore, SourceCluster: pro-onelab, TargetCluster: staging-onelab}
Backup is still in progress. Waiting
2025/05/27 13:04:19 Backup synced, proceding to restore the workload
Restore is still in progress. Waiting
2025/05/27 13:04:39 No deployments found in the namespace
2025/05/27 13:04:39 No pods found in the namespace.
2025/05/27 13:04:39 Migration completed successfully
2025/05/27 13:04:39 Message published to exchange 'nemo.api.workload' with routing key 'update'
2025/05/27 13:04:39 Message published to exchange 'nemo.api.workload' with routing key 'intent-notify'
2025/05/27 13:04:39 {WorkloadID: caedfa4f-daa6-4049-a89b-d722d14317b9, Type: migration, SourceCluster: pro-onelab, TargetCluster: staging-onelab, Timestamp: 20250527130439}

t | tail -n 3 | 'SourceCluster'





Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	80 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



2. Scaling:



5.4 NEMO workload lifecycle management

The workload lifecycle management scenario illustrates the monitoring of a workload during its lifecycle in NEMO meta-OS collecting workload intents, resources consumption, performance metrics and health status of the workload. Figure 64 shows the sequence diagram of this scenario.



Figure 64. Workload Lifecycle Sequence Diagram

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						81 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



5.4.1 Verification scenario

Test 4: NEMO wo	rkload lifecycle management
Objective	To verify the workload lifecycle management process in NEMO covering all the steps identified.
Components	• LCM
	• Intent-API
	• PPEF
	• CMDT
	RabbitMQ
Features to be tested	This integration scenario aims to validate the workload lifecycle management. The NEMO workload intents and complementary measurements that concern the resources' consumption and the resulting performance and liveness of a workload are collected by the PPEF and the CMDT components. From there they are communicated through the RabbitMQ to the LCM UI where there are visualized to the NEMO user.
Test setup	The associated components are deployed in OneLab facilities at NEMO dev cluster 1 The CFDRL component which is undergoes its final stages of development.
Steps	 The NEMO user accesses the LCN UT NEMO workload monitoring collects metrics that correspond to the NEMO workload (PPEF) The collected workload metrics are communicated to the Intent-API The NEMO Cluster monitoring collects measurements that concern the NEMO meta-OS operated clusters (RPEF) The collected cluster metrics are communicated to the RabbitMQ NEMO workload complementary monitoring (CMDT) The collected metrics are communicated to the Rabbit MQ The collected metrics are communicated to the Rabbit MQ The collected metrics are communicated to the Rabbit MQ The LCM aggregates the collected metrics and visualize them to the NEMO user The PREF report intent violations to the Intent-based API

5.4.2 Results

This section outlines the step-by-step process described in the scenario above.

5.4.2.1 The NEWO user accesses the LCM UI

The NEMO user accesses the LCM interface by adding its credentials and gains access to the monitoring dashboard through NEMO Keycloak identity manager. Depending on its role, the user has access to its own workloads or workloads he manages.

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	82 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Sign in to your account Username or email		
Password	۲	
Sign In		\leq

Figure 65. LCM User access

5.4.2.2 NEMO workload monitoring collects metrics that correspond to the NEMO workload (PPEF)

Within the NEMO framework there is a periodic task that evaluates the Intent target conditions for all types of workload intents that are created though the LCM UI. For example, a "DeliverComputingWorkload" Intent allows to the user to set conditions regarding the workload resource usage (CPU, RAM, disk). Figure 66 shows an example of an intent requiring for the workload to be executed in a cluster that can provide more than 2 milliseconds of CPU time.

			►	• •			
ID Name Feasibility Fulfilment Sta Not Fulfilled S Last Updated	atus State I Time	: 128 : DeliverComputingWorklos : FEASIBLE : FULFILLED : COMPLIANT : 2025-05-14T09:24:13.35	ad 9393Z				
v	Vorkload Inst	ance ID *		Start Da	ate-time		
	128			30/0	4/2025, 00:00:00	Ē	
Ir	ntent Type *			End Dat	te-time		
	DeliverCor	nputingWorkload		30/0	5/2025, 00:00:00	Ē	
т	arget Name		Target Condition		Target Value		
	cpuUsage		IS_GREATER_TH	AN	2	FULFILLED	Image: A state of the state

Figure 66. Example of intent conditions

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	83 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Figure 67 (a) and (b) show the logs of the intent evaluator periodic task. The evaluator checks though the PPEF for the values corresponding to the intents' targets and determines if the required conditions are met. The report produced is sent through RabbitMQ for the Intent API to consume.



5.4.2.3 The collected workload metrics are communicated to the Intent-API

Figure 68 demonstrates the logs of the Intent API task which is responsible for receiving the produced intent evaluation reports. The Intent API receives this report and uses it to update its intents' status ("Fulfilled" or "Not Fullfilled").

	Autoscroll:0	n FullScree	en:Off Time	stamps:Off	Wrap:Of		
intent.collector IN	IFO 2025-05-29	14:10:14,586 col	llector Parsed	#17 valid	intents, #0	invalid	intents
intent.collector IN	IFO 2025-05-29	14:11:15,195 col	llector Parsed	#17 valid	intents, #0	invalid	intents
intent.collector IN	IFO 2025-05-29	14:12:15,058 col	llector Parsed	#17 valid	intents, #0	invalid	intents
intent.collector IN	IFO 2025-05-29	14:13:15,319 col	llector Parsed	#17 valid	intents, #0	invalid	intents
intent.collector IN	IFO 2025-05-29	14:14:14,986 col	llector Parsed	#17 valid	intents, #0	invalid	intents

Figure 68. Intent API receives the Intent Evaluation report

The NEMO Cluster monitoring collects measurements that concern the NEMO meta-OS operated clusters (PPEF)

In NEMO a monitoring mechanism has been established that periodically communicates with PPEF to retrieve the CPU, RAM and disk usage metrics for a time window of 1 minute to keep track of the state of the NEMO meta-OS operated clusters. Figure 69 shows the logs of the monitoring task. The message includes the resource usage metrics for all NEMO clusters. This message is published to NEMO RabbitMQ, for the LCM to consume the information and display the changes in the dashboard.

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	84 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



→ kubectl logs nemo-cluster-metrics-collector-29141928-pjwvv -n nemo-svc
[{"cluster": "staging-onelab", "cpu_usage": 3040.962, "memory_usage": 48.4, "disk_usage": 1.1, "timestamp": "2025-05-29 10:48:03.547723+00:00"}, {"clust
er": "dev-onelab", "cpu_usage": 4153.883, "memory_usage": 79.8, "disk_usage": 1.2, "timestamp": "2025-05-29 10:48:03.577198+00:00"}, {"cluster": "pro-on
elab", "cpu_usage": 3327.596, "memory_usage": 40.5, "disk_usage": 1.4, "timestamp": "2025-05-29 10:48:03.618761+00:00"}, {"cluster": "nemo-energy", "cpu
_usage": 1019.052, "memory_usage": 16.5, "disk_usage": 0.8, "timestamp": "2025-05-29 10:48:03.750547+00:00"}, {"cluster": "nemo-smart-media", "cpu_usage
": 737.063, "memory_usage": 14.3, "disk_usage": 0.8, "timestamp": "2025-05-29 10:48:03.952942+00:00"}]
[x] Sent cluster metrics!

Figure 69. Cluster metrics from PPEF

5.4.2.5 The collected cluster metrics are communicated to the RabbitMQ

Figure 70 gives a closer look at the details of the message that is pushed to the RabbitMQ.



Figure 70. Cluster metrics RabbitMQ paylog

5.4.2.6 NEMO workload complementary monitoring (CMDT)

The CMDT component is thoroughly presented in NEMO D2.3 and relies on information gathered by two distinct services:

- kube-state-metrics, collecting pod-related metadata,
- Linkerd and Linkerd-viz, collecting network-related traffic metrics.

Both services forward collected metrics to Prometheus/Thanos services. Finally, the CMDT components internally utilize Prometheus Query Language (PromQL) to retrieve data and expose enriched data directly through the API, as well as via RebbitMQ communication channels, Figure 71.

Readyness	
GET /readyness Health Check	
Liveness	
GET /liveness Health Check	
default	
GET /api/vl/workloads Get Workloads List	
GET /api/vl/pods/{pod_name_or_uid} Get Pod Details	
GET /api/v1/pods/{pod_name_or_uid}/traffic Get Traffic Stats	
GET /api/v1/tree Get Deployment Tree	
GET /api/vl/tree/workloads Get Pruned Deployment Tree	
Schemas	
HTTPValidationError > Expand all object	
PodInfoExtended > Expand all object	
TrafficStats > Expand all object	
ValidationError > Fynand all object	

Figure 71. The CMDT component includes a SwaggerUI endpoint with up-to-date documentation and examples.

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	85 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



The CMDT provides the following complementary information:

- 1. The `/api/v1/workloads` endpoint provides a list of workloads on the NEMO cluster. Each NEMO workload has a UUID, and CMDT provides a list of pods from the corresponding NEMO workload. Pod information includes the pod's location (NEMO cluster, region, node, tenant), namespace, service, job, pod's name, and pod's unique ID.
- 2. The `/api/v1/pods/{pod_name_or_uid}` endpoint provides detailed information with all the data provided previous point and adds pod's current status (running, terminated, restarting, etc), number status changes within last hour (if pod was restarting), number of replicas of the same pod, and traffic summary. The traffic summary provides the incoming and outgoing bytes of network traffic, the traffic rate (requests per second), response status codes and their rate, and response times presented in percentiles (in milliseconds), Figure 72.
- 3. The `*/api/v1/pods/{pod_name_or_uid}/traffic*` endpoint provides a pod's detailed traffic view, where CMDT provides bytes of incoming or outgoing traffic and provides the name of external pods communicating with the corresponding pod, Figure 73.
- 4. The `*/api/v1/tree*` endpoint provides a tree-like representation of the relationships between all deployments, replicasets, and pods.
- 5. The `*/api/v1/tree/workloads*` endpoint provides a tree-like representation of all deployments, replicasets, and pods, excluding non-NEMO workload components, Figure 74.



Document name:	D4.3 A Final ve	14.3 Advanced NEMO platform & laboratory testing results. inal version					86 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final





Figure 73. Detailed information about the pod's traffic. Additionally, it displays its internal traffic, as well as inbound and outbound traffic with threign pods.

GET	/api/v1/tree/workloads Get Pruned Deployment Tree	^
Parameters	8	Cancel
No parame	ters	
	Execute	Clear
Responses	5	
Curl		
curl -X 'G 'http://	SET' \ localhost: 5000/api/v1/tree/workloads' \ web. amd fortfor/fern'	
-A alle	br. shkrrernul/jour	6
Request URI	L	
Server respo	олое	
Code	Details	
200	Response body	
3	<pre>"dellow-test-cebs_arrow": { "isid": "Doubyment", "resplicatest": { "dellow-test-cebs_arrow": { "dellow-test-cebs_arrow": { "dellow-test-cebs_arrow": 3 "dellow-test-cebs_arrow": 3 "dellow-test-cebs_arrow": 3 "dellow-test-cebs_arrow": 3 "dellow-test-cebs_arrow": 3 "dellow-test-cebs-arrow": 3 "dellow-test-c</pre>	Download,
Code	Description	Links
200	Successful Response	No links
	Media type application/json	

Figure 74. Tree-like representation of all deployments, replicasets, and pods, excluding non-NEMO workload components

Document name:	D4.3 Ao Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	87 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



5.4.2.7 The LCM aggregates the collected metrics and visualizes them to the NEMO user

The LCM gathers data from the previously mentioned sources and consolidates it into a unified user interface, providing a comprehensive view of the workload lifecycle. This includes historical data, near real-time performance insights, and management capabilities, Figure 75.

Sea	Irch									
Create V	Workload									
ID ↓	Workload Name	Version	Status	Ingress Support	User	Balance	Actions			
59	Smart4Game WL1	1.0.0	PENDING	•	nenaddstojanovic@gmail.com	-	Þ	∢		
58	hello-world-api	0.2.0	ACCEPTED	•	seitanidis@synelixis.com	-	•	⊘	Î	
56	hello-world	0.1.1	ACCEPTED	•	seitanidis@synelixis.com	-	•	⊘	Î	
54	hello-world	0.1.0	PENDING	•	pedro.branco@xilbi.com	-	Þ	⊳	Î	
48	echo-server	0.5.17	PENDING	•	lakka@synelixis.com	-	•	⊳	Î	
45	smart-xr-use-case	0.1.6	ACCEPTED	•	lakka@synelixis.com	-	A	⊛	Î	
44	smart4game	0.0.1	PENDING	•	nenaddstojanovic@gmail.com	-	Image: A start of the start	\bigcirc	Î	
43	smart-xr-use-case	0.1.5	ACCEPTED	•	lakka@synelixis.com	-	•	⊘	Î	
42	smart-xr-use-case	0.1.4	ACCEPTED	٠	lakka@synelixis.com	-	•	€	Î	
41	smart-xr-use-case	0.1.3	ACCEPTED	٠	lakka@synelixis.com	-	The second se	€	Î	



Each workload is associated with its instances in NEMO meta-OS. The user can manage workload instances, Figure 76, and get detailed information on its activity in a status timeline, Figure 77.

			□ Wo	rkload Instances			
Sear	rch						
ID ↓	Release Name	Workload ID	Status	Cluster	Instance ID	Actions	
105	smart-xr-7	43	DELETED	dev-onelab	2ea5bcc	~ 🗈	* =
104	smart-xr-6	43	DELETED	dev-onelab	7a7d5e7	~ B	*
103	smart-xr-5	43	RENDERED	dev-onelab	3cdeda9	~ B	*
102	smart-xr-4	43	RENDERED	dev-onelab	1c2c09c	~ B	A



Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborate	ory testing r	esults.	Page:	88 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Sea	irch		
ID ↓	Release Name	May 2025	
105	smart-xr-7		* =
104	smart-xr-6	May Deletion 12 Deployment_CI	* 1
103	smart-xr-5	uster rev- onelab, 2025	*
102	smart-xr-4		* =
		May Deployment	1 = 4 of 4 <
		12 ster: dev-onelab, 09:32 PM	
		2025	
		May Rendered	
		12 ¹ 09:32 PM	
		2025	

Figure 77. Workload Instance lifecycle timeline

The LCM collects information about the deployment of each workload instance, the ReplicaSets and the number of pods as it is provided by CMDT. Based on this data LCM displays the deployment tree of the workload visualizing the Deployment controller used to manage the ReplicaSets and pods in a tree-like structure, Figure 78.

75 test-echo2	25	DELETED	dev-onelab	439089d	~ 🗎	A 1	i –
Deployment ID : 8b72c5f1-88 Name : echo-server-i	34-410-3869-6266989cf1 hegration-300	Replica Set Name :ecto-server-integr Cluster :staging-anelab	aton-300-6df5d9fc8b	Pod ID : 3877d15c:bde9-4bbe-b Name :echo-server-integration- Cluster Cluster :staging-onelab Pod ID ID : 622e9d281-5874-4800-9; Name : echo-server-integration- Cluster Cluster : staging-onelab Pod ID ID : c3305c27-04ab-4018-40 Name : echo-server-integration- Cluster Cluster : staging-onelab Pod ID ID : 43a4dabc-a041-46db-bc Name : echo-server-integration- Cluster Cluster : staging-onelab	16-66569650c8 300-64549fc6b-2s164 448-668e0139075 300-64549fc6b-2s164 448-954c8778061d 300-64549fc6b-mxx9 464-954c8778061 300-64549fc6b-xx375		

Figure 78. LCM workload deployment tree

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	results.	Page:	89 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



While each workload instance is associated with one or more intents, the LCM dashboard provides information about the attached intents and their status, Figure 79. This includes intent violations as presented in the next section.

Search				
4955528				
Create Inten				
	•			
ID ↓	Status	Intent Type	Workload Instance ID	Actions
114	FULFILLED	Availability	4955528	💢 > 🔳 🖀
113	FULFILLED	Availability	4955528	∭ → = ≡
	FULFILLED	Availability	4955528	ji → = i
112				H =
112 111	FULFILLED	Availability	4955528	<u>, , , , , , , , , , , , , , , , , , , </u>

Figure 79. Workload intents management

Finally, the LCM interface shows resource provisioning monitoring and management, Figure 80, providing useful insights in cluster performance (CPU usage, memory usage, disk usage) both in historical and runtime, Figure 81.

0									
Search									
eate Cluster									
) †	Name	Status	Endpoint	CPUs	Storage	Memory	Vram	Actions	
933189e	dev-onelab	ОК	https://api.main.nemo.onelab.eu:6443	68	1620	172	0		
8d63cc0	test-cluster	ОК	https://api.main.nemo.onelab.eu:6443	10	300	120	0	5 E	
67c21fc	pro-onelab	ОК	https://api.prod.nemo.onelab.eu:6443	32	1350	62	14	S 🕯	
60a2ee7	k3s-onelab	ОК	https://api.s2.nemo.onelab.eu:6443	16	1080	32	0	5 E	
1e58cae	nemo-smart-farming	ОК	https://83.235.169.221:35443	6	1000	8	0	S 🕯	
179592b	staging-onelab	ОК	https://api.staging1.nemo.onelab.eu:6443	32	1080	62	0	s =	
142e334	energy	ОК	https://comsensus.eu:13387	5	5	5	0	S 🕯	
						ltems pe	r page: 10		

Figure 80. Resource provisioning

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	ory testing r	esults.	Page:	90 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final





5.4.2.8 The PPEF report intent violations to the Intent-based API

The PPEF reports the intents status for each workload instance and the LCM displays the status of each intent along with detailed information as presented in Figure 82 if the intent is fulfilled and Figure 83 if the intent has been violated.

125	NC							1
124	ID Name	: 120 : Availability						-
123	Feasibility Fulfilment Status	: FEASIBLE : FULFILLED					•	
122	Not Fulfilled State Last Updated Time	: COMPLIANT : 2025-05-28T04:43:14.0	80468Z					•
121	NO	natanco ID X		Start Data time				I
120	FU	iistance iD		Start Date-time				
119	120			15/04/2025, 00:00:00				-
118	NC Intent Type	*		End Date-time				-
117	Availab	bility		30/04/2025, 00:00:00	×.			
116	Target Nam	ne	Target Condition	Target Value				
115	availab	ility	IS_EQUAL_TO	90	FULFILLED	2 🖻		1
114	FU							1

Figure 82. Intent Fulfilled

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	91 of 111		
Reference:	D4.3	Dissemination:	PU	D4.3 Dissemination: PU Version: 1.0					



127	NOT_FULFILLED FULFILMEN	TFAILED	MachineLearning	f0407bb	۲.	•	•
126	NC		►	•			Î.
125	NC						Î
124	Name	: 119 : cloud_continuum					i
123	Feasibility Fulfilment Status	: FEASIBLE : NOT_FULFILLED					i
122	Not Fulfilled State Last Updated Time	: COMPLIANT : 2025-04-14T12:50:08	882046Z				i
121	NC						Î
120	FU 110	istance ID *					i
119	NC						Î
118	Intent Type						
117		continuum					Î.
116	Target Nam	e	Target Condition	Target Value			
115	isolate	dConnectivity	IS_EQUAL_TO	true			1
114	FU						

Figure 83. Intent Violation

5.5 NEMO Secure Execution Environment

The following scenario focuses on the deployment of unikernets through the Secure Execution Environment (SEE) component, starting from the meta-Orchestrator (MO). The creation, modification, and deletion of unikernels are essentially the same at this level.

Unikernel Creation



Figure 84. Unikernel creation sequence diagram

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	92 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Test 5: Unikernels	;
Objective	Deploy Unikernels inside NEMO premises
Components	Meta-Orchestrator (MO)
	• RabbitMQ
	• SEE
Features to be	This scenario validates the deployment of unikernel into a specific node. The SEE interface is
testeu	RabbitMQ. In particular, MO triggers this scenario to deploy the hermit technology solution (unikernel).
Test setup	SEE, RabbitMQ and MO are deployed in OneLab.
Steps	 a. Meta-Orchestrator sends a <i>create unikernel</i> message to RabbitMQ. b. SEE reads the creation request from RabbitMQ. c. SEE performs the creation process internally. d. SEE deploys the created unikernel to the Note. e. If the deployment is successful, ShE sends an ok message to RabbitMQ. f. If there is an error, SEE sends an error message to RabbitMQ.

5.5.1 Verification scenario

5.5.2 Results





Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	93 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final





Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	94 of 111
Reference:	D4.3	Dissemination:	Status:	Final			



see		^
POST /	ublishToSee Using the next verbs: create, delete, apply	^ ii
Parameters		Try it out
Name	Description	
info * required object (body)	Info to trigger the SEE component Example Value Model	
	<pre>{ "body": { "additionalProp1": () }, reply_to": "string", "verb": "create" }</pre>	
	Parameter content type application/json v	
Responses		Response content type application/json
Code (Pescription	
200)k	
200 (
200 (xample Value Model	

Figure 85. publishToSee endpoint

Depending on the result, the SEE generates and queues a message into RabbitMQ. If something is wrong, the message appears as in Listing 13. Otherwise, the server is sending a message with *a correlation_id* but without a payload Figure 86.

{
correlation_id: 4fa6082a-7479-4a5d-b137-475237d24972,
"ErrStatus": {
"apiVersion": "v1",
"code": 404,
"details": {
"kind": "pods",
"name": "nginx"
"kind": "Status",
"message": "pods \"nginx/" not found",
"metadata":
}, Feas on" "NotFound",
"status": "Failure"

Listing 13. SEE error response

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	95 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Message 1	
The server repor	ted 0 messages remaining.
Exchange	(AMQP default)
Routing Key	see-interface-response
Redelivered	0
Properties	correlation_id: 4aba1ccb-e203-49ee-a208-9089a759b4f0 content_type: application/json
Payload 0 bytes Encoding: string	
	Figure 86. SEE success message
In the end, depending on the a or delete), the message is nemo.see.delete).	action represented as the "verb" field in the payload endpoint (create, apply, published in different queues (nemo.see.create, nemo.see.apply, or
	queue: nemo.see.create nemo.see.apply nemo.see.delete



Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	96 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



6 Conclusions

Renge

In this document the latest updates regarding the WP4 tools (LCM, IBMC, MOCA and IBA) were presented. The major highlight of this document is the set of integration workflow pipelines presented in section 5. In this document the integration pipelines across WP2, WP3 and WP4 were extended to support the complete integration and validation tests of the final NEMO meta-OS. This provides the reader with a complete and up-to-date vision of what has been developed within the NEMO framework as well as how the components interact with each other. In addition, this document is intended to be used as the handbook of NEMO for other existing and future projects, a detailed user guidelines of the NEMO meta-OS was presented.

With the completion of the NEMO components' development, a thorough testing and integration procedure took place to guarantee that the initial objectives of the project as well as the various KPIs defined have been successfully met. The work carried out within the framework of WP4 lead to a fully functional and complete toolset that provides an easy-to-use ZeroDevOpa platform for managing services in the form of workload across the IoT-Edge-Cloud continuum.

Finally, the resulting final version of the NEMO meta-OS will be further validated and verified in the framework of NEMO pilots. The corresponding information will be included in deliverable D5.4 *"NEMO Living Labs use cases evaluation results. Final version"*. Moreover, D5.4 is also forseen to contain results stemming from the integration activities of the NEMO Open Call Projects in the context of NEMO meta-OS pilots.

 Document name:
 D4.3 Advanced NEMO platform & laboratory testing results. Final version
 Page:
 97 of 111

 Reference:
 D4.3
 Dissemination:
 PU
 Version:
 1.0
 Status:
 Final



7 References

- [1] NEMO, "D4.2 Advanced NEMO platform & laboratory testing results. Intermediate version," HORIZON 101070118 NEMO Deliverable Report, 2024.
- [2] NEMO, "D4.1 Advanced NEMO platform & laboratory testing results. Initial version," HORIZON - 101070118 - NEMO Deliverable Report, 2023.
- [3] NEMO, "D1.2 NEMO meta-architecture, components and benchmarking. Initial version," HORIZON - 101070118 - NEMO Deliverable Report, 2023.
- [4] GitLab, "The Role of AI in DevOps," 2023. [Online]. Available: https://about.gitlab.com/topics/devops/the-role-of-ai-in-devops/.
- [5] Docker. [Online]. Available: https://www.docker.com/. [Accessed 23 06 2023]
- [6] Flux, "Flux the GitOps family of projects," 2023. [Online]. Available, https://fluxcd.io.
- [7] Flux, "Flux multi-tenancy," 2025. [Online]. Available: https://fluxcd.io/flux/installation/configuration/multitenancy.
- [8] B. Ruecker, "The Microservices Workflow Automation Cheat Sheet: The Role of the Workflow Engine," 2020. [Online]. Available: https://camuda.com/blog/2020/02/the-microservicesworkflow-automation-cheat-sheet-the-role-of-the-workflow-engine/.
- [9] Kubernetes, "Good practices for Kubernetes Secrets," 2023. [Online]. Available: https://kubernetes.io/docs/concepts/security/secrets-good-practices/.
- [10] Kubernetes, "Controlling Access to the Kubernetes API," 2023. [Online]. Available: https://kubernetes.io/docs/concepts/security/controlling-access/.
- [11] "Kustomize Kubernetes Native Configuration Management," [Online]. Available: https://kustomize.io/. [Accessed 22023].
- [12] NEMO, "D1.3 NEMO meta-architecture, components and benchmarking. Final version," HORIZON - 101070118 - VEMO Deliverable Report, 2024.
- [13] Keycloak, "https://www.keycloak.org," [Online]. Available: https://www.keycloak.org.
- [14] "Helm," [Onine]. Available: https://helm.sh/. [Accessed 12 2023].
- [15] OPENAVI Initiative, "OpenAPI Specification," [Online]. Available: https://www.openapis.org. [Accessed 2024].
- [16] RabbitMQ, "RabbitMQ," [Online]. Available: https://www.rabbitmq.com.
- [17] Argo CD, "Introduction to ApplicationSet controller," 2023. [Online]. Available: https://argocd.readthedocs.io/en/stable/operator-manual/applicationset/.
- [18] Argo, "Argo CD Declarative GitOps CD for Kubernetes," 2023. [Online]. Available: https://argo-cd.readthedocs.io/en/stable/.
- [19] Argo, "Argo Workflows," 2025. [Online]. Available: https://argoproj.github.io/argo-workflows/.

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	98 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



- [20] NEMO, "D3.3 NEMO Kernel. Final version," HORIZON 101070118 NEMO Deliverable Report, 2025.
- [21] gRPC, "gRPC A high performance, open source universal RPC framework," 2023. [Online]. Available: https://grpc.io/.
- [22] Google Cloud, "Cloud Endpoints for gRPC," 2023. [Online]. Available: Cloud Endpoints for gRPC.
- [23] R. Gancarz, "Why LinkedIn chose gRPC+Protobul over REST+JSON: Q&A with Karthik Ramgopal and Min Chen," 2023. [Online]. Available: https://www.infoq.com/news/2023/12/linkedin-grpc-protobuf-rest-json/.
- [24] Camunda, "Zeebe: Cloud Native Workflow and Decision Engine," 2023. [Online]. Available: https://camunda.com/platform/zeebe/.
- [25] D. Goel, "Elevate Kubernetes Security with Zero Trust," 2023. [Online]. Available: https://d2iq.com/blog/elevate-kubernetes-security-zero-trust.
- [26] A. Syed, "Zero Trust Security for Kubernetes with a Service Mesh," 2022. [Online]. Available: https://www.hashicorp.com/blog/zero-trust-security-for-kubernetes with-a-service-mesh.
- [27] KongHQ, "Kong Mesh Modernized service mesh for development and governance," 2023. [Online]. Available: https://konghq.com/products/kong-mesh.
- [28] M. Palladino, "Service Mesh vs. API Gateway: What's The Difference?," 2020. [Online]. Available: https://konghq.com/blog/enterprise/the-difference-between-api-gateways-and-servicemesh.
- [29] I. Krutov, "Architecting Zero Trust Security for Kabernetes Apps with NGINX," 2022. [Online]. Available: https://www.nginx.com/blog/architecting-zero-trust-security-for-kubernetes-appswith-nginx/.
- [30] Kata Containers, "About Kata Containers," [Online]. Available: https://katacontainers.io.
- [31] Kepler Contributors, "Kubernetes Efficient Power Level Exporter (Kepler)," 2023. [Online]. Available: https://sustainable-computing.io/.
- [32] Pixie, "Open source Kubernetes observability for developers," 2023. [Online]. Available: https://px.dev/.
- [33] The Linux Foundation, "Kubernetes," [Online]. Available: https://kubernetes.io/. [Accessed 6 2025].
- [34] The Linux Foundation, "Cloud Native Computing Foundation," [Online]. Available: https://www.cncf.io/. [Accessed 6 2025].
- [35] The Linux Foundation, "Falco," [Online]. Available: https://falco.org/. [Accessed 6 2025].
- [36] "trivy-operator," [Online]. Available: https://github.com/aquasecurity/trivy-operator. [Accessed 6 2025].
- [37] Flower Labs GmbH, "Flower," open source, [Online]. Available: https://flower.dev. [Accessed 6 2025].
- [38] AsynvAPI Initiative, "Building the future of Event-Driven Architectures (EDA)," [Online]. Available: https://www.asyncapi.com. [Accessed 6 2025].

Document name:	D4.3 A Final ve	.3 Advanced NEMO platform & laboratory testing results. al version					99 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



- [39] Swagger, "Open API Specification," [Online]. Available: https://swagger.io/resources/open-api/. [Accessed 6 2025].
- [40] Swagger, "Swagger Editor," [Online]. Available: https://swagger.io/tools/swagger-editor/download/. [Accessed 6 2025].
- [41] Robot Framework Foundation, "Robot Framework," [Online]. Available: https://robotframework.org. [Accessed 6 2025].
- [42] BlazeMeter, "Taurus," [Online]. Available: https://gettaurus.org. [Accessed 6 2025].
- [43] Selenium, "Selenium," 2025. [Online]. Available: https://www.selenium.dev/.

Pendine

- [44] Istio, "The Istio service mesh," 2025. [Online]. Available: https://istio.io/latest/about/service-mesh/.
- [45] Envoy, "Envoy Gateway," 2025. [Online]. Available: https://gateway.envoyproxy.io.
- [46] KongHQ, "Kong Gateway," 2025. [Online]. Available: https://konghq.com/products/kong-gateway.

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	100 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Annex 1-Network Intent

Intent:

```
id: '\''mncc l2sm 1'\''
userLabel: '\''cloud continuum'\''
intentExpectations: -
expectationId: '\''1'\''
expectationVerb: '\''DELIVER'\''
                                                                 or of
expectationObject:
objectType: '\''K8S L2 NETWORK'\''
objectInstance: '\''spain-network-8'\''
objectContexts: -
contextAttribute: '\''name'\''
contextCondition: '\''IS_EQUAL_TO'\''
contextValueRange: '\''spain-network-8'\''
contextAttribute: '\''providerName'\''
contextCondition: '\''IS_EQUAL_TO'\''
contextValueRange: '\''default-slice'\''
                                                       P?
contextAttribute: '\''domain'\''
contextCondition: '\''IS_EQUAL_TO'\''
contextValueRange: '\''api.main.nemo.onelab.eu
contextAttribute: '\''pod cidr'\''
contextCondition: '\''IS EQUAL TO'\''
contextValueRange: '\''10.1.0.0/16'\'
expectationTargets:
targetName: '\''isolatedConnectivity
targetCondition: '\''IS EQUAL TO'\''
targetValueRange: '\''true'\'
expectationId: '\''2'\''
expectationVerb: '\'
                     'DELI
expectationObject:
objectType: '\''K8S
                              NFIG'
objectInstance:
                            hetwork-8'\''
objectContexts
                      name'\''
contextAttribut
                     'IS EQUAL TO'\''
contextCor
                     ''staging-2'\''
           eRai
con
        ttribute:
                 '\''bearer token'\''
CO
       ondition: '\''IS EQUAL TO'\''
conte
context alueRange:
```

'\''eyJhbGciOiJSUzI1NiISImtpZCI6IjZGbm42VVhCWVp4dVhRVEYwSGZTdnJsZEdXX1VsXOlSSGIGZz1FUGpua0kifQ .eyJhdWQiOlsiaHR0cHM6Ly9rdWJlcm5ldGVzLmRlZmF1bHQuc3ZjLmNsdXN0ZXIubG9jYWwiLCJrM3MiXSwiZXhwIjoxN zc3NzIyODUyLCJpYXQiOjE3NDE3MjI4NTISIm1zcyI6Imh0dHBzOi8va3ViZXJuZXRlcy5kZWZhdWx0LnN2Yy5jbHVzdGV yLmxvY2FsIiwianRpIjoiZWI2OTRiMjYtZDg2ZS00NDIyLWJmZTktZWVkYzgwOTg3OGI3Iiwia3ViZXJuZXRlcy5byI6e yJuYW1lc3BhY2UiOiJuZWlvLW5ldCIsInNlcnZpY2VhY2NvdW50Ijp7Im5hbWUiOiJsMnNtLWNvbnRyb2xsZXItbWFuYWd lciIsInVpZCI6IjdiYzIyMTRkLWZmOGEtNDlmYi04YzYxLTczNjE5YmM0MDhiYyJ9fSwibmJmIjoxNzQxNzIyODUyLCJzd WIiOiJzeXN0ZW06c2Vydm1jZWFjY291bnQ6bmVtby1uZXQ6bDJzbS1jb250cm9sbGVyLW1hbmFnZXIifQ.aWq9F6wrwslk nWryEWfkd7lONizjc47WW1-hyn-

WshkyjXKAqUJF86WFD5arbNIyGOglsNop92yN_fSA_FnHCv5F5S7oYXEx0BLlB3NDn3MCICAtklfWFpZ_hHlc2PW1A18Ua 2jvogQyDhx4LFx8jR7xBJatY6JkxXKrhhLBIZCx7VVhIIEz2B0Niy90CpU2rvsN8YCB9CfTnDwimVFtQjZ5eG6TXZr8N6n

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	latform & laborato	esults.	Page:	101 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



```
xE1Uo68pqfoO7z34GJVWWWJdwLlqzS5UAdLmlccx0XZutr9-9R1dFfj6hvMnxKP-LXOYLkKVidtzeqr-
X8IrkAHDuM4IqOY1_c4fdC89Un6L-y0hOg'\'
contextAttribute: '\''api key'\''
contextCondition: '\''IS EQUAL TO'\''
contextValueRange: '\''<u>https://api.s2.nemo.onelab.eu:6443</u>'\''
expectationTargets:
 targetName: '\''isolatedConnectivity'\''
targetCondition: '\''IS_EQUAL_TO'\''
targetValueRange: '\''true'\''
expectationContexts:
                                                              or or
contextAttribute: '\''k8s_l2_network'\''
contextCondition: '\''IS EQUAL TO'\''
contextValueRange: '\''spain-network-8'\''
contextAttribute: '\''namespace'\''
contextCondition: '\''IS_EQUAL_TO'\''
contextValueRange: '\''0937609e-a6c4-43d3-b0dd-2d7c3641ee0c'\''
intentContexts:
contextAttribute: '\''NEMO WORKLOAD'\''
contextCondition: '\''IS_EQUAL_TO'\''
contextValueRange: '\''0937609e-a6c4-43d3-b0dd-2d7c3641ee0c
intentPriority: 1
observationPeriod: 60
intentAdminState: '\''ACTIVATED'\'''
     en chine
```

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version					Page:	102 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Annex 2-LCM subcomponent



Document name:	D4.3 A Final ve	D4.3 Advanced NEMO platform & laboratory testing results. Final version					103 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



```
timestamp=timestamp,
            )
            workload document instance.cluster name = message['targetCluster']
            workload document instance.status
WorkloadDocumentInstance.WorkloadDocumentInstanceStatus.DELETED
            workload document instance.save()
         case _:
            LOG.warning('Got message type: %s', message type)
Pentine Approx
             return
```


 Document name:
 D4.3 Advanced NEMO platform & laboratory testing results. Final version
 Page:
 104 of 111

 Reference:
 D4.3
 Dissemination:
 PU
 Version:
 1.0
 Status:
 Final



Annex 3 – ServiceProviderModel



Document name:	D4.3 A Final ve	dvanced NEMO p ersion	esults.	Page:	105 of 111		
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



```
greenEnergyRewards["0%"] = 0;
  greenEnergyRewards["20%"] = 20000000;
  greenEnergyRewards["40%"] = 40000000;
  greenEnergyRewards["80%"] = 80000000;
  greenEnergyRewards["100%"] = 10000000;
}
modifier checkRegistration(string memory serviceId) {
  require(
    !nemoFunds.isCustomerRegistered(serviceId),
                                                                                1012
     "The customer is already registered!"
  );
}
modifier checkRegionData(string memory region) {
  require(
     nemoTokenEstimationSetup.isRegionSet(region),
     "Data for region must be set before calling this function."
  );
}
function register(
  string memory serviceld
) public checkRegistration(serviceId) {
  string memory _identifier = "ServiceProvider"
  nemoFunds.registerCustomer(serviceId, _identifi
}
                                         nemory serviceId) public {
function removeServiceProviderInfo(string
  require(
     nemoFunds.isCustomerRegistered(serviceId),
                           gistered!"
     "The customer is
                     not r
  );
  nemoF
                     Customer(serviceId);
function enabledIngress(string memory serviceId) public {
  require(
    nemoFunds.isCustomerRegistered(serviceId),
     "The customer is not registered!"
  );
  nemoFunds.enabledIngress(serviceId);
}
```

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						106 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



```
function computeCpuCredits(
  ServiceMetrics memory _metrics,
  uint256 _regionalCpuCosts,
  uint256 _cpuBaseRate
) public pure returns (uint256) {
  // CPU
  uint256 _cpuTokens = 0;
  uint256 _cpuUsage = _metrics.cpuUsage * 10 ** 3;
                                                                          or of
  if (_cpuUsage > _regionalCpuCosts) {
    _cpuTokens =
       ((_cpuUsage) / _metrics.clusterCpuUsage) *
       10 ** 8 +
       (_cpuBaseRate * 10 ** 5);
  }
  return _cpuTokens;
}
                                                   function computeMemoryCredits(
  ServiceMetrics memory _metrics,
  uint256 _regionalMemoryCosts,
  uint256 memoryBaseRate
) public pure returns (uint256) {
  // RAM
  uint256 ramTokens = 0;
  uint256 _ramUsage = _metrics.memoryUsage
                                              10 *
  if (_ramUsage > _regionalMemoryCo
    _ramTokens =
       (_ramUsage / _metrics.clu
                                            age)
                                  Men
       10 ** 8 +
       ( memoryBase
  }
  return ram7dken
}
  nction computeCredits(
   ServiceMetrics memory _metrics
) public checkRegionData(_metrics.region) {
  require(
    nemoFunds.isCustomerRegistered(_metrics.clusterId),
    "The cluster is not registered!"
  );
  require(
    nemoFunds.isCustomerRegistered( metrics.serviceId),
    "The service is not registered!"
```

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version						107 of 111
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



```
);
(
  bool _highDemand,
  uint256 _highDemandCost,
  uint256 regionalCpuCosts,
  uint256 _regionalMemoryCosts
) = nemoTokenEstimationSetup.getRegionInfo(_metrics.region);
                                                                        oron o
(
  uint256 _cpuBaseRate,
  uint256 _memoryBaseRate,
  string memory _greenEnergy
) = infrastructure.getInfrastructureCostsAndGreenEnergy(
    _metrics.clusterId
  );
uint256 _tokens = 0;
uint256 _greenEnergyReward = 0;
// CPU
uint256 _cpuUsage = _metrics.cpuUsage * 10 ** 3;
//RAM
uint256 _ramUsage = _metrics.memoryUsage
                                           10
_tokens =
                                                _cpuBaseRate) +
  computeCpuCredits(_metrics, _regionalCpuCo
  computeMemoryCredits(
    _metrics,
    _regionalMemoryCosts
    _memoryBaseRate
  );
if (_highDemand) {
  _tokens +
             _highDemandCost;
// Check infrastructure green energy percentage and charge less
   eenEnergyReward = greenEnergyRewards[_greenEnergy];
if (_tokens >= _greenEnergyReward) {
  _tokens -= _greenEnergyReward;
}
_tokens = _tokens / 1000;
nemoFunds.makeTransaction(
```

Document name:	D4.3 A Final ve	dvanced NEMO p ersion	Page:	108 of 111			
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final


metrics.serviceId,	
metrics.clusterId,	
tokens	
);	
emit ServiceComputeTokens(
_metrics.serviceId,	
_metrics.clusterId,	•
_cpuUsage,	
_ramUsage,	
_tokens	
);	$\sim 10^{-1}$
}	
}	^ `
	\rightarrow V
1	
X	
*	

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	109 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



Annex 4-Intent for network connectivity request

Intent for network connectivity request. L2S-M network request example.

Intent:	
id: 'mncc_l2sm_1'	
userLabel: 'cloud_continuum'	
intentExpectations:	
- expectationId: '1'	
expectationVerb: 'DELIVER'	\wedge
expectationObject:	
objectType: 'K8S_L2_NETWORK'	
objectInstance: 'spain-network'	
objectContexts:	
- contextAttribute: 'name'	×
contextCondition: 'IS_EQUAL_TO'	
contextValueRange: 'spain-network'	
- contextAttribute: 'providerName'	
contextCondition: 'IS_EQUAL_TO'	
contextValueRange: 'default-slice'	
- contextAttribute: 'domain'	
contextCondition: 'IS_EQUAL_TO'	
contextValueRange: 'api.main.nemo.onelab.eu'	
- contextAttribute: 'pod_cidr'	
contextCondition: 'IS_EQUAL_TO'	
contextValueRange: '10.1.0.0/16'	
expectationTarget:	
- targetAttribute: 'isolatedConnectivity'	
targetCondition: 'IS_EQUAL_TO	
targetValueRange: 'true'	
- expectationId: '2'	
expectationVerb: 'DELDVER	
expectationObject:	
objectType: 'K85_CLUSTER_CONFIG'	
objectInstance: 'spain-fietwork'	
objectContexts:	
- contextAttribute: 'name'	
contextCondition: 'IS_EQUAL_TO'	
context ValueRange: 'staging-1'	
- contextAttribute: 'bearer_token'	
contextCondition: 'IS_EQUAL_TO'	
context valueRange: 'eyJhbGciO	
- contextAttribute: "api_key"	
contextCondition: 'IS_EQUAL_IU'	
context v alueRange: 'https://api.staging1.nemo.onelab.eu:6443'	
expectation I arget:	

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	110 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final



- targetAttribute: 'isolatedConnectivity' targetCondition: 'IS_EQUAL_TO' targetValueRange: 'true' expectationContexts: - contextAttribute: 'k8s 12 network' contextCondition: 'IS_EQUAL_TO' contextValueRange: 'spain-network' - contextAttribute: 'namespace' contextCondition: 'IS EQUAL TO' orov? contextValueRange: 'cbcb208a-d535-434b-bb35-217a64bd516b' - expectationId: '3' expectationVerb: 'DELIVER' expectationObject: objectType: 'K8S_CLUSTER_CONFIG' objectInstance: 'spain-network' objectContexts: - contextAttribute: 'name' contextCondition: 'IS_EQUAL_TO' contextValueRange: 'staging-2' - contextAttribute: 'bearer_token' contextCondition: 'IS_EQUAL_TO' contextValueRange: 'eyJhbGc...' - contextAttribute: 'api_key' contextCondition: 'IS_EQUAL_TO' contextValueRange: 'https://api.s2.nemo.onelab.eu expectationTarget: - targetAttribute: 'isolatedConnectivity' targetCondition: 'IS_EQUAL_TO' targetValueRange: 'true' expectationContexts: - contextAttribute: 'k8s_l2_networ contextCondition: 'IS EQUAL' contextValueRange: spain-ne - contextAttribute: 'namesp contextCondition: 'IS FOUAL TO' contextValueRange: nemo-workload' intentContext tAttributer 'NEMO_WORKLOAD' ntextCondition: 'IS EQUAL TO' xtValueRange: 'cbcb208a-d535-434b-bb35-217a64bd516b' con intentPriority: 1 observationPeriod: 60 intentAdminState: 'ACTIVATED'

Document name:	D4.3 Advanced NEMO platform & laboratory testing results. Final version				Page:	111 of 111	
Reference:	D4.3	Dissemination:	PU	Version:	1.0	Status:	Final