

# NEMO

## Next Generation Meta Operating System

### D2.3 Enhancing NEMO Underlying Technology

Document Identification			
Status	Final	Due Date	30/12/2024
Version	1.0	Submission Date	31/12/2024

Related WP	WP2	Document Reference	D2.3
Related Deliverable(s)	D2.1, D2.2, D3.2	Dissemination Level (*)	PU
Lead Participant	TID	Lead Author	Luis M Contreras
Contributors	TID, TSG, COMS, INTRA, UC3M, WIND3, CMC, AEGIS, SU, UPM, STS, ATOS, SYN	Reviewers	Dimitrios Skias
			Gianluca Rizzi

#### Keywords:

Cybersecure Micro-services' Digital Twins, Cybersecure Federated Deep Reinforcement Learning, Federated Machine Learning, federated meta Network Cluster Controller, Time Sensitive Networking

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
Luis M. Contreras, Guillermo Sánchez, Alejandro Muñiz	TID
Gregor Cerar, Blaž Bertalani, Matija Cankar	COMS
Victor Gabillon	TSG
Spyridon Vantolas, Konstantinos Raftopoulos, Nikolaos Papadakis, Manos Karampinakis	AEGIS
Mika Skarp, Jose Costa-Requena	CMC
Skias Dimitrios	INTRA
G.Spanoudakis, Y.Papaefstathiou	STS
Mohamed Legheraba	SU
Borja Nogales, Iván Vidal, Francisco Valera, Luis F. González	UC3M
Alberto del Rio, Javier Serrano, David Jimenez	UPM
Gianluca Rizzi	WIND3

Document History			
Version	Date	Change editors	Changes
0.1	18/10/2024	TID	First ToC draft and work assignments
0.2	22/10/2024	TID	First Inputs in Executive Summary and section 1
	22/10/2024	COMS, TSG, TID, CMC	Definition of the final structure, first inputs and work assignment in sections 2, 3, 4 and 5.
0.3	5/11/2024	ALL	First iteration
0.4	19/11/2024	ALL	Second iteration
0.5	26/11/2024	ALL	Final contributions in sections 2, 3, 4 and 5.
	2/12/2024	ALL	Final contributions in sections 6 and 7.
0.6	17/12/2024	WIND3	First Review
0.7	20/12/2024	INTRA	Second Review
0.9	23/12/2024	TID, ALL	Review's corrections pending to be accepted
0.91	26/12/2024	TID	FINAL VERSION TO BE SUBMITTED
1.0	30/12/2024	ATOS	Quality review and final version

Quality Control		
Role	Who (Partner short name)	Approval Date
Deliverable leader	Luis M Contreras (TID)	30/12/2024
Quality manager	Rosana Valle Soriano (ATOS)	30/12/2024
Project Coordinator	Enric Pages Montanera (ATOS)	30/12/2024
Technical Manager	Terpsi Velivassaki (SYN)	30/12/2024

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	2 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

# Table of Contents

---

Document Information .....	2
Table of Contents .....	3
List of Tables .....	5
List of Figures .....	6
List of Acronyms .....	7
Executive Summary .....	10
1 Introduction .....	11
1.1 Purpose of the document .....	11
1.2 Relation to other project work .....	11
1.3 Structure of the document .....	11
2 Cybersecure Micro-services Digital Twins (CMDT) .....	12
2.1 Overview .....	12
2.2 Architecture and Approach .....	12
2.3 Internal and External Interfacing .....	15
2.4 Conclusion .....	17
3 Cybersecure Federated Deep Reinforcement Learning (CF-DRL) .....	18
3.1 Overview .....	18
3.2 Architecture and Approach .....	18
3.2.1 CFDRL and Meta Orchestration .....	18
3.3 Conclusion .....	19
4 Federated meta-Network Cluster Controller (mNCC) .....	20
4.1 Overview .....	20
4.2 Architecture and Approach .....	21
4.2.1 Intent-Based System .....	21
4.2.2 Network Metrics Exposure .....	24
4.2.3 Technology Connectivity Adaptors .....	29
4.2.4 Network Performance Monitoring .....	34
4.3 Conclusion .....	35
5 Time Sensitive Networking .....	36
5.1 Overview .....	36
5.2 Architecture and Approach .....	36
5.3 Internal and External Interfacing .....	37
5.4 Conclusion .....	37
6 Proof of Concepts .....	38

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	3 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

6.1	CMDT first stable version.....	38
6.1.1	Detailed description.....	38
6.2	CF-DRL first stable version.....	41
6.3	mNCC hybrid clusters and intent networking.....	43
6.3.1	Single Cluster Connectivity throughout L2SM overlay adaptor.....	43
6.3.2	Intent-Based System.....	46
6.3.3	Multi-Cluster monitoring and exposure.....	52
7	Conclusions.....	55
8	References.....	56
9	Annexes.....	57
9.1	TFS interface and example intent.....	57
9.2	L2SM example intent.....	59

PENDING EC APPROVAL

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	4 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

## List of Tables

---

Table 1: <i>objectContext</i> for 5G adapter .....	21
Table 2: <i>expectationTarget</i> for 5G adapter .....	22
Table 3: <i>targetContext</i> for 5G adapter .....	22
Table 4: <i>expectationContext</i> for 5G adapter.....	22
Table 5: <i>intentContext</i> for all adapters.....	22
Table 6: <i>objectContext</i> for L2SM adapter .....	22
Table 7: <i>objrctContext</i> for L2SM adapter.....	23
Table 8: <i>expectationContext</i> for L2SM adapter .....	23
Table 9. <i>NeMeX</i> input metrics.....	28
Table 10: <i>objectContext</i> for L2VPN.....	32
Table 11: <i>expextationTargets</i> for L2VPN .....	32
Table 12: <i>objectContext</i> for BGP session .....	34
Table 13: <i>expectationContext</i> for BGP session.....	34

PENDING EC APPROVAL

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	5 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

## List of Figures

Figure 1 CMDT architecture and integration points.	13
Figure 2 CMDT data collection and passing to RabbitMQ	14
Figure 3 Current version of CMDT AMPQ message	16
Figure 4 Intent-based API interface that contains PPEF information	18
Figure 5 Result of querying the API (filtering RAM and CPU usage of the workload of interest)	18
Figure 6 Retrieving the information from the Meta orchestrator through Rabbit MQ	19
Figure 7 Retrieving the information from the Meta orchestrator through Rabbit MQ	20
Figure 8 Intent workflow schema	23
Figure 9 NeMeX internal workflow.	27
Figure 10 mNCC metrics data model.	29
Figure 11 Towards integration workflow 5G adapter	31
Figure 12: Private peering interactions	33
Figure 13: TSN Architecture	36
Figure 14: List of resources per workload.	39
Figure 15 The CMDT message example which is transported over RabbitMQ	40
Figure 16 CFDRDL communication architecture for scaling the replicas of the workloads in the Meta-Orchestrator	41
Figure 17: Code snippet for request	42
Figure 18: RabbitMQ connectivity snippet	43
Figure 19: Metric retrieving	43
Figure 20 K8s clusters of this PoC	43
Figure 21 Resources present in the clusters in the PoC	44
Figure 22 API endpoint generation	44
Figure 23 K8s cluster token generation	44
Figure 24 Configuration file used to create a new slice with the L2S-M-md component	45
Figure 25 CLI used to create the slice between the clusters	45
Figure 26 L2S-M resources successfully installed in both clusters	46
Figure 27 IBS importing libraries	46
Figure 28 IBS importing executioners	47
Figure 29 IBS input intent	47
Figure 30 IBS first classification	47
Figure 31: IBS subintent generation.	48
Figure 32 IBS second classification	48
Figure 33 IBS tfs_l2vpn translation	49
Figure 34 IBS "sys_out" executioner	49
Figure 35 IBS "tfs_connector" executioner	49
Figure 36 IBS http request to TFS	50
Figure 37 TFS webui L2VPN service status	50
Figure 38 TFS webui L2VPN service details	50
Figure 39: Intent classification. L2SM case.	51
Figure 40: Final L2SM translation.	51
Figure 41: L2SM intercluster network active.	52
Figure 42 Logs from the network probes in the first cluster.	52
Figure 43 Logs from the network probes in the second cluster.	53
Figure 44 Metrics saved in the internal Prometheus in the first cluster.	53
Figure 45 On the top, the logs of two Nemex components deployed in two clusters. In the bottom, the RabbitMQ updates.	54

## List of Acronyms

Abbreviation / acronym	Description
AAA	Authentication, Authorization, and Accounting
AD	Architecture Description
AF	Application Function
AI	Artificial Intelligence
AMF	Access and Mobility Management Function
AMPQ	Advanced Message Queuing Protocol
API	Application Programming Interface
CFDRL	Cybersecure Federated Deep Reinforcement Learning
CI/CD	Continuous Integration & Continuous Delivery/Continuous Deployment
CMDT	Cybersecure Micro services Digital Twin
CN	Core Network
CNC	Centralised Network Configuration element
CNCF	Cloud-Native Computing Foundation
CNI	Container Network Interface
CPU	Central Processing Unit
CR	Custom Resource
DAST	Dynamic Application Security Testing
DLT	Distributed Ledger Technology
DQN	Deep Q-Network
DS TT	Device Side TSN translator
Dx.y	Deliverable number y belonging to WP x
E2E	End-to-End
EC	European Commission
GDPR	General Data Protection Regulation
gPTP	generic Precision Time Protocol
gRPC	Google Remote Procedure Call
HLA	High-Level Architecture
HW	Hardware
IAM	Identity and Access Management
IAST	Interactive Application Security Testing
IBMC	Intent-Based Migration Controller
IBS	Intent-Based System
IDCO	Inter-Cluster Connectivity Orchestrator
IMC	Intent-based Migration Controller
IoT	Internet of Things
K8s	Kubernetes
KPI	Key Performance Indicator

LCM	Life Cycle Management
LCM UI	Life Cycle Management User Interface
LLDP	Link Layer Discovery Protocol
MANO	Management and Orchestration
MEC	Multi-access Edge Computing
meta-OS	Meta-Operating System
ML	Machine Learning
mNCC	meta–Network Cluster Controller
mRA	meta-Reference Architecture
MO	meta-Orchestrator
MOCA	Monetization and Consensus-based Accountability
NAC	NEMO Access Control
NED	Network Edge Device
NFV	Network Function Virtualization
NGIoT	Next-Generation IoT
NW TT	Network TSN Translator
OS	Operating System
PA-LCM	Plugin & Apps Lifecycle Manager
PCF	Policy Control Function
PoC	Proof of Concept
PPEF	PRESS and Policies Framework
PRESS	Privacy, data pRotection, Ethics, Security & Societal
PSFP	Per-stream filtering and policing
RAN	Radio Access Network
RASP	Run-time Application Security Protection
REST	Representational State Transfer
RBAC	Role-Based Access Control
RL	Reinforcement Learning
SDK	Software Development Kit
SDN	Software Defined Networking
SDO	Standard Development Organization
SD-WAN	Software-Defined Wide Area Network
SLA	Service Level Agreement
SLO	Service Level Objective
SMF	Session Management Function
TEE	Trusted Execution Environment
TSCAI	Time Sensitive Communication Assistance Information
UPF	User Plane Function
UUID	Universally Unique Identifier
V&V	Validation & Verification
VM	Virtual Machine
VNF	Virtual Network Function



WP	Work Package
YAML	Yet Another Markup Language

PENDING EC APPROVAL

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	9 of 60				
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

## Executive Summary

---

The main goal for this document is to outline the work conducted in the three tasks of WP2, emphasizing the components that form the underlay layer of the meta-OS of NEMO (i.e., three components plus the TSN capabilities). In here it is provided a thorough description, focusing on overcoming the initial challenges related to the development of these components as they begin to coexist within the same ecosystem. Through this discussion, the reader will gain a clear understanding of the advancements proposed by NEMO partners relative to the deployment and advances for the technologies utilized.

The document highlights the achievements made thus far during the Enhance phase, with particular attention to the overall deployments and internal integrations employed to develop and report the Proof of Concepts (PoC) for each of the main subcomponent. This phase finalizes with three of the four components (CMDT, mNCC and TSN) in an integration stage, with some functional and non-functional improvements from the previous version, and with the final stable version of the CF-DRL.

To completely understand this work it is important to have had a previous reading of the *D2.2 Enhancing NEMO Underlying Technology* [1] as this document, due to the fact that is an extension of it and it works as a continuation for the initial integration done in the previous phase.

Once finalized, the reader should end with a big-picture knowledge of the four main components depicted in here, being able to understand the motivation and functionality of each of these components and have a clear view about what to expect for the final developing phase.

The main results for this document included:

- A final approach to a stable version for the Cybersecure Micro-Services Digital Twin, with an overlook about how it is deployed and how it works.
- A final stable version of the Cybersecure Federated Deep Reinforcement Learning with a comprehension of the internal architecture and a description of each of the algorithms available to be run in the first stable version.
- A final approach to a stable version for the Federated meta-Network Cluster Controller, with an initial deployment of the service provisioning and the network monitoring and exposure capabilities.
- A final approach to a stable version of the Time Sensitive Network component for 5G networks.

As conclusion, this document depicts the main advances in the deployment of the NEMO technological enables, providing a deep understanding of how each of them works and the updated contributions already done in a concluded development and deployment phase. During the elaboration of this document some initial approaches proposed in the previous stage have been enhanced and some others have been dropped as there were considered as unnecessary or redundant. Also, the presence of external difficulties has made CMDT and mNCC suffer changes to adapt to the new consortium status, impacting these difficulties in the completeness of these components but not in a critical way.

In the current phase, it is expected to develop the final deployment and validation for each of the four components, and an improvement on the integration inside the full NEMO meta-OS ecosystem.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	10 of 60				
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

# 1 Introduction

---

## 1.1 Purpose of the document

---

This document works as an update from *D2.2 Enhancing NEMO Underlying Technology* [1], providing an updated view of each of the three components defined in WP2. The main goal is to describe the latest version of each component and its parts, providing a complete knowledge about each of them and how is the final stable version. Additionally, it is also depicted how these components interact with other NEMO components, focusing on the functionalities provided and the interfaces exposed, being this crucial for the overall NEMO stack integration steps. Moreover, the document will serve to evaluate the latest status of the components against the NEMO KPIs. The three components described (each of them carried out in one of the Task within the Work Package) are going to focus on topics as cloud and edge computing, networking, AI operations, and microservices-oriented architectures, being important to have a background on these technologies.

## 1.2 Relation to other project work

---

*D2.3 Enhancing NEMO Underlying Technology* works as an update to *D2.2 Enhancing NEMO Underlying Technology*, providing advancements on the technologies defined in there and giving the latest view of its implementation. D2.3 will follow the same structure and objectives as D3.2 and D4.2 providing a jointed view of NEMO architecture and implementation and working as a description of the integration work and the evaluation of Use Cases described in D5.2 NEMO Living Labs use cases evaluation results at Milestones 8 and 11.

## 1.3 Structure of the document

---

This document starts with an introduction and explanation of the context and goals defined to be addressed here. Then, there will be a section for each of the three WP2 components, and a last one for the Time Sensitive Networking, due to its characteristics. In each of these sections dedicated to the components it is provided an overview of what is being pursued in each of them, the final improved architecture, how each of these components works and the interfaces to communicate with the rest of NEMO components.

Once components are depicted in the theoretical way, there is also a section 6 for the explanation and validation of each of them, showing proofs of concept that evaluate the functionalities and bring a view on how the integration with the rest of NEMO WP2 architecture it's done. For the pending integrations with the rest of the NEMO components, proper descriptions and reporting will be included in the ongoing D4.3.

Finally, all the information expressed in sections 2 to 6 will be summarized in the Conclusions section, focusing on the results obtained related to the goals and challenges expected for the technologies depicted. Also, as this is the final document for the work package, this conclusion section will provide a big picture about the work realized and the benefits that these technologies bring not only to the rest of the NEMO architecture, also to the general State-of-the-Art.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	11 of 60				
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

## 2 Cybersecure Micro-services Digital Twins (CMDT)

### 2.1 Overview

The Deliverable D2.2 [1] describes all potential expected functionalities of the Cybersecure Micro-services Digital Twins (CMDT) component, which plays a pivotal role in monitoring and assessing the status and availability of each workload defined as a Digital Twin (DT). This service is responsible for aggregating data from all available clusters, organizing and enriching it with crucial metadata—such as service ownership details—and providing consolidated, actionable insights to other services.

The role and functionalities of the CMDT have improved the previous CMDT version presented in D2.2. The main changes are in the distinction of the components' goal inside the NEMO MetaOS framework, removing unnecessary functionalities and clarifying the integration interfaces of the tool. With these changes, CMDT became near real-time data fusion and distribution service, capturing and DT life-cycle anomalies and passing them to the user interface. This means that the CMDT dashboards will remain, but will be used only for NEMO super-users, while the DT information, crucial for the end user is integrated in the NEMO LCM UI.

Additionally, the CMDT integration is finally detailed, including the definition of course of the data that is passed over the AMPQ message bus (e.g., RabbitMQ). This information is crucial to transmit the status from the distributed NEMO environment to the LCM UI in a near real-time manner. The messages include the following data:

- **Timestamp:** Depicts the exact time when the data was prepared and message sent, ensuring chronological order.
- **Workload ID:** Represents a unique identifier for the workload (e.g. Digital Twin).
- **UserID:** Represents a unique identifier for the user.
- **LinkerD retrieved info** such as network traffic to/from pod, traffic rate, and response stats.
- **Kube-state-metrics info** such as memory, and CPU utilization,
- **Number of replicas:** Represents a current number of replica instances (k8s pods) running for a workload. It indicates the number of instances or replica instances.
- **Status:** Describes the current and past state of the workload or k8s pods, which can be:
  - Running
  - Restart
  - Ready
  - Failure

The message details currently cover all KPIs defined in section 2.2 of previous deliverable “D2.2 Enhancing NEMO Underlying Technology”. The list of message data attributes is not fixed and could be updated in the future due to the stakeholder requests for better user experience. In the following subsections we point out the specific changes made on CMDT architecture in the last development and integration period.

### 2.2 Architecture and Approach

The initial iteration of the CMDT implementation was ambitious, incorporating numerous data sources—some of which were redundant—to gather as much information as possible and gain a comprehensive understanding of the system and its Digital Twin deployments. Following the initial deployment, we evaluated the collected data and compared the CMDT’s capabilities with other NEMO services. This led us to reshape the architecture to better align with the NEMO meta-OS framework, ensuring it complemented the functionality of other services effectively.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	12 of 60				
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

The new version of CMDT still provides a comprehensive approach for in detail management of DT, however, the architecture is restructured in a way to guarantee seamless and real-time provision of DT status to the LCM graphical user interface (LCM UI), organized by each workload or user. In essence, this version still includes the same core functionality, without the parts as Grafana and Prometheus, which were used only to showcase the CMDT results in the past. In this final version of CMDT, which is fully integrated into the NEMO operating system, the interfaces mentioned became redundant and were removed from the core CMDT architecture.

The revised CMDT architecture is presented in Figure 1. The core services consist of **k8s Agent aggregation engine, Service Mesh Network Aggregation Engine, CMDT data fusion and broker.** Mentioned services and their processes streamline the DT status message delivery from the infrastructure and software layer to the user space in LCM UI. Both aggregation engines rely on the data gathered through special agents locally collecting the data on each pod. The reader can find these agents shown on the bottom of Figure 1, where each cluster is monitored by Service mesh monitoring and Pod monitor. The gathered and CMDT processed data is shared on the AMPQ service (RabbitMQ). The data is consumed by NEMO services, like the LCM – UI, which uses the data to represent and status of the Digital twin health.

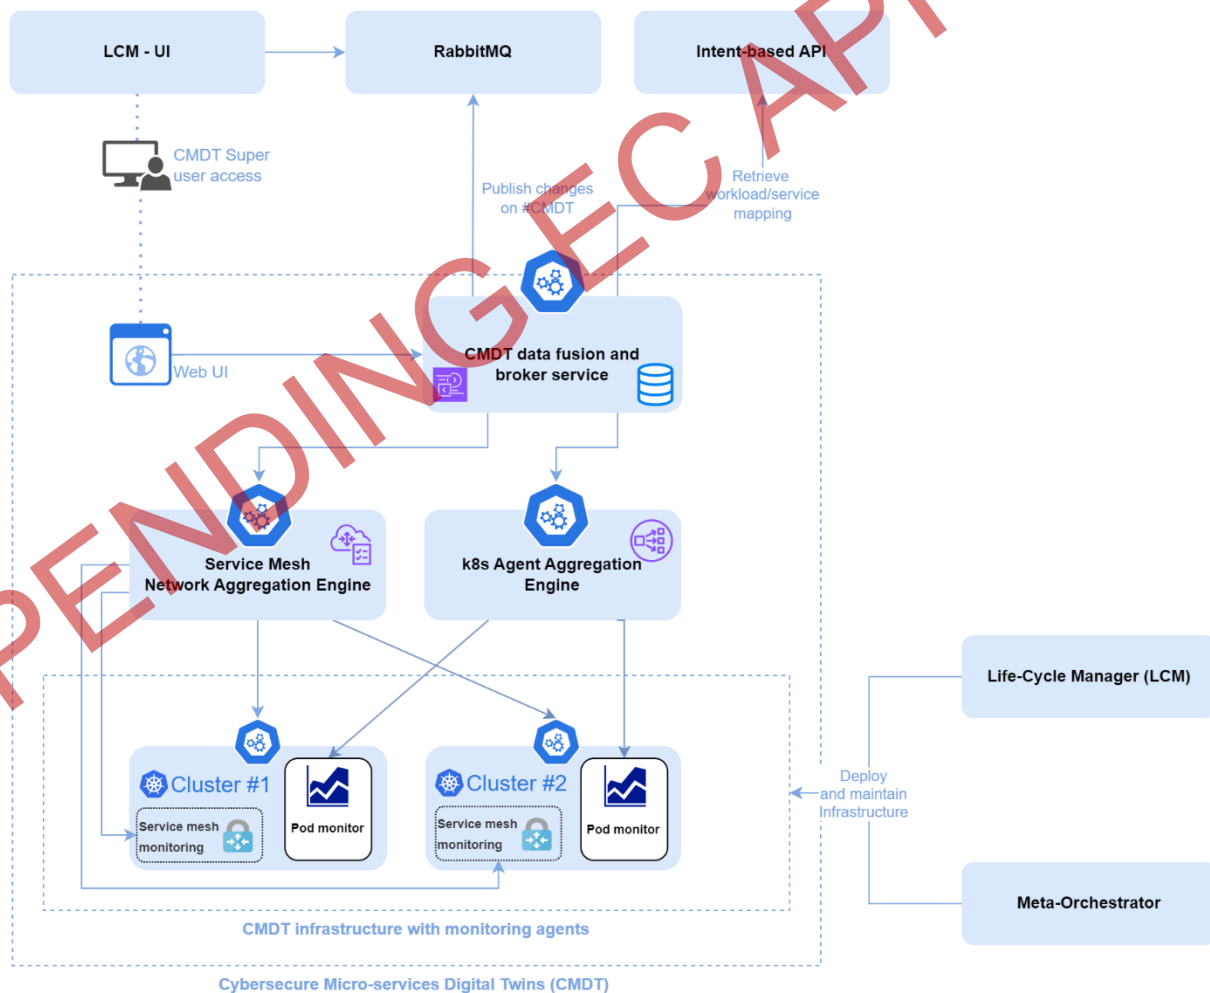


Figure 1 CMDT architecture and integration points.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	13 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
<b>Version:</b>	1.0	<b>Status:</b>	Final

The sequence diagram from Figure 2 outlines the process of CMDT near real-time notifications for the LCM-UI within a distributed infrastructure. The LCM-UI monitors the #CMDT channel via RabbitMQ to receive updates. CMDT Core Services interact with multiple clusters (Cluster 1 to Cluster N) distributed across the infrastructure, gathering network mesh and service data from pods. This DT data is processed, aligned, annotated, and relayed to RabbitMQ as cluster status updates. RabbitMQ acts as a mediator, ensuring that LCM-UI remains informed of the latest statuses from each cluster, enabling near real-time notifications and streamlined monitoring capabilities for the system.

CMDT near real-time notifications for LCM-UI

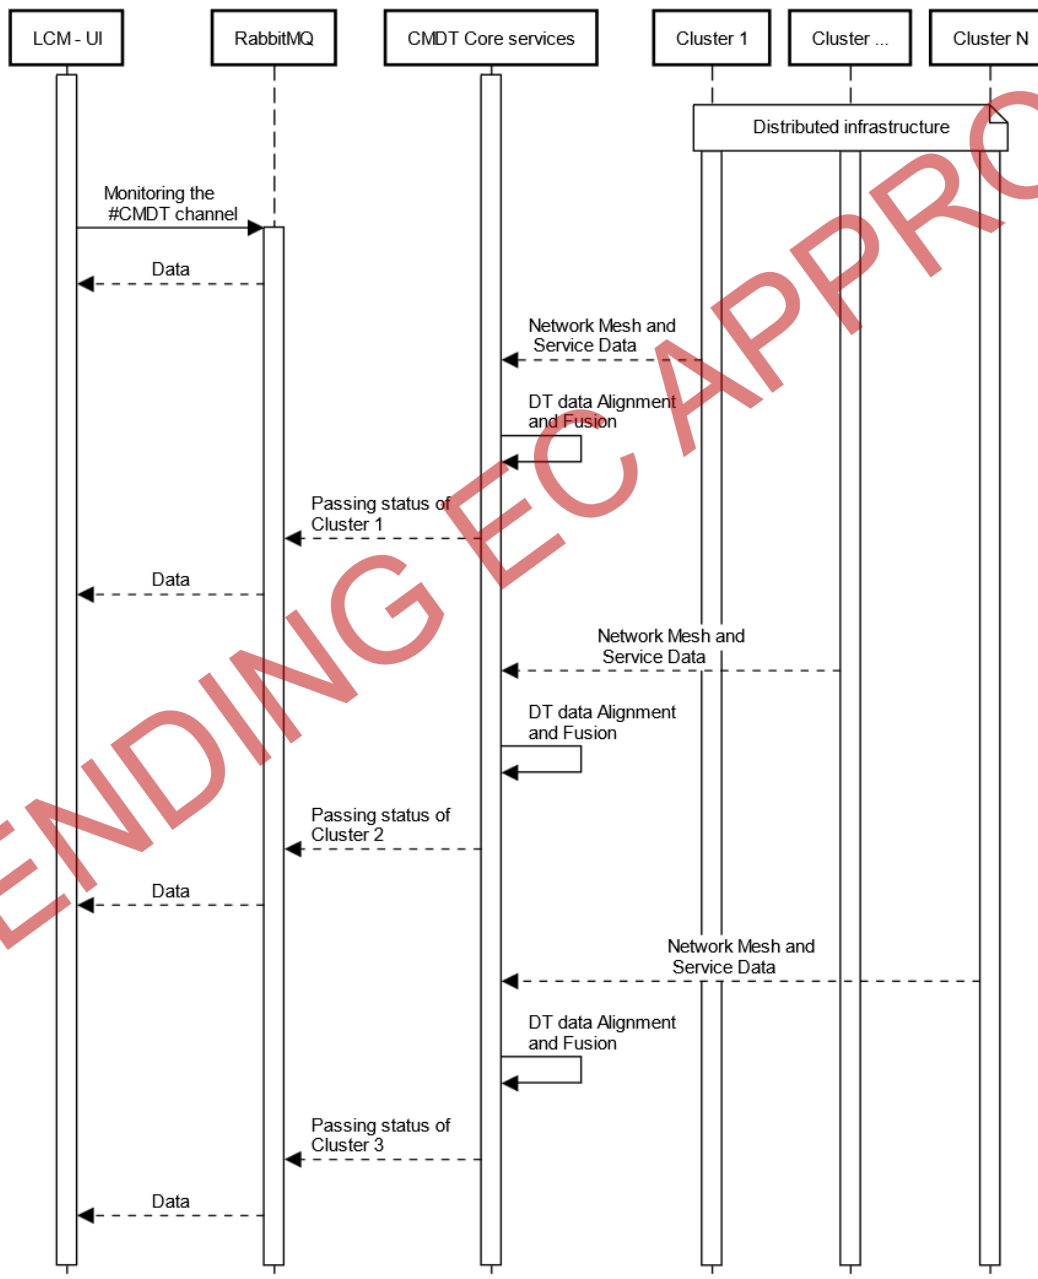


Figure 2 CMDT data collection and passing to RabbitMQ

## 2.3 Internal and External Interfacing

CMDT uses only internal interfaces to communicate with the rest of the NEMO components. The input and output interfaces include the following information:

- **INPUT:**
  - Internal tools:
    - Linkerd (<https://linkerd.io/>) API to gather internal and external network traffic stats.
    - Kube-state-metrics (<https://github.com/kubernetes/kube-state-metrics>) API to gather Kubernetes nodes' stats (CPU, memory, storage utilization).
    - Prometheus (<https://prometheus.io/>) API to access internal storage to query retrospective stats.
    - Kubernetes API to gather services' and pods' state and logs.
  - Intent-based API to access mapping between pods and workload UUID.
  - User/workload relations.
  
- **OUTPUT:**
  - **AMPQ messages:** RabbitMQ periodic and event-driven (on events/changes) messages to be consumed by other services, such as LCM and CF-DRL.
  - **Optional GUI dashboard for NEMO.**

Message destination	Message format
<AMPQ url on RabbitMQ>	Payload (Screenshot of JSON)
Channel: #CMDT	The message format:



```

1  {
2    "pod_name": "app-bot-555c556d76-cdvh9",
3    "timestamp": "2024-12-11T08:32:13.956429Z",
4    "workload_id": "455ccf73-6d05-4f30-b05e-6a43dc211181",
5    "status": [
6      {
7        "status_type": "Ready",
8        "status": "True",
9        "last_transition_time": "2024-12-11T08:01:55Z",
10       "message": null
11      },
12      {
13        "status_type": "ContainersReady",
14        "status": "True",
15        "last_transition_time": "2024-12-11T08:01:55Z",
16        "message": null
17      },
18      {
19        "status_type": "PodReadyToStartContainers",
20        "status": "True",
21        "last_transition_time": "2024-12-11T08:00:10Z",
22        "message": null
23      },
24      {
25        "status_type": "Initialized",
26        "status": "True",
27        "last_transition_time": "2024-11-15T12:24:47Z",
28        "message": null
29      },
30    ],
31    "status_type": "PodScheduled",
32    "status": "True",
33    "last_transition_time": "2024-11-15T12:24:46Z",
34    "message": null
35  },
36  ],
37  "labels": {
38    "app": "app-bot",
39    "linkerd.io/control-plane-ns": "linkerd",
40    "linkerd.io/proxy-deployment": "app-bot",
41    "linkerd.io/workload-ns": "demoapp",
42    "nemo.eu/workload": "455ccf73-6d05-4f30-b05e-6a43dc211181",
43    "pod-template-hash": "555c556d76",
44    "version": "v11"
45  },
46  "traffic_stats": {
47    "req_rate": 0.32079628767407103,
48    "res_rate": 2.1029978858633545,
49    "res_rate_by_code": {
50      "500": 0.17153690382571854,
51      "200": 1.9314609820376358
52    },
53    "res_time_quantiles_ms": {
54      "p99": 18.702803738317762,
55      "p95": 11.644859813084098,
56      "p75": 7.683397683397683,
57      "p50": 4.876404494382022
58    }
59  }
60 }

```

Figure 3 Current version of CMDT AMPQ message

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	16 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
		<b>Version:</b>	1.0
		<b>Status:</b>	Final



## 2.4 Conclusion

---

The CMDT has significantly changed in the last development period. The key improvement was the definition of the essential content gathered and annotated by CMDT and the most efficient and useful integration point through AMPQ.

Current simulations show that the LCM-UI will be able to gain responsiveness and near real-time digital twin visualization by using the CMDT notifications. For advanced NEMO administration users, the CMDT will still maintain the backdoor access for easier status debugging.

PENDING EC APPROVAL

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	17 of 60				
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

# 3 Cybersecure Federated Deep Reinforcement Learning (CF-DRL)

## 3.1 Overview

One of the main applications of CF-DRL within NEMO is to help orchestrate microservices in the Meta-Orchestrator as described in D2.2

## 3.2 Architecture and Approach

### 3.2.1 CFDRL and Meta Orchestration

- **Designing a Deep Q-Network (DQN) Algorithm.** Because the Thales library for reinforcement learning is closed source, it was not possible to share the library as is on the One lab cluster. However, it has been recoded the specific DQN algorithm for the CFDRL-NEMO application.
- **Connection to PPEF:** The PPEF can provide information about the status of the workload that concern the workload’s intents’ expectations set by the NEMO user. CFDRL needs to query the intent-based API automatically using a Keycloak identifier. The CFDRL queries the API at *intent-api.nemo.onelab.eu/api*. Then the information is retrieved and filtered in order to extract the relevant information. In the case of the PPEF, the relevant information<sup>1</sup> is the usage of CPU and RAM and whether it exceeds the limits set for a normal execution.

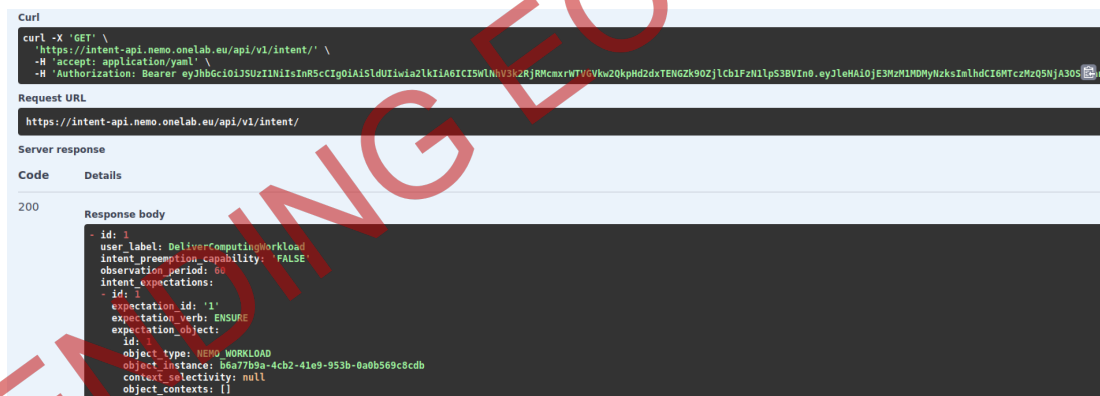


Figure 4 Intent-based API interface that contains PPEF information





Figure 5 Result of querying the API (filtering RAM and CPU usage of the workload of interest)

<sup>1</sup> The information from the PPEF is not only limited to the CPU and RAM, but intents also include other details such as compute expectations (CPU, RAM), energy expectations, security, etc.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	18 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

It's possible to follow over time the variations of the values of the CPU usage and RAM usage with respect to the maximum target values in the following figures. One thing to notice is that the usage values in blue here do not go over the limit target value in orange.



- Connection with the MO:** The connection with the meta-orchestrator is through RabbitMQ. On the CFDRl side the connection is made in python using the pika library. A queue was set in the RabbitMQ so that CFDRl can query it for general info about static cluster resource/specification information (such as its CPU, RAM, storage, etc), as illustrated in Figure 6 below. CFDRl obtains information about the cluster id, capacity such as CPU and memory.

```

NAME                                READY   STATUS    RESTARTS   AGE
cfdr1-rabbit-deployment9-867768685-tjxdb  1/1     Running   0           68d
(base) victor@cpuclin:~$ kubectl logs cfdr1-rabbit-deployment9-867768685-tjxdb
hello bitcoiner to rab
Starting Consuming
received new message b[{"response":{"name":"k3s-cluster","uid":"1a058930-8dd4-4986-80c4-04c0f9831cbb","creationTimestamp":"2024-10-03T15:11:38Z","hubAcceptsClient":true,"capacity":{"cpu":"12","ephemeral-storage":"2029975572Kl","memory":"23983932Kl","pods":"330"},"allocatable":{"cpu":"12","ephemeral-storage":"1974708234894","memory":"23983932Kl","pods":"330"},"version":"v1.30.5+k3s1"},"name":"rasptest","uid":"26ab1655-f353-4874-9c02-7c4d6d694dde","creationTimestamp":"2024-09-10T08:08:03Z","hubAcceptsClient":true,"capacity":{"cpu":"4","ephemeral-storage":"29485028Kl","memory":"4143584Kl","pods":"110"},"allocatable":{"cpu":"4","ephemeral-storage":"28683035216","memory":"4143584Kl","pods":"110"},"version":"v1.30.4+k3s1"},"status":"ok"}]
received new message b[{"response":{"name":"k3s-cluster","uid":"1a058930-8dd4-4986-80c4-04c0f9831cbb","creationTimestamp":"2024-10-03T15:11:38Z","hubAcceptsClient":true,"capacity":{"cpu":"12","ephemeral-storage":"2029975572Kl","memory":"23983932Kl","pods":"330"},"allocatable":{"cpu":"12","ephemeral-storage":"1974708234894","memory":"23983932Kl","pods":"330"},"version":"v1.30.5+k3s1"},"name":"rasptest","uid":"26ab1655-f353-4874-9c02-7c4d6d694dde","creationTimestamp":"2024-09-10T08:08:03Z","hubAcceptsClient":true,"capacity":{"cpu":"4","ephemeral-storage":"29485028Kl","memory":"4143584Kl","pods":"110"},"allocatable":{"cpu":"4","ephemeral-storage":"28683035216","memory":"4143584Kl","pods":"110"},"version":"v1.30.4+k3s1"},"status":"ok"}]
received new message b[{"response":{"name":"k3s-cluster","uid":"1a058930-8dd4-4986-80c4-04c0f9831cbb","creationTimestamp":"2024-10-03T15:11:38Z","hubAcceptsClient":true,"capacity":{"cpu":"12","ephemeral-storage":"2029975572Kl","memory":"23983932Kl","pods":"330"},"allocatable":{"cpu":"12","ephemeral-storage":"1974708234894","memory":"23983932Kl","pods":"330"},"version":"v1.30.5+k3s1"},"name":"rasptest","uid":"26ab1655-f353-4874-9c02-7c4d6d694dde","creationTimestamp":"2024-09-10T08:08:03Z","hubAcceptsClient":true,"capacity":{"cpu":"4","ephemeral-storage":"29485028Kl","memory":"4143584Kl","pods":"110"},"allocatable":{"cpu":"4","ephemeral-storage":"28683035216","memory":"4143584Kl","pods":"110"},"version":"v1.30.4+k3s1"},"status":"ok"}]

```

Figure 6 Retrieving the information from the Meta orchestrator through Rabbit MQ

### 3.3 Conclusion

The CFDRl was uploaded to the Onelab cluster and connected to the Meta Orchestrator and the PPEF so that the RL model could learn from interacting with the monitoring data.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	19 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

# 4 Federated meta-Network Cluster Controller (mNCC)

## 4.1 Overview

The mNCC module is tasked with managing network connectivity and the assessment and exportation of network characteristics within the NEMO system. It serves as an intermediary between NEMO and the underlying physical network topology, providing connectivity and abstracting network services from the deployment specifics. The mNCC enables real-time access to network characteristics, allowing services to autonomously update their network perspective.

Recent updates to the mNCC have significantly enhanced its capabilities and integration within the NEMO ecosystem. The first stable version has been successfully integrated, marking a major milestone in the module's development. The intent-based library has been expanded with additional service libraries, broadening the range of network functionalities that can be expressed through high-level intents. Furthermore, the integration of expected technology adaptors has been completed, enabling the mNCC to interface with a wider array of network technologies seamlessly.

Moreover, the integration of the metric reporting and exposure capabilities has been completed, allowing for more comprehensive network monitoring and analysis. This update enables the mNCC to provide more accurate and detailed network characteristics to other NEMO components. The achievement of automated deployment, supported by the MO and the CI/CD frameworks in NEMO, represents a significant step forward in operational efficiency, streamlining the process of implementing and scaling the mNCC across various network environments.

The interaction of the different modules of the component is represented in Figure 7.

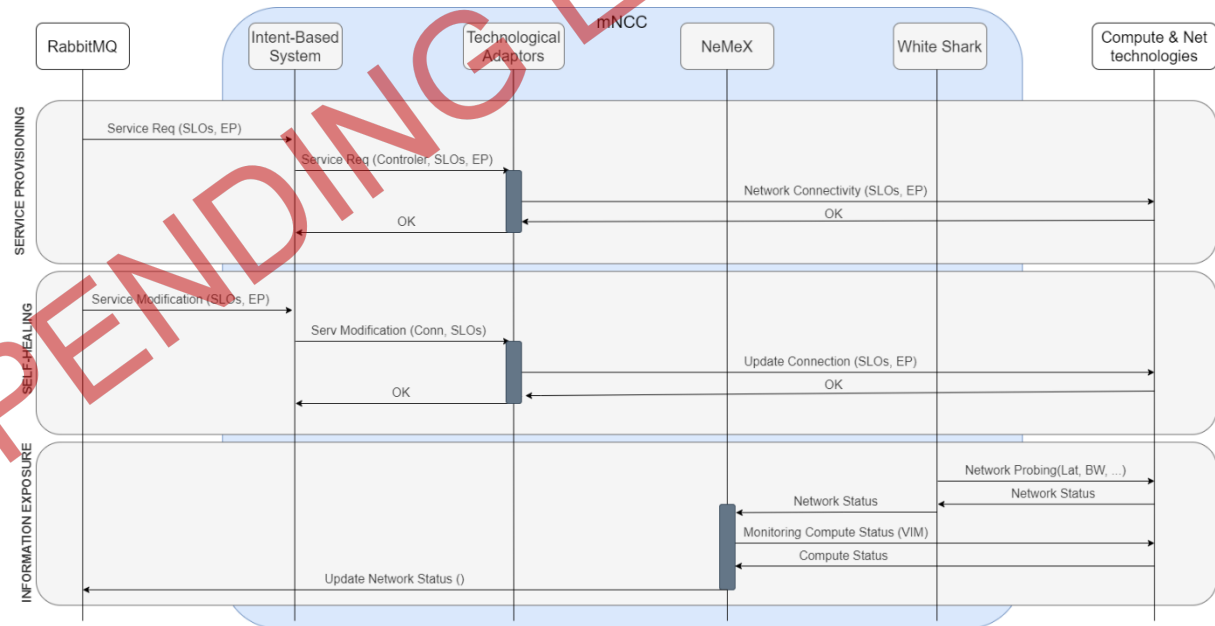


Figure 7 Retrieving the information from the Meta orchestrator through Rabbit MQ

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	20 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
		<b>Version:</b>	1.0
		<b>Status:</b>	Final

## 4.2 Architecture and Approach

In this section, are described the components of the mNCC. These include the Intent-Based System (IBS), which translates network requests into specific technology implementations, and the Network Metrics Exposure (NeMeX), which integrates and exports network topology and performance metrics. The mNCC also incorporates Technology Adaptors, such as the 5G Adapter and Teraflow SDN Controller, which manage specific network technologies. The Connectivity Controller, based on L2SM technology, provides secure link-layer connectivity between microservices within Kubernetes clusters.

### 4.2.1 Intent-Based System

#### 4.2.1.1 Introduction

The Intent Based System (IBS) is a translator between different technologies, abstracting the specific details of network components from the general directives of higher-level components. This abstract order is known as Intent (an intention). It expresses an expectation of what the underlying technologies should do without knowing “the how”. This concept adapts nicely with the paradigm of the mNCC to be able to scale with new Network Adaptors. Just with the addition of a new intent library, the IBS can translate an intention to a new network technology.

In the following section, there is a general overview of recent updates and the current status of the software. Additionally, how the Meta-Orchestrator (MO) utilizes intents to request different types of networks and the work in progress of the integration will be explored. This insight into the system's evolution and functionality will demonstrate how the Intent Based System continues to enhance network management and orchestration capabilities.

#### 4.2.1.2 Intent structure

In this section, the different targets, context and expectations that will trigger the different network adaptors in the mNCC will be described. Notice that the intent structure has not received any update from the version reported in [D2.2], the structure defined the 3GPP specification [3gpp 28.321]. For the sake of simplicity, not all the objects will be described. Also, the targets, contexts and expectations will be divided in tables, although the intent requests must have the complete structure defined in the Gitlab repository<sup>2</sup>.

Table 1: objectContext for 5G adapter

objectContexts	objectType	contextAttribute	contextCondition	contextValueRange
5G_SLICE_FLOW		ip4Address	IS_EQUAL_TO	ip string
		type	IS_EQUAL_TO	'BOTH'
		portNumber	IS_EQUAL_TO	port string
		portType	IS_EQUAL_TO	'UDP' / 'TCP'

<sup>2</sup>Gitlab: <https://gitlab.eclipse.org/eclipse-research-labs/nemo-project/nemo-infrastructure-management/federated-meta-network-cluster-controller/intent-based-system>

Table 2: expectationTarget for 5G adapter

expectationTargets	objectType	targetName	targetCondition	targetValueRange
	5G_SLICE_FLOW	ulCapacity	IS_EQUAL_TO	ip string
		dlCapacity	IS_EQUAL_TO	'BOTH'

Table 3: targetContext for 5G adapter

targetContexts	targetName	contextAttribute	contextCondition	contextValueRange
	dlCapacity/ulCapacity	profile	IS_EQUAL_TO	profile string (audio4k)

Table 4: expectationContext for 5G adapter

expectatinContexts	objectType	contextAttribute	contextCondition	contextValueRange
	5G_SLICE_FLOW	startTime	IS_EQUAL_TO	RFC 3339 – Date/time format
		stopTime	IS_EQUAL_TO	RFC 3339 – Date/time format
		url	IS_EQUAL_TO	url string

Table 5: intentContext for all adapters

intentContexts	userLabel	contextAttribute	contextCondition	contextValueRange
	Cloud_continuum (NEMO)	NEMO_WORKLOAD	IS_EQUAL_TO	NEMO uuid

Table 6: objectContext for L2SM adapter

objectContexts	objectType	contextAttribute	contextCondition	contextValueRange
	K8S_L2_NETWORK	name	IS_EQUAL_TO	network string
		providerName	IS_EQUAL_TO	provider string
		uc3m	IS_EQUAL_TO	domain string

Table 7: objrectContext for L2SM adapter

objectContexts	objectType	contextAttribute	contextCondition	contextValueRange
objectContexts	K8S_CLUSTER_CONFIG	name	IS_EQUAL_TO	network string
		beare_token	IS_EQUAL_TO	token string
		api_key	IS_EQUAL_TO	key string

Table 8: expectationContext for L2SM adapter

expectationContexts	objectType	contextAttribute	contextCondition	contextValueRange
expectationContexts	K8S_CLUSTER_CONFIG	k8s_l2_network	IS_EQUAL_TO	network string
		url	IS_EQUAL_TO	url string

The complete set of possible intents and examples continue in the annexes 9.1 and 9.2.

#### 4.2.1.3 Intent lifecycle

The intent lifecycle starts with the request of the Meta-Orchestrator via the RabbitMQ queue. Once the IBS consumes the intent, it starts the classification phase. In this phase, the classifier module of the IBS compares the key attributes of the intent and matches them with one of the Intent libraries installed [D2.2]. The classification has been updated with a better subdivision of incoming intents into sub-intents. Also, the IBS can process several intents in the same file separated with the usual "yaml" separator "---". This enables in the MO more flexibility and scalability in the requests. In the current version, the IBS can classify an intent in the cloud continuum library and then sub divide this intent to be processed by one of the libraries capable of translating the intent into the specific technological network adaptor.

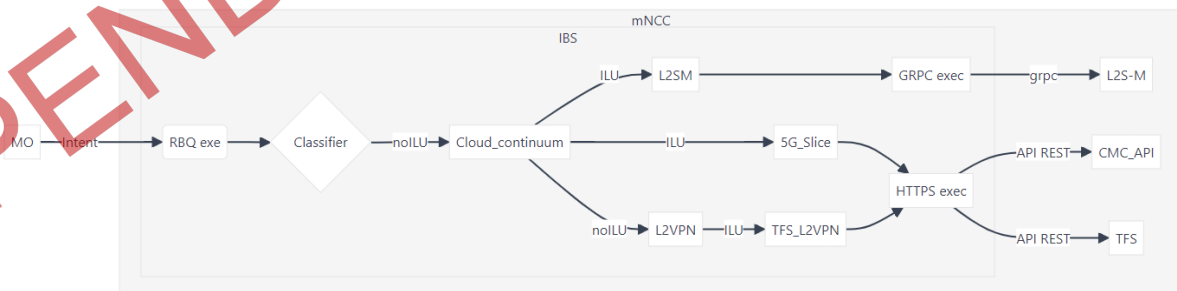


Figure 8 Intent workflow schema

Then, once the intent is translated by the corresponding library, the order is sent to the corresponding executioners. As they are programmed to be interfaces towards the different technologies, if the interfaces are the same, the executioners can be reused, giving the IBS flexibility and independence between the different steps in the lifecycle.



#### 4.2.1.4 Intent model in NEMO cloud continuum paradigm

The implementation of the Intent Based System (IBS) as a micro-kernel architecture significantly enhances the development and integration of new network adaptors. This modular design separates core functionality from specific network adaptor implementations, allowing for seamless integration of new technologies without modifying the central IBS components. As new network technologies emerge, corresponding adaptors can be developed and plugged into the system with minimal disruption to existing operations, ensuring scalability and future proofing. The micro-kernel approach also simplifies maintenance and updates of individual components, including network adaptors. Furthermore, it offers organizations the flexibility to develop or modify network adaptors to suit their specific needs without altering the core IBS functionality. This architectural choice ultimately ensures that the IBS remains highly adaptable, capable of evolving alongside the rapidly changing landscape of network technologies in the cloud continuum paradigm.

#### 4.2.2 Network Metrics Exposure

As defined in previous deliverables, the **Network Metrics Exposure (NeMeX)** is the component serving as the North-Bound Interface (NBI) for mNCC. It interacts with the internal subcomponents of mNCC and the RabbitMQ queue to export network topology and performance metrics for the managed technological domains.

NeMeX has two main functions:

1. **Topology view integration:** it gathers information from multiple sources and combines it into a unified, coherent topology view.
2. **Data modeling and performance metrics delivery:** it uses RabbitMQ to share information with other NEMO components, ensuring scalability and interoperability for the integration of technological controllers or add-ons in a standardized manner.

In its initial deployment, NeMeX establishes a connection with NEMO's RabbitMQ broker. While this version is manually configured, future versions aim to automate this process to enhance scalability, availability, and robustness. Once connected, NeMeX maps multi-domain resources to create an integrated view that combines performance values with domain-specific resources.

The process transforms three initial views into a unified topology by assigning common identifiers to the same nodes. Every  $D$  seconds, updated values are sent to the designated RabbitMQ queue if changes occur.

##### 4.2.2.1 Technological dependencies and requisites

The NeMeX sub-component is based on the next technologies:

- **Python3<sup>3</sup>**. This is the programming language used for almost all code in this component. Python 3 is a high-level interpreted language designed to facilitate code usability, thereby enhancing the potential for code improvements and extensions. Currently version Python3.12 is used.
- **Kubernetes (K8S)<sup>4</sup>**. is used for managing distributed systems, specifically for monitoring computational capacities required for virtualizing network resources. This choice is based on the strong presence of Kubernetes in current systems and its compatibility with other

<sup>3</sup> Python3 official release: <https://www.python.org/download/releases/3.0/>

<sup>4</sup> Official K8S documentation: <https://kubernetes.io/docs/home/>

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	24 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b>
			Final



components of the NEMO meta-OS, from Kubernetes API we extract Node information and launch the sub-component. Currently [version 1.31](#) is used.

Internally with NEMO, there are also some dependencies:

- **Underlay Network Probes.** NeMeX requires the deployment of the Underlay Network Probes to obtain the metrics extracted from the network. If probes are not deployed or are failing, there would be no content to expose. This would not cause a crash in the system, but the behavior would not be the expected one for the main workflow.
- **RabbitMQ.** Like what happens with the Underlay Network Probes, if there is no broker to exchange the metrics, NeMeX will be unable to push them, and therefore the behavior will not be the desired one. Also, it would neither kill NeMeX, just stopping the exposition service.

Also, there are some Python libraries also used and needed in this deployment. Although there is no need for a specific version in each of those technologies, it may be interesting to mention them just in case new implementations of those technologies include incompatibilities:

```
requests~=2.25
pika~=1.0
kubernetes~=31.0
```

About the System dependencies, all code has been tested and executed in Linux (Ubuntu 22.04 LTS) but, as the code is deployed using Kubernetes and Python3 (an interpreted language) there is no special dependency on the system to be used. Also, there are neither dependencies with hardware capabilities.

Finally, there are also configuration dependencies. These dependencies are related to the arguments to be provided during the deployment of the component:

- **delay:** Time between updates (default: 3600)
- **cluster:** Cluster id (default: 'cluster-1')
- **rmq\_ip:** IP where the RabbitMQ broker is available (default: '132.227.122.23')
- **rmq\_port:** Port where the RabbitMQ broker is available (default: '30403')
- **exchange:** Exchange used in the Rabbit MQ (default: '')
- **key:** Key used in RabbitMQ (default: 'costs-map')
- **user:** Nemo user for pushing RabbitMQ messages (default: 'nemo-user')
- **passwd:** Password to push RabbitMQ messages (default: '1234')

These arguments should be passed using the Kubernetes manifest available in the NeMeX official repository <sup>5</sup>.

#### 4.2.2.2 Internal workflow

When the NeMeX execution starts, this subcomponent evaluates if it has any topology stored locally (for example, in case tests are performed and the function in charge of reading the metrics is launched more than once without deleting the local data). In case there is no topology already stored, one will be loaded natively. Initially this topology was enhanced through a graph management, to provide path

<sup>5</sup> [kubefiles/02\\_nemex-deployment.yaml · main · Eclipse Research Labs / NEMO Project / NEMO Infrastructure Management / Federated Meta-Network Cluster Controller / Network Exposure module · GitLab](#)

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	25 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
		<b>Version:</b>	1.0
		<b>Status:</b>	Final

Computation capabilities, but due to the lack of demand for this service, we have proceeded to use local dictionaries, which consume a smaller amount of memory and have a faster management.

Once the topology update is initialized, two data sources will be evaluated: Kubernetes to obtain the mapping between the Kubernetes nodes and the IPs that represent them and the *Underlay Network Probes*. From the *Kubernetes API* (`client.CoreV1Api()`) the node information is loaded, extracting the name and IP values from the `metadata.name` and `status.addresses` field.

Afterwards, and having these measurements, the different metrics repositories of the *Underlay Network Probes*, which are present in each node of the cluster through a Prometheus service on port 5001, are analyzed. For each node, the metrics and properties are read and parsed into the format defined for the output. If any node does not have the service available or does not provide these metrics, the error will be detected, and it will be discarded from the process.

Once we have all the metrics processed, they are integrated in the same JSON, including the timestamp of when they have been generated and the cluster from which they come, to allow the identification of these. This resulting JSON is sent to the **RabbitMQ** broker indicated in the deployment and the process is paused until the next iteration.

This process is depicted in Figure 9.

PENDING EC APPROVAL

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	26 of 60				
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

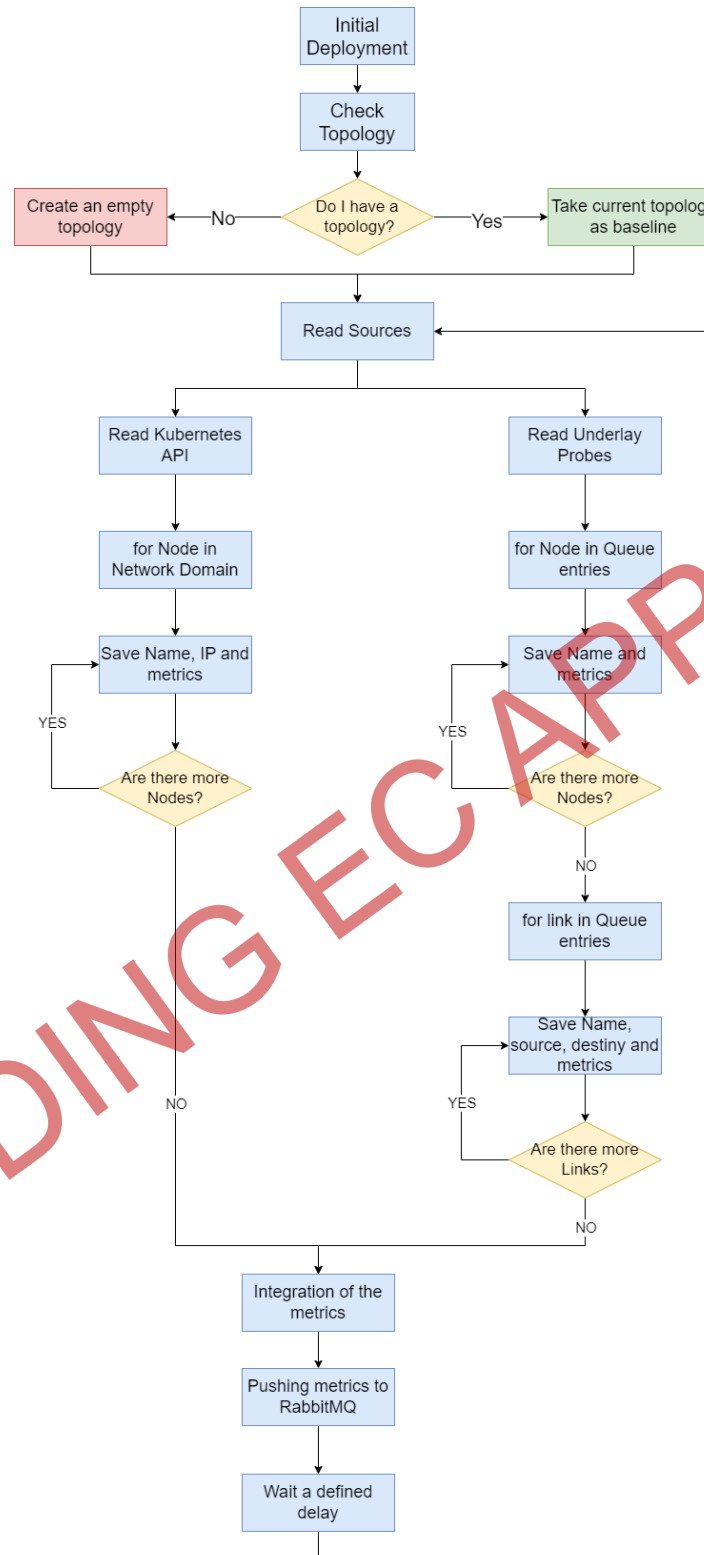


Figure 9 NeMeX internal workflow.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	27 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
		<b>Version:</b>	1.0
		<b>Status:</b>	Final

#### 4.2.2.3 Deployment and permissions.

The NeMeX component is designed to be deployed in a Kubernetes environment, in which write permissions are not required but read permissions are required. In the deployment process there are two steps that must be performed sequentially: on the one hand the definition of the *Namespace* where they will be executed (in the case of the Nemo project this is **nemo-net**) and the *Deployment* of the NeMeX resources. Both documents are available in the eclipse repository <sup>6</sup>.

Within NeMeX resources, we have 4 necessary deployments:

1. **Deployment:** this is the deployment of the code. Here the image to be instantiated is identified as well as the parameters that will serve as variables. These parameters can be parameterized or hardcoded. The user and password must be managed through Kubernetes secrets, which has not been uploaded to the repository for security reasons.
2. **ServiceAccount:** Creates an identity in Kubernetes for the deployment to manage the necessary resources for monitoring nodes.
3. **ClusterRole:** Role that NeMeX will perform within Kubernetes. In this section, read permissions are requested for the pods at namespace level and for the nodes at cluster level. This component can be replaced by a *Role*, but then it will only be able to monitor the pods at namespace level and not the nodes.
4. **ClusterRoleBinding:** Maps the permissions requested in the *Cluster Role* with the *Service Account* created. If a *Role* is used instead of a *Cluster Role*, a *RoleBinding* should be used instead.

#### 4.2.2.4 Inputs and outputs.

We can identify two main inputs: on the one hand we have the compute information obtained using the **Kubernetes API**; this information will help to match the nodes with their IPs. This input is managed using the Python3 Kubernetes API and the value obtained is a string with IP format per each node.

On the other hand, it is also received metric information from the network probes; these metrics will be associated with nodes and links and will be summarized in the following table:

Table 9. NeMeX input metrics.

Name	Description	Type	Valid values
<b>network_metric_latency</b>	Delay between two probes at IP level. Unit: ms	Float	0.01 – N
<b>network_metric_throughput</b>	Data rate available in the link. Unit: Mbps	Float	0.00 – N
<b>network_metric_packet_loss</b>	Packet loss rate since last measurement. Unit: %	Int	0 – 100
<b>network_metric_link_energy</b>	Average power consumption. Unit: W	Int	0 – N
<b>network_metric_link_failure</b>	Number of times the link drops since last measure. Unit: fails	Int	0 – N

<sup>6</sup> <https://gitlab.eclipse.org/eclipse-research-labs/nemo-project/nemo-infrastructure-management/federated-meta-network-cluster-controller/network-exposure-module/-/tree/main/kubefiles>

These metrics are processed and formatted to be aligned with the node names and IPs used by the Kubernetes resources to be managed. This integration returns the following data model:

```

1  {
2    "metadata": {
3      "name": "string",
4      "cluster": "string",
5      "timestamp": "integer",
6      "metrics": [
7        {
8          "name": "string",
9          "type": "string",
10         "metric": "string"
11        }
12      ],
13    },
14    "network-metrics": {
15      "nodes": [
16        {
17          "name": "string",
18          "ip": "string"
19        }
20      ],
21    },
22    "links": [
23      {
24        "name": "string",
25        "source": "string",
26        "destination": "string",
27        "metrics": [
28          {
29            "link-failure": "integer",
30            "link-energy": "integer",
31            "packet-loss": "integer",
32            "bandwidth": "float",
33            "latency": "float"
34          }
35        ]
36      }
37    ]
38  }

```

Figure 10 mNCC metrics data model

In Section 6 there will be available an example of the deployment and exposure of metrics, showing a demo scenario with three nodes.

### 4.2.3 Technology Connectivity Adaptors

The final working network adaptors. There are two types of connectivity adaptors: the ones part of the NEMO project, and the ones external to the project but used as a proof of concept for the versatility and flexibility of the mNCC architecture. In section 6 the demonstration for two cases will be shown.

#### 4.2.3.1 L2SM connectivity solution

Link Layer Secure Microservices (L2SM) serves as a secure connectivity solution within the mNCC, offering advanced network functionalities through an extension of the Kubernetes (K8s) API. It enables the creation, management, and deletion of virtual networks and supports the use of custom overlay networks. By taking advantage of the Software Defined Networking (SDN) capabilities, L2SM provides an API that deploys a customized data plane and enforces rules within Open Virtual Switches (OVS) as part of the overlay network topology.

L2SM has undergone significant advancements to align with the mNCC's evolving requirements. Initially, as presented in D2.2, L2SM handled internal cluster connectivity via overlays and introduced the concept of Network Edge Devices (NEDs) to enable the interconnection of overlay networks across different clusters. This concept has been further developed and implemented, facilitating communication between workloads deployed through NEMO OS using a standardized interface.

#### L2SM MD

The creation and management of L2SM resources are now handled via the L2SM multi-Domain (L2SM MD) client—a gRPC server that processes requests to create inter-cluster overlay topologies, add or delete clusters, and manage inter-cluster virtual networks atop these overlays. This modular approach simplifies the expansion of L2SM's functionalities, integrating existing concepts of virtual networks and overlay networks with the new NEDs and inter-cluster networking capabilities through a standardized API. This enhances the efficiency of the mNCC's connectivity adaptor.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	29 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

Specifically, L2SM MD implements the following methods:

- **Create and Delete Networks:** Allows the creation and deletion of virtual networks across clusters, given a list of clusters and the name of the network.
- **Create and Delete Overlays:** Manages overlay network topologies that define the interconnection of clusters, provided with a list of clusters and a name for the overlay.
- **Add and Remove Clusters:** Enables dynamic addition or removal of clusters within an overlay network, specified by the overlay name and the cluster to modify.

To securely manage these operations, L2SM MD requires cluster information, including the K8s API endpoint and a user's bearer token with permissions to manage L2SM Custom Resource Definitions (CRDs), Overlays, NEDs, and L2Networks in the managed clusters. Since the K8s API endpoint uses HTTPS, L2SM MD needs access to the TLS certificate to establish secure communication. Therefore, copies of the public certificates are stored securely within the control plane cluster as Kubernetes secrets.

To facilitate the setup of the gRPC server and ease the management of certificates and resources, L2SM MD provides a Command Line Interface (CLI) with a set of tools:

- **Apply-cert:** Takes a .key file and a cluster name as input, containing the public certificate, and stores it as a secret, allowing L2SM MD to access it securely.
- **Generate-cr:** Automates the generation of CRs by receiving input values for each resource. For example, when creating an Overlay, you input the names of clusters and links, and it generates the complete CR, including the necessary NED CRs required when connecting multiple clusters. This CLI tool mirrors the utility of the gRPC server but without actually creating the resources, serving as an alternative for quick fixes in case of errors in the workflow or for testing the adaptor in a resource-constrained environment.

#### 4.2.3.2 5G adapter: functions, capabilities and workflow update

The 5G adapter is responsible for managing data flows and providing monitoring and analytics capabilities within the 5G network. Here are its key functions with an extended definition in D2.2:

Function	Description
<b>Applying Data Flows</b>	Sets up new data transmission paths for efficient routing.
<b>Modifying Data Flows</b>	Adjusts existing data flows due to changing network conditions or to optimize performance.
<b>Deleting Data Flows</b>	Removes unnecessary data flows to free up network resources.
<b>Performance Monitoring</b>	Tracks metrics like latency, throughput, and error rates to ensure slice performance.
<b>Utilization Monitoring</b>	Observes resource usage to identify underutilized or overburdened slices.

Moreover, the deployment and request of 5G network flows has been updated to achieve an automated deployment by the meta-orchestrator and the following API exposure towards the IBS. In such a way

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	30 of 60	
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	
	<b>Version:</b>	1.0	<b>Status:</b>	Final

that, the MO can request a new network flow using the unified interface provided by the IBS. This workflow is not fully implemented at the date of this deliverable since MO tasks is part of WP3 and ends later in time. The general workflow is depicted on Figure 11.

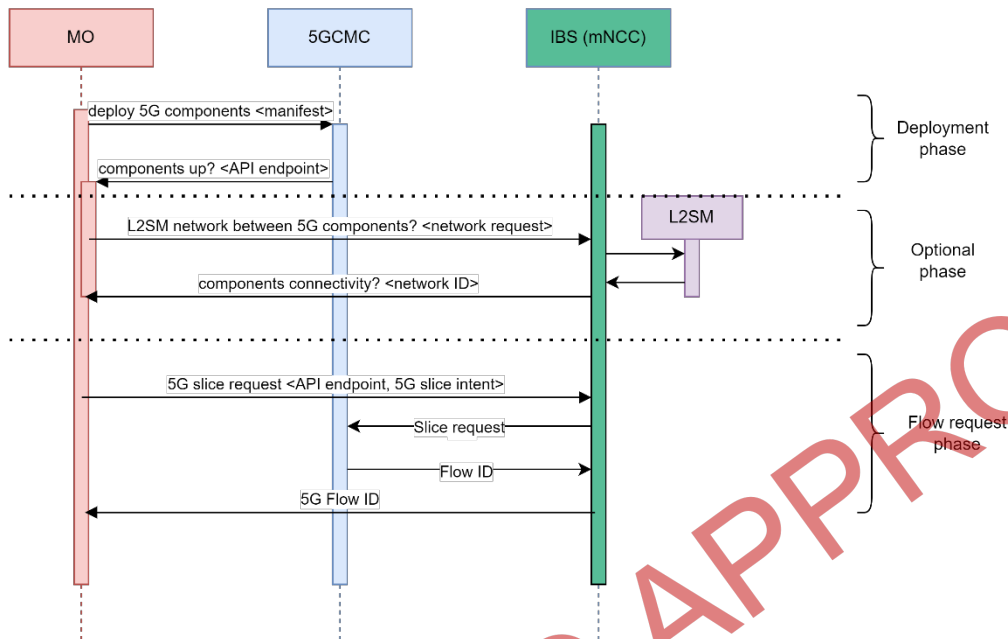


Figure 11 Towards integration workflow 5G adapter

In the Figure 11 there are shown three main phases internal interactions the network adapter:

- Deployment phase: is the one orchestrated by the MO, in charge of the deployment of the 5G core components such as the UPF, AMF, SMF...
- Optional phase: This phase has been theoretically described, as there is the possibility of connecting the different components of the core if deployed in different clusters through the L2SM adapter. From the L2SM perspective, this could be treated as a common service being connected through a L2SM network.
- Flow request phase: Once the adapter is deployed, the MO can start requesting new 5G flows through the intent-based interface connected to the RabbitMQ queue specific for the mNCC.

#### 4.2.3.3 IETF TeraFlow SDN-based connectivity adapter. (External)

The ETSI TeraFlowSDN<sup>7</sup> is an innovative, open-source, cloud-native Software-Defined Networking (SDN) controller and orchestrator designed to support smart networks and services for beyond 5G (B5G) and 6G networks. Developed by the ETSI Software Development Group TeraFlowSDN, it employs a micro-services architecture, enabling seamless integration with other ETSI initiatives such as OpenSourceMANO and compliance with standards from bodies like IETF and [2].

TeraFlowSDN is characterized by its high performance, advanced SDN automation, and support for transport network slicing, multi-tenancy, and cyberthreat analysis using machine learning (ML) and deep learning (DL) components. It also incorporates features like distributed ledger technology and smart contracts for secure network management. The controller is highly scalable, allowing for rapid

<sup>7</sup> TFS controller GitLab: <https://labs.etsi.org/rep/tfs/controller/-/wikis/home>

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	31 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final



prototyping and experimentation, and it supports the management of heterogeneous network equipment, including packet optical, IP and microwave networks [3].

The latest release, TeraFlowSDN 2.1, has been used by the mNCC IBS library to create a VPN layer 2 translation. A complete example of a request and the translated version in TFS controller format is shown in Annex 7. Briefly, the following context can be used to build an intent expectation to request a layer 2 VPN.

Table 10: objectContext for L2VPN

objectContexts	objectType	contextAttribute	contextCondition	contextValueRange	
	L2VPN	nodeSrc		IS_EQUAL_TO	ip string
		nodeDst		IS_EQUAL_TO	ip string
		endpointSrc		IS_EQUAL_TO	interface string
		vlanId		IS_EQUAL_TO	interface string
		niName		IS_EQUAL_TO	name string

Table 11: expectationTargets for L2VPN

expectationTargets	objectType	targetAttribute	targetCondition	targetValueRange	
	L2VPN	bandwidth		IS_EQUAL_TO_OR_GREATER_THAN	float
		latency		IS_LESS_THAN	float

#### 4.2.3.4 Network Slice Controller (External)

The IETF-based Network Slice Controller (NSC) processes end-to-end network slice requests initiated by 5G customers. These requests are coordinated by the 5G end-to-end orchestrator, which configures the Radio Access Network (RAN), and Core Network elements as needed and forwards the request to the NSC for execution. The NSC subsequently collaborates with the relevant network controllers to deploy the network slice within the transport network. [4]

There are some implementations of a network slice controller following the approach of the IETF draft (e.g., in-house Telefónica development in progress, not yet fully release as open source<sup>8</sup>). In the context of the project, the main characteristics that could have an intent for a slice request has been described as followed:

<sup>8</sup> Github: [Telefonica/network\\_slice\\_controller: Repository created for the Network Slice Controller \(NSC\) development](https://github.com/Telefonica/network_slice_controller)



objectType	contextAttribute	contextCondition	contextValueRange	
objectContexts	SLICE_SERVICE_SDP	slo-sle-policy	IS_EQUAL_TO	name string
		sdp-id	IS_EQUAL_TO	id string
		node-id	IS_EQUAL_TO	name string
		sdp-ip-address	IS_EQUAL_TO	ip string
		niName	IS_EQUAL_TO	name string

objectType	contextAttribute	contextCondition	contextValueRange	
objectTarget	SLICE_SERVICE_SDP	one-way-bandwidth	IS_EQUAL_TO	bandwidth string
		one-way-delay-maximum	IS_EQUAL_TO	delay string

#### 4.2.3.5 Auto peering API. (External)

Auto peering, as defined in this draft [5], refers to the automated process of establishing and managing peering relationships between different Autonomous Systems (ASes). This involves using an API to exchange necessary information, such as network details, routing policies, and contact information, to facilitate the setup and maintenance of peering connections without the need for manual intervention. The goal of auto peering is to simplify and streamline the peering process, reducing the time and effort required to establish new peering relationships and ensuring more efficient and reliable interconnectivity between networks. Although literally could not be perceived as a network connectivity adaptor, it is very useful towards future adaptors that could involve BGP session to access new routes or link state parameters [6] [7].

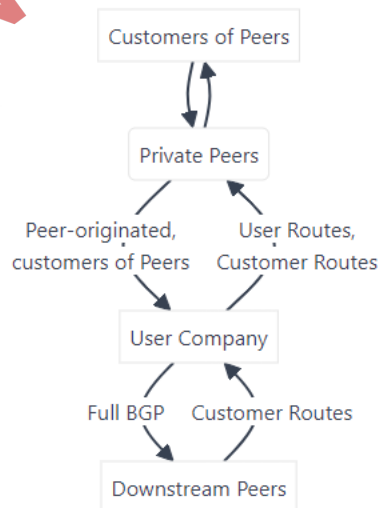


Figure 12: Private peering interactions

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	33 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

There is not a workable implementation of the process but, theoretically, the intents can be described accordingly<sup>9</sup>. The intentExpectation objects to request a new peering could contain mainly the following objectContexts:

Table 12: objectContext for BGP session

objectContexts	objectType	contextAttribute	contextCondition	contextValueRange
objectContexts	BGP_SESSION	localAsn	IS_EQUAL_TO	ASn integer
		localIp	IS_EQUAL_TO	ip string
		peerAsn	IS_EQUAL_TO	ASn integer
		peerType	IS_EQUAL_TO	private
		md5	IS_EQUAL_TO	key string

Table 13: expectationContext for BGP session

expectationContexts	objectType	contextAttribute	contextCondition	contextValueRange
expectationContexts	BGP_SESSION	location	IS_EQUAL_TO	string

#### 4.2.4 Network Performance Monitoring

The Network Performance Monitoring module within the mNCC provides the basic tools to evaluate real-time network conditions. This component leverages a Python-based network performance probe to measure key metrics across diverse network environments. By utilizing socket communication, it collects granular data on throughput, latency, and packet loss. These metrics are critical for ensuring efficient network operation and supporting the autonomous decision-making capabilities of other NEMO components.

The monitoring probe supports configurable parameters, allowing for fine-tuned evaluations tailored to specific network conditions or service requirements. For example, Users can define test duration, target hosts, and ports to suit their operational needs. The monitoring tool seamlessly integrates with the broader mNCC architecture, feeding network metrics directly into the meta-OS for actionable insights. Its modular design supports deployment via containerized environments, ensuring scalability and compatibility with multi-cluster scenarios. Leveraging the CI/CD framework of NEMO meta-OS, the monitoring tool is deployed and updated automatically, aligning with the intent-based management principles of the mNCC.

<sup>9</sup> In fact, from NEMO Project standardization contributions are being proposed for the definition of interconnection intents, e.g. <https://datatracker.ietf.org/doc/draft-contreras-nmrg-interconnection-intents/>.

### 4.3 Conclusion

The Federated Meta-Network Cluster Controller represents a key module in the NEMO meta-operating system, bringing an innovative solution for network management for edge-cloud environments. Throughout its development, the mNCC has evolved into a robust and flexible framework designed to address the challenges of scalability, adaptability and interoperability across diverse network technologies and multi-cluster scenarios.

The technological progression of the mNCC highlights its work as an infrastructure manager that simplifies and abstracts the complexities of network management. Conceptualized to cover the gap between the NEMO meta-Orchestrator and underlying network infrastructures, the mNCC has deployed its capabilities through the integration of intent-based networking and advanced metric exposure mechanisms with a secure solution for network connectivity, that works as the main technological adaptor, having also the capability to use up to five different technological adaptors according to the needed scenario (KPI 3.3, at least 5 technological adaptors). These developments have allowed the mNCC to not only provide seamless connectivity but also enable the capability of NEMO meta-OS to make dynamic and data-driven decisions.

The mNCC's architecture is built around the principles of abstraction, interoperability, and automation, enabling it to efficiently manage and monitor network operations across diverse environments. By leveraging an intent-based system, the mNCC translates high-level service expectations into detailed network configurations, allowing services to operate with greater autonomy and reduced reliance on manual intervention. This capability significantly enhances scalability, a critical requirement in distributed and multi-cluster settings, making the deployment of 50 Network Segments in 3.531s (KPI 3.1, < 5ms). Additionally, the integration of advanced metric collection and exposure mechanisms provides real-time insights into network performance. This feature supports network optimization and facilitates seamless interoperability between the mNCC and other NEMO components, such as the Meta-Orchestrator.

Looking forward, several areas for improvement and expansion have been identified to ensure the continued relevance and effectiveness of the mNCC. One priority is the integration of advanced security features, such as Zero Trust solutions, which will strengthen the module's defense mechanisms in distributed network scenarios. Additionally, optimizing resource allocation in dynamic and resource-constrained environments is an ongoing focus. Addressing these challenges, with the potential introduction of AI-based engines, is one of the main paths to follow during the next steps of the component.

In conclusion, the mNCC has completed the main goals expected for the component, even it is still at an early Technology Readiness Level (TRL), the proofs of concept done show a potential solution for multi-environment scenarios, allowing quick reactions and an abstraction level that reduces the complexity to integrate it with external modules, being the first step for seamless operation in next-generation network ecosystems.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	35 of 60				
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final



synchronized data transfer from the Grand Master to each UE. The TSN Application Function (AF) interfaces with the 5G control plane, translating TN parameters to 5GS-specific settings and reporting the capabilities of the 5GS bridge (e.g., delay specifications and topology details).

- **5G QoS Mapping:** Ethernet and TSN traffic flows are mapped to 5G QoS flows to maintain service quality. This mapping allows the CNC to optimize flows' handling across the network, supporting per-stream filtering, traffic policing, and VLAN configuration in compliance with IEEE 802.1Q.
- **Seamless Redundancy and Preemption:** TSN framework offers Seamless Redundancy for ultra-reliable communication and Frame Preemption to prioritize time-critical information.
- **Time Synchronization:** Through the TSN Translator (DS-TT) on UEs, precise time synchronization is achieved using gPTP messages, adhering to 3GPP specifications (24.535).

### 5.3 Internal and External Interfacing

**TSN AF** is part of 5GC and provides the control plane translator functionality for the integration of the 5GS with a TSN network, e.g. the interactions with the CNC. The TSN AF interfaces towards the CNC for the PSFP (IEEE Std 802.1Q) managed objects that correspond to the PSFP functionality implemented by the DS-TT and the NW-TT. Thus, when PSFP information is provided by the CNC, the TSN AF may extract relevant parameters from the PSFP configuration. The TSN AF calculates traffic pattern parameters (such as burst arrival time with reference to the ingress port and periodicity). TSN AF also obtains the flow direction. TSN AF is responsible for forwarding these parameters in TSC Assistance Information (TSCAI) Container to the SMF (via PCF). TSN AF may enable aggregation of TSN streams if the TSN streams belong to the same traffic class, terminate in the same egress port and have the same periodicity and compatible Burst arrival time. One set of parameters and one container are calculated by the TSN AF for multiple TSN streams to enable aggregation of TSN streams to the same QoS Flow.

**NW TT** (Network TSN Translator): supports link layer connectivity discovery and reporting as defined in IEEE Std 802.1AB for discovery of Ethernet devices attached to NW-TT. When integrating normal devices, we cannot assume that would be a DS-TT does not support link layer connectivity discovery and reporting, then NW-TT performs link layer connectivity discovery and reporting as defined in IEEE Std 802.1AB for discovery of Ethernet devices attached to DS-TT on behalf of DS-TT. If NW-TT performs link layer connectivity discovery and reporting on behalf of DS-TT, it is assumed that LLDP frames are transmitted between NW-TT and UE on the QoS Flow with the default QoS rule. Alternatively, SMF can establish a dedicated QoS Flow matching on the Ethertype defined for LLDP (IEEE Std 802.1AB).

**DS-TT** (Device Side TSN translator): translate the TSN 802.1 protocols on top of 5G networks but cannot be assumed that would be a DS-TT capable UEs.

### 5.4 Conclusion

As it points out, TSN main characteristics are performed: Seamless Redundancy that permits ultra-reliability due to the duplicity of buffers, Frame Pre-emption that includes a priority of the information that produces, and time critical communication. In addition to the Time Aware Shaper that allows scheduling of the outputs and Time Synchronization.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	37 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

## 6 Proof of Concepts

---

The Proof of Concepts chapter demonstrates the practical implementation and functionality of the key components developed within the NEMO project. This section showcases the integration and performance of the Cybersecure Micro-services Digital Twins (CMDT), Cybersecure Federated Deep Reinforcement Learning (CF-DRL), and the meta-Network Cluster Controller (mNCC) in various scenarios. Through detailed demonstrations and analyses, we evaluate the components against their defined Key Performance Indicators (KPIs) and illustrate their interactions within the project ecosystem. The following subsections provide in-depth insights into each component's deployment, operational workflows, and achieved results, highlighting the advancements made from the initial versions.

### 6.1 CMDT first stable version

---

#### 6.1.1 Detailed description

The easiest way of demonstrating the functionality of CMDT is to showcase it on a real example. For this reason, we took an application developed for NEMO Meta-OS, deploy it as digital twin and present the CMDT actions in live environment. For the app we selected the isolated ASM Terni deployment of PMU devices and applications that are controlling it. This will be our Digital Twin for this section. The simplified workload includes PMU devices deployed in the network, edge device for pre-processing the data locally on the site and cloud service, which is used to aggregate the data from edge devices and initiate actions towards the devices on the field. The application allows reviewing the status, gathering data and Firmware update over the air (FOTA). The digital twin of the workload comprises all mentioned device logs that can describe application footprint and behavior.

One crucial message that a user would request from the CMDT would be a list of services (Digital twin components) that represent a digital twin, and a corresponding status. A list of all workloads that CMDT can get from NEMO infrastructure is presented in Figure 14.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	38 of 60				
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

```

workloads:
  cbc208a-d535-434b-bb35-217a64bd516c:
    pods:
      0:
        name: "myrelease2-echo-server-5f7fd7b9f8-mbbdz"
        namespace: "default"
        labels:
          app.kubernetes.io/instance: "myrelease2"
          app.kubernetes.io/name: "echo-server"
          nemo.eu/workload: "cbc208a-d535-434b-bb35-217a64bd516c"
          pod-template-hash: "5f7fd7b9f8"
      a3177d01-863d-415b-a998-180c87113z50:
        pods:
          0:
            name: "migration-workload"
            namespace: "nemo-kernel"
            labels:
              nemo.eu/workload: "a3177d01-863d-415b-a998-180c87113z50"
              velero.io/backup-name: "one1ab-backup20241205154728"
              velero.io/restore-name: "one1ab-k3s-restore20241205161559"
      b6a77b9a-4cb2-41e9-953b-0a0b569c8cdb:
        pods:
          0:
            name: "demo-nginx-b87758879-g2bgh"
            namespace: "nemo-workloads"
            labels:
              app.kubernetes.io/instance: "demo-nginx"
              app.kubernetes.io/name: "demo"
              nemo.eu/workload: "b6a77b9a-4cb2-41e9-953b-0a0b569c8cdb"
              pod-template-hash: "b87758879"
          1:
            name: "demo-nginx-b87758879-g5h1c"
            namespace: "nemo-workloads"
            labels:
              app.kubernetes.io/instance: "demo-nginx"
              app.kubernetes.io/name: "demo"
              nemo.eu/workload: "b6a77b9a-4cb2-41e9-953b-0a0b569c8cdb"
              pod-template-hash: "b87758879"
          2:
            name: "demo-nginx-b87758879-ww51l"
            namespace: "nemo-workloads"
            labels:
              app.kubernetes.io/instance: "demo-nginx"
              app.kubernetes.io/name: "demo"
              nemo.eu/workload: "b6a77b9a-4cb2-41e9-953b-0a0b569c8cdb"
              pod-template-hash: "b87758879"
          3:
            name: "krallis-echo-server-7f7f56d759-7x6tb"
            namespace: "nemo-workloads"

```

Figure 14: List of resources per workload.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	39 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
		<b>Version:</b>	1.0
		<b>Status:</b>	Final



```

1  {
2    "pod_name": "app-bot-555c556d76-cdvh9",
3    "timestamp": "2024-12-11T08:32:13.956429Z",
4    "workload_id": "455ccf73-6d05-4f30-b05e-6a43dc211181",
5    "status": [
6      {
7        "status_type": "Ready",
8        "status": "True",
9        "last_transition_time": "2024-12-11T08:01:55Z",
10       "message": null
11      },
12      {
13        "status_type": "ContainersReady",
14        "status": "True",
15        "last_transition_time": "2024-12-11T08:01:55Z",
16        "message": null
17      },
18      {
19        "status_type": "PodReadyToStartContainers",
20        "status": "True",
21        "last_transition_time": "2024-12-11T08:00:10Z",
22        "message": null
23      },
24      {
25        "status_type": "Initialized",
26        "status": "True",
27        "last_transition_time": "2024-11-15T12:24:47Z",
28        "message": null
29      },
30      {
31        "status_type": "PodScheduled",
32        "status": "True",
33        "last_transition_time": "2024-11-15T12:24:46Z",
34        "message": null
35      }
36    ],
37    "labels": {
38      "app": "app-bot",
39      "linkerd.io/control-plane-ns": "linkerd",
40      "linkerd.io/proxy-deployment": "app-bot",
41      "linkerd.io/workload-ns": "demoapp",
42      "nemo.eu/workload": "455ccf73-6d05-4f30-b05e-6a43dc211181",
43      "pod-template-hash": "555c556d76",
44      "version": "v11"
45    },
46    "traffic_stats": {
47      "req_rate": 0.32079628767407103,
48      "res_rate": 2.1029978858633545,
49      "res_rate_by_code": {
50        "500": 0.17153690382571854,
51        "200": 1.9314609820376358
52      },
53      "res_time_quantiles_ms": {
54        "p99": 18.702803738317762,
55        "p95": 11.644859813084098,
56        "p75": 7.683397683397683,
57        "p50": 4.876404494382022
58      }
59    }
60  }

```

Figure 15 The CMTD message example which is transported over RabbitMQ

The CMTD monitors the status of each workload component continuously and distributes this information accordingly. This represents workloads DT status, which is timestamped for each event and transmitted the LCM and LCM – UI service. This notification path is crucial to notify the user about the status of any component in the workload. The messages that are passed over the RabbitMQ are sent for each component on the #CMTD channel. Example of such CMTD message structure (Figure 15) includes all information that was already explained and defined in the Deliverable D2.2, e.i. rates of times in quantiles and rates of response codes, e.g. 200 for OK and 500 for Server Error. From “res\_rate\_by\_code” parameter (Figure 15) user can find that our demo application two requests per minute with return the content successfully (code 200) and once per 6 minutes application responded with server error (Code 500).

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	40 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final



An additional improvement from the last version of CMDT is the introduction of the pod status block. For higher reliability and avoiding missing the information when some services are restarted, the block includes last N status changes attributed with the timestamps. As already presented in this document, the recipient of this data is LCM UI, which creates graphics and continuously updates user with the status of his digital twin workload living on NEMO infrastructure.

In this proof of concept, we presented the CMDT operation on real-world example that is deployed in NEMO. The workload status is constantly monitored, processed, collected and sent to the LCM UI. With the demonstrated example we showed the effectiveness of the CMDT service and outline the progress from the previous deliverable D2.2.

## 6.2 CF-DRL first stable version

We demonstrate the run of the CF-DRL for the scaling the replicas of the workloads in the Meta-Orchestrator.

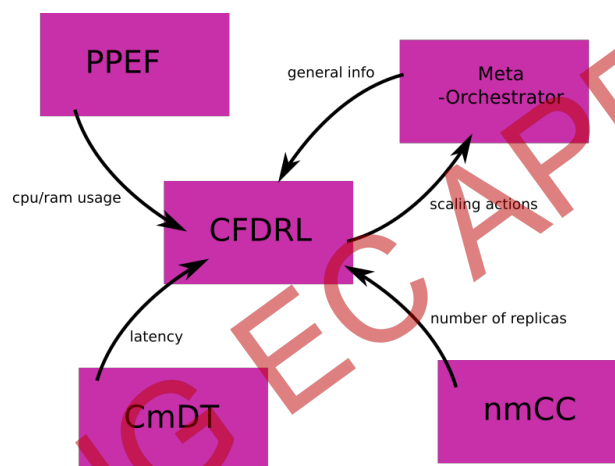


Figure 16 CF-DRL communication architecture for scaling the replicas of the workloads in the Meta-Orchestrator

The first step is to deploy the CF-DRL docker container to the Onelab cluster. Once the code is containerized. It is sent to the cluster using the commands:

```
export KUBECONFIG="/home/victor/nemo/victor.gabillon-kubeconfig.yaml"
kubectl apply -f -n nemo-ai
```

Then the run of the CF-DRL can be monitored using the kubernetes functions:

```
kubectl get pods -n nemo-ai
kubectl logs cf-drl-rabbit-deployment9-8677686865-tjxdb
```

The CF-DRL automatically connects to the Meta Orchestrator with RabbitMQ and with the PPEF through the intent-based API.

```
response = requests.get( url: 'https://intent-api.nemo.onelab.eu/api/v1/intent/', headers=headers)
print('response', response, type(response), response.status_code)
```

Figure 17: Code snippet for request


In the code the connection in Python to the intent-based API is made with a request to <https://intent-api.nemo.onelab.eu/api/v1/intent/>

The connection to the intent-based API requires an identification token (bearer token) that expires every 2 hours. A curl request periodically is sent to the NEMO Identity Management component to re-initiate the token.

```
if block_content['intent_report_reference']['intent_fulfillment_report'][
    'expectation_fulfillment_results']:
    # print('-----', len(block_content['intent_report_
    for block_ex_res in block_content['intent_report_reference']['intent_fulfillment_report'][
        'expectation_fulfillment_results']:
            assert (2 == len(block_ex_res['target_fulfillment_results']))
            # print('block_content', json.dumps(block_ex_res['target_fulfillment_results'], indent=4))
            cpu_achieved = block_ex_res['target_fulfillment_results'][0]['target_achieved_value']
            ram_achieved = block_ex_res['target_fulfillment_results'][1]['target_achieved_value']
            if cpu_achieved is not None:
                print(
                    f'workload_id {workload_id} cpu_achieved {cpu_achieved} ram_achieved {ram_achieved}
                    f' cpu_usage_target {cpu_usage_target} ram_usage_target {ram_usage_target}'
                )
                value_one_workload = [workload_id, cpu_achieved, ram_achieved, cpu_usage_target,
                                     ram_usage_target]
                value_at_time.append(value_one_workload)
```

The information received from the PPEF (through the intent-based API) is filtered in order to focus on the CPU and RAM of each workloads.

This allow us to display the change of the CPU and RAM though time:



#### Meta Orchestrator connection:

The connection with the meta-orchestrator is through RabbitMQ. On the CFDRL side the connection is made in python using the pika library.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	42 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

```
def reading():
    connection = pika.BlockingConnection(connection_parameters)
    channel = connection.channel()
    # channel.queue_declare(queue='mo_cfdrl')
    channel.basic_consume(queue='mo_cfdrl', auto_ack=True, on_message_callback=on_message_received)
    print(f'Starting Consuming')
    logger.info(f'Starting Consuming')
    channel.start_consuming()
```

Figure 18: RabbitMQ connectivity snipped

A queue was set in the RabbitMq so that CFDRl can query it for general info about the network as illustrated in the Figure below. CFDRl obtains information about the cluster id, capacity such as CPU and memory.

```
NAME                 READY   STATUS    RESTARTS   AGE
cfdr1-rabbit-deployment-867768685-tjydb   1/1     Running   0           60s
(base) victor@capucin:~$ kubectl logs cfdr1-rabbit-deployment-867768685-tjydb
hello bltcoonec to rab
Starting Consuming
received new message b'{"response":[{"name":"k3s-cluster","uid":"1a858938-8dda-4986-86c4-04c0f9831cbb","creationTimestamp":"2024-10-03T15:11:38Z","hubAcceptsClient":true,"capacity":{"cpu":"12","ephemeral-storage":"2029975572k","memory":"23983932k"},"pods":"330"},"allocatable":{"cpu":"12","ephemeral-storage":"1974769234894","memory":"23983932k"},"pods":"330"},"version":"v1.30.4+k3s1"}],"name":"rasptest","uid":"26ab1655-f353-4874-9c02-7c4d4dd694dda","creationTimestamp":"2024-09-10T08:08:03Z","hubAcceptsClient":true,"capacity":{"cpu":"4","ephemeral-storage":"29485028k"},"memory":"4143584k"},"pods":"110"},"allocatable":{"cpu":"4","ephemeral-storage":"286830352k"},"memory":"4143584k"},"pods":"110"},"version":"v1.30.4+k3s1"},"status":"ok"}]
received new message b'{"response":[{"name":"k3s-cluster","uid":"1a858938-8dda-4986-86c4-04c0f9831cbb","creationTimestamp":"2024-10-03T15:11:38Z","hubAcceptsClient":true,"capacity":{"cpu":"12","ephemeral-storage":"2029975572k"},"memory":"23983932k"},"pods":"330"},"allocatable":{"cpu":"12","ephemeral-storage":"1974769234894","memory":"23983932k"},"pods":"330"},"version":"v1.30.5+k3s1"}],"name":"rasptest","uid":"26ab1655-f353-4874-9c02-7c4d4dd694dda","creationTimestamp":"2024-09-10T08:08:03Z","hubAcceptsClient":true,"capacity":{"cpu":"4","ephemeral-storage":"29485028k"},"memory":"4143584k"},"pods":"110"},"allocatable":{"cpu":"4","ephemeral-storage":"286830352k"},"memory":"4143584k"},"pods":"110"},"version":"v1.30.4+k3s1"},"status":"ok"}]
received new message b'{"response":[{"name":"k3s-cluster","uid":"1a858938-8dda-4986-86c4-04c0f9831cbb","creationTimestamp":"2024-10-03T15:11:38Z","hubAcceptsClient":true,"capacity":{"cpu":"12","ephemeral-storage":"2029975572k"},"memory":"23983932k"},"pods":"330"},"allocatable":{"cpu":"12","ephemeral-storage":"1974769234894","memory":"23983932k"},"pods":"330"},"version":"v1.30.5+k3s1"}],"name":"rasptest","uid":"26ab1655-f353-4874-9c02-7c4d4dd694dda","creationTimestamp":"2024-09-10T08:08:03Z","hubAcceptsClient":true,"capacity":{"cpu":"4","ephemeral-storage":"29485028k"},"memory":"4143584k"},"pods":"110"},"allocatable":{"cpu":"4","ephemeral-storage":"286830352k"},"memory":"4143584k"},"pods":"110"},"version":"v1.30.4+k3s1"},"status":"ok"}]'
```

Figure 19: Metric retrieving

## 6.3 mNCC hybrid clusters and intent networking

### 6.3.1 Single Cluster Connectivity throughout L2SM overlay adaptor

In order to demonstrate the advances of the L2S-M connectivity adaptor in the NEMO project, this new PoC showcases the automatic setup of the adaptor in a NEMO inter-domain cluster environment in order to enable the deployment of virtualized workloads within the infrastructure of the project. Since the internal components that enabled this connectivity were previously shown in D.2.2, this new PoC will focus on showcasing the ability of the new L2S-M-MD module (described in Section 4) to properly configure new k8s clusters to the NEMO project, ensuring their ability to create inter-domain virtual networks that enable the connectivity of network functions in multiple K8s clusters.

Particularly, this setup consists of three separate K8s clusters: one controller cluster and two worker clusters, all following the guidelines of the NEMO project hierarchy and its characteristics. In this regard, for the K8s control plane cluster, the L2SM-MD-server, the IDCO provider and multidomain DNS components were installed, while the L2S-M operator and its components were installed in the L2S-M namespace of the K8s worker clusters. This setup can be seen in Figure 20: K8s clusters of this PoC.

```
alex@alex-pc:~/Documents/l2sm/l2sm-md$ kind get clusters
control-plane
worker-cluster-1
worker-cluster-2
```

Figure 20 K8s clusters of this PoC

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	43 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final





Once all the elements have been created to enable the management of the K8s resources in the clusters that are going to be managed by the mNCC component, the system administrator can proceed to create a slice (i.e., inter cluster overlay that will logically connect both clusters) between both clusters. To do so, it is necessary to introduce the API endpoints and the generated tokens to perform a gRPC call, as it can be seen in Figure 24:

```
const (
    serverAddress      = "localhost:30000"
    networkName        = "ping-network"
    providerName        = "test-slice"
    providerDomain      = "172.20.0.2"
    networkType        = "vnet"
    clusterName1        = "kind-worker-cluster-1"
    clusterName2        = "kind-worker-cluster-2"
    clusterBearerToken1 = "eyJhbGciOiJIUzI1NiIsImtpZCI6InIaYlMxJTnR4MktRa33dnhbJg"
    clusterBearerToken2 = "eyJhbGciOiJIUzI1NiIsImtpZCI6InIxdElrON6F3bVF0BfVBk0FR"
    clusterApiKey1      = "https://172.20.0.3:6443"
    clusterApiKey2      = "https://172.20.0.4:6443"
    node1Cluster1       = "worker-cluster-1-control-plane"
    node1Cluster2       = "worker-cluster-2-control-plane"
)
```

Figure 24 Configuration file used to create a new slice with the L2S-M-md component

With these values, the administrator can proceed with the creation using the available API in the manager cluster, which will install both NEDs and the overlay network (slice) in the managed K8s clusters as their own resources, as seen in Figure 25:

```
alex@alex-pc:~/Documentos/l2sn/l2sn-md$ go run test/client.go --test-slice-create
Creating slice...
CreateSlice response: slice 'test-slice' created successfully.
```

Figure 25 CLI used to create the slice between the clusters

```

alex@galex-pc:~/Documentos/l2sm/l2sm-md$ kubectl get l2networks --context kind-worker-cluster-1
NAME                AVAILABILITY  CONNECTED_PODS  AGE
ping-network        Available     23s
alex@galex-pc:~/Documentos/l2sm/l2sm-md$ kubectl get l2networks --context kind-worker-cluster-2
NAME                AVAILABILITY  CONNECTED_PODS  AGE
ping-network        Available     22s
alex@galex-pc:~/Documentos/l2sm/l2sm-md$ kubectl describe l2networks ping-network --context kind-worker-cluster-2
Name:                ping-network
Namespace:           default
Labels:              <none>
Annotations:         <none>
API Version:         l2sm.l2sm.k8s.local/v1
Kind:                L2Network
Metadata:
  Creation Timestamp: 2024-12-12T13:32:28Z
  Finalizers:
    l2sm.operator.io/finalizer
  Generation:        1
  Resource Version:  12305
  UID:               3b74dfb2-d93f-42b4-9509-c033a704443b
Spec:
  Provider:
    Domain: 172.20.0.2:8181
    Name:   test-slice
    Type:   vnet
Status:
  Internal Connectivity: Available
  Provider Connectivity: Available
Events:                  <none>

```

Figure 26 L2S-M resources successfully installed in both clusters

After this installation, the mNCC can interconnect NEMO workloads in these managed K8s clusters, which showcases how the new L2SM-MD component can greatly assist the procedure of setting up the inter-cluster communication between K8s clusters belonging to a NEMO infrastructure.

### 6.3.2 Intent-Based System

In order to showcase the advancements in the Intent-Based System, two different cases will be described. The purpose of the first one will be to describe the full intent lifecycle that the IBS follows to translate an Intent to the specific network technology. In this case it will be using the TFS adaptor. The second one, will show the integration with the L2SM component and making use of the cloud continuum intents for intercluster connectivity.

#### 6.3.2.1 Intent lifecycle TFS L2VPN adaptor.

As explained in previous sections and deliverables (4.2.1, [1]), the Intent lifecycle is divided mainly in classification, translation and execution. The next logs snippets are part of an execution of the IBS managing the request of a new L2VPN creation in the intent format.

At the very beginning, the IBS detects all the intent libraries installed and loads all of them to later be able to translate to the different underlying technologies. This phase is divided in two parts: the one importing libraries Figure 27, and the one importing the executioners Figure 28.

```

[DEBUG - 14:18:24] root: [MainThread] Logger in DEBUG mode
----- Importer-----
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] Importing from: /home/ubuntu/Repos/gitlab/intent-based-system/intent_engine
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] module_name: intent_engine.catalogue.l2sm
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] imported_module: <module 'intent_engine.catalogue.l2sm' from '/home/ubuntu/R
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] module_name: intent_engine.catalogue.nemo
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] imported_module: <module 'intent_engine.catalogue.nemo' from '/home/ubuntu/R
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] module_name: intent_engine.catalogue.l2vpn
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] imported_module: <module 'intent_engine.catalogue.l2vpn' from '/home/ubuntu/t
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] module_name: intent_engine.catalogue.tfs_controller
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] imported_module: <module 'intent_engine.catalogue.tfs_controller' from '/home

```

Figure 27 IBS importing libraries

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	46 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final



```
[DEBUG - 14:18:24] intent_engine.core.intent_core: [MainThread] None
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] Importing from: /home/ubuntu/Repos/gitlab/intent-based-s
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] module_name: intent_engine.executioners.sys_out
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] imported_module: <module 'intent_engine.executioners.sys
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] module_name: intent_engine.executioners.http_handler
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] imported_module: <module 'intent_engine.executioners.htt
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] module_name: intent_engine.executioners.tfs_connector
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] imported_module: <module 'intent_engine.executioners.tfs
[DEBUG - 14:18:24] intent_engine.core.importer: [MainThread] module_name: intent_engine.executioners.grpc_connector
/home/ubuntu/.local/lib/python3.10/site-packages/google/protobuf/runtime_version.py:112: UserWarning: Protobuf gencod
```

Figure 28 IBS importing executioners

All the executioners depicted in the previous figure are later available for the different libraries in case they are required. There is a specific executioner created for debugging purposes called “sys\_out”, used to show the output of the intent translation. Any of them can be used at the same time, so at the end, the one handling TFS and the one handling the system out print will be activated. The fact that this is possible is showcasing also the flexibility of the architecture and the easy integration with different executioners and libraries depending on the needs of the user.

The next screenshot shows the intent incoming and having a first format check. In Figure 29, an arrow is showing the objectType of the intent, in this case “L2VPN”. It is important to remark this because the intent will be understood as a layer two VPN but not yet as a TFS L2VPN.

```
----- Wating Intent -----
[DEBUG - 14:18:24] intent_engine.core.ib_model: [MainThread] Schema Type Expectation: <class 'intent_engine.core.IntentNm.NewL2VPNIntentExpectation'>
{
  'id': 'new_intent_2',
  'userLabel': 'cloud_continuum',
  'intentExpectations': [
    {
      'expectationId': '1',
      'expectationVerb': 'DELIVER',
      'expectationObject': {
        'objectType': 'L2VPN',
        'objectInstance': 'l2vpn_tfs_1',
        'objectContexts': [

```

Figure 29 IBS input intent

In the Figure 30, it is showed how the comparison between the different libraries is traversed, and at the end, how the IBS is able to classify the intent as “L2VPN”. This is due to the fact that, in the classification tree used, when the system reaches a leaf, it means it fully matches all the keywords for the library.

```
----- Intent classifier -----
[DEBUG - 14:18:24] intent_engine.core.intent_classifier: [MainThread] Keywords: ['cloud_continuum', 'DELIVER', 'L2VPN']
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] cloud_continuum : {'DELIVER': {'L2SM_NETWORK': 'l2sm',
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] DELIVER : {'L2SM_NETWORK': 'l2sm', 'L2SM_NEW_CLUSTER'
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] L2SM_NETWORK : l2sm
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] L2SM_NEW_CLUSTER : l2sm
[DEBUG - 14:18:24] intent_engine.core.intent_classifier: [MainThread] Keywords: ['cloud_continuum', 'DELIVER', 'L2VPN']
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] cloud_continuum : {'DELIVER': {'K8S_L2_NETWORK': 'nemo
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] DELIVER : {'K8S_L2_NETWORK': 'nemo', '5G_NETWORK_SLICE
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] K8S_L2_NETWORK : nemo
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] 5G_NETWORK_SLICE : nemo
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] K8S_CLUSTER_CONFIG : nemo
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] L2_VPN : l2vpn
[DEBUG - 14:18:24] intent_engine.core.intent_classifier: [MainThread] Keywords: ['cloud_continuum', 'DELIVER', 'L2VPN']
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] cloud_continuum : {'DELIVER': {'L2VPN': 'l2vpn'}, 'mir
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] DELIVER : {'L2VPN': 'l2vpn'}
[INFO - 14:18:24] intent_engine.core.intent_classifier: [MainThread] L2VPN : l2vpn
is leaf : l2vpn
```

Figure 30 IBS first classification

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	47 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final



The shows the IBS passing the intent to the L2VPN library. As this library doesn't generate ILUs (Intent Logic Units), the library reclassifies the intent (now called sub-intent), into a less abstracted library the TFS controller one. This step is important to abstract less abstracted libraries from higher ones, this means that the TFS controller library has no need to be aware of more abstracted L2VPN request, and the L2VPN must be aware of the specific technology library but not of the full translation to the technology. The step of reclassifying the intent is done until a library capable of generating an ILU is executed.

```
[DEBUG - 14:18:24] intent_engine.core.intent_classifier: [MainThread] new iter subintent module: l2vpn
[DEBUG - 14:18:24] intent_engine.core.intent_classifier: [MainThread] reclassify...
[DEBUG - 14:18:24] intent_engine.catalogue.l2vpn: [MainThread] Generating L2VPN TFS subintent...
{
  'id': 'new_intent_2',
  'userLabel': 'cloud_continuum',
  'intentExpectations': [
    {
      'expectationId': '1',
      'expectationVerb': 'DELIVER',
      'expectationObject': {
        'objectType': 'L2VPN_TFS',
        'objectInstance': 'l2vpn_tfs_1',
        'objectContexts': [
```

Figure 31: IBS subintent generation.

Also, in the reclassification step, the library can make adjustments to the intent so that the following library is able to detect the intent as own. As shown in Figure 31 the “*objectType*” of the intent has changed. Now the sub-intent is passed through the classification again and detected as “*tfs\_controller*” intent.

```
[DEBUG - 14:36:56] intent_engine.core.intent_classifier: [MainThread] Keywords: ['cloud_continuum', 'DELIVER', 'L2VPN_TFS']
[INFO - 14:36:56] intent_engine.core.intent_classifier: [MainThread] cloud_continuum : {'DELIVER': {'L2VPN_TFS': 'tfs_controller'}}
[INFO - 14:36:56] intent_engine.core.intent_classifier: [MainThread] DELIVER : {'L2VPN_TFS': 'tfs_controller'}
[INFO - 14:36:56] intent_engine.core.intent_classifier: [MainThread] L2VPN_TFS : tfs_controller
is leaf : tfs_controller
[DEBUG - 14:36:56] intent_engine.core.intent_classifier: [MainThread] Keywords: ['cloud_continuum', 'DELIVER', 'L2VPN_TFS']
[INFO - 14:36:56] intent_engine.core.intent_classifier: [MainThread] cloud_continuum : {'CREATE': {'5G_SLICE_FLOW': 'slice_5g'}}
[INFO - 14:36:56] intent_engine.core.intent_classifier: [MainThread] CREATE : {'5G_SLICE_FLOW': 'slice_5g'}
[DEBUG - 14:36:56] intent_engine.core.intent_classifier: [MainThread] unique list: ['tfs_controller']
[DEBUG - 14:36:56] intent_engine.core.intent_classifier: [MainThread] new iter subintent module: tfs_controller
[DEBUG - 14:36:56] intent_engine.core.intent_classifier: [MainThread] Module is ilu: <intent_engine.catalogue.tfs_controller.tfs_contr
[DEBUG - 14:36:56] intent_engine.core.intent_classifier: [MainThread] translators iteration ILU: ['tfs_controller']
```

Figure 32 IBS second classification

In the Figure 32, it is also shown how now the intent is classified as ILU. From this point, it starts the translation phase. In the Figure 33, it is shown the parsing of the intent, and how it traverses all the contexts, targets and expectations.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	48 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

```

----- Intent translator -----
[INFO - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] Translating TFS connector...
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] debug TFS connector...
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] expectation case DELIVER
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] request case obj: objectType='L2VPN_TFS', contextValueRange='3.3.3.3', NodeDstContext(contextAttribute='nodeDst', contextCondition='IS_EQUAL_TO', contextValueRange='GigabitEthernet0/0/1'), EndpointDstContext(contextAttribute='endpointDst', contextCondition='IS_EQUAL_TO', contextValueRange='999'), NiNameContext(contextAttribute='niName', contextCondition='IS_EQUAL_TO', contextValueRange='3.3.3.3')
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] objectctx case contextAttribute='nodeSrc'
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] node_src case 3.3.3.3
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] params after {'context_uid': 'admin', 'service_id': '12-ac1-svc-intent', 'endpoint_dst': 'GigabitEthernet0/0/1', 'bandwidth': '10.0', 'latency': '15.2', 'vlan_id': 999, 'circuit_id': '12-ac1-svc-intent'}
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] objectctx case contextAttribute='nodeDst'
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] node_dst case
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] objectctx case contextAttribute='endpointDst'
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] objectctx case contextAttribute='endpointId'
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] objectctx case contextAttribute='vlanId'
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] vlan_id case
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] objectctx case contextAttribute='niName'
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] bandwidth case
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] latency case
[DEBUG - 14:18:24] intent_engine.catalogue.tfs_controller: [MainThread] url case
----- Execution platform -----

```

Figure 33 IBS tfs\_l2vpn translation

Finally, when the intent is translated, the library in charge of the translation calls the executioners needed to execute the intent. As explained earlier this case shows a first one calling “*sys\_out*” () and a second one to the real “*tfs\_connector*” ().

```

----- Execution platform -----
[INFO - 14:18:24] intent_engine.core.intent_core: [MainThread] Running sys_out
-> -> Executing:
[
  {
    'services': [
      {
        'service_id': {
          'context_id': {
            'context_uid': {
              'uid': 'admin',
            },
          },
          'service_uid': {
            'uid': '12-ac1-svc-intent',
          },
        },
      },
    ],
  },
]

```

Figure 34 IBS "sys\_out" executioner

```

[
  {
    'url': 'http://192.168.165.10/webui',
    'headers': {
      'Content-Type': 'multipart/form-data',
    },
    'connect_type': 'get',
  },
]
<- <-
[INFO - 14:18:24] intent_engine.core.intent_core: [MainThread] Running tfs_connector

```

Figure 35 IBS "tfs\_connector" executioner

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	49 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

Now, the executioner starts the proper protocol to create a connection to TFS north bound interface and makes a request for a new L2VPN connection.

```
[DEBUG - 14:18:24] urllib3.connectionpool: [MainThread] Starting new HTTP connection (1): 192.168.165.10:80
[DEBUG - 14:18:24] urllib3.connectionpool: [MainThread] http://192.168.165.10:80 "GET /webui HTTP/1.1" 200 9465
[DEBUG - 14:18:24] intent_engine.executioners.tfs_connector: [MainThread] csrf token IjZkNTNkNmIyODkxMTk4MDI0ZjE4MD
[DEBUG - 14:18:24] intent_engine.executioners.tfs_connector: [MainThread] Data to send: {'services': [{'service_id':
2, 'service_status': {'service_status': 1}, 'service_endpoint_ids': [{'device_id': {'device_uuid': {'uuid': '3.3.3.
'endpoint_uuid': {'uuid': '0/0/1-GigabitEthernet0/0/0/1'}}}], 'service_constraints': [{'custom': {'constraint_type':
}}], 'service_config': {'config_rules': [{'action': 1, 'custom': {'resource_key': '/settings', 'resource_value': {
urce_value': {'sub_interface_index': 0, 'ni_name': 'ninametfsssss', 'vlan_id': 999, 'circuit_id': '999', 'remote_r
]/settings', 'resource_value': {'sub_interface_index': 0, 'ni_name': 'ninametfsssss', 'vlan_id': 999, 'circuit_id'
[DEBUG - 14:18:54] urllib3.connectionpool: [MainThread] http://192.168.165.10:80 "POST /webui HTTP/1.1" 302 201
[DEBUG - 14:18:54] urllib3.connectionpool: [MainThread] http://192.168.165.10:80 "GET /webui/ HTTP/1.1" 200 10111
[INFO - 14:18:54] intent_engine.executioners.tfs_connector: [MainThread] Http response: <!--
```

Figure 36 IBS http request to TFS

As a result of the process, a VPN service in TFS controller is created. In Figure 37, notice how the service is active on the nodes requested, and in Figure 38, are shown all the details translated from the *intentContext* and the *intentTargets* specified in the input intent.

### Services

+ Add New Service

1 services found in context *admin*

UUID	Name	Type	End points	Status
e35ab28c-5474-5adb-aaf0-7f426ff5d5b6	l2-acl-svc-intent	L2NM	<ul style="list-style-type: none"> <li>0/0/1-GigabitEthernet0/0/0/1 / Device: <a href="#">3.3.3.3</a></li> <li>0/0/1-GigabitEthernet0/0/0/1 / Device: <a href="#">5.5.5.5</a></li> </ul>	ACTIVE

Figure 37 TFS webui L2VPN service status

**Context:** 43813baf-195e-5da6-af20-b3d0922e71a7  
**UUID:** e35ab28c-5474-5adb-aaf0-7f426ff5d5b6  
**Name:** l2-acl-svc-intent  
**Type:** L2NM  
**Status:** ACTIVE

Endpoint UUID	Name	Device	Endpoint Type
90ee9eb3-0c93-5d9b-b726-5628cf44f517	0/0/1-GigabitEthernet0/0/0/1	<a href="#">3.3.3.3</a>	-
057b26e8-53cf-5ec9-9169-cf2298f2781e	0/0/1-GigabitEthernet0/0/0/1	<a href="#">5.5.5.5</a>	-

**Constraints:**

Kind	Key/Type	Value
Custom	bandwidth[gbps]	10
Custom	latency[ms]	15

**Configurations:**

Key	Value
/settings	<ul style="list-style-type: none"> <li>mtu: 1450</li> </ul>
/device[3.3.3.3]/endpoint[0/0/1-GigabitEthernet0/0/0/1]/settings	<ul style="list-style-type: none"> <li>circuit_id: 999</li> <li>ni_name: ninametfsssss</li> <li>remote_router: 5.5.5.5</li> <li>sub_interface_index: 0</li> <li>vlan_id: 999</li> </ul>
/device[5.5.5.5]/endpoint[0/0/1-GigabitEthernet0/0/0/1]/settings	<ul style="list-style-type: none"> <li>circuit_id: 999</li> <li>ni_name: ninametfsssss</li> <li>remote_router: 3.3.3.3</li> <li>sub_interface_index: 0</li> <li>vlan_id: 999</li> </ul>

Connection Id	Sub-Service	Path
58bfee2b-6b07-49c0-b9b6-16f5ad721b1f	<a href="#">3.3.3.3</a> / 0/0/1-GigabitEthernet0/0/0/1	<a href="#">5.5.5.5</a> / 0/0/1-GigabitEthernet0/0/0/1

Figure 38 TFS webui L2VPN service details

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	50 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final



```
ubuntu@netma-test-1:~/nemo/intent-based-system$ kubectl get l2networks --context kind-worker-cluster-1 -A
NAMESPACE   NAME           AVAILABILITY   CONNECTED_PODS   AGE
default     spain-network   Available      1                 2m18s
```

Figure 41: L2SM intercluster network active.

### 6.3.3 Multi-Cluster monitoring and exposure

In order to perform the PoC of the multi-cluster monitoring and exposure, the mNCC components in charge of monitoring and exposure of the mNCC have been deployed in a scenario with 3 nodes. The mNCC monitoring and exposure process is explained in section 4, so we will not extend with this explanation, but it is worth remembering that the monitoring is based on the communication between pods deployed in each node and then shared with the exposure module through a local Prometheus. Once the exposure module has those metrics, it formats them and indicates to which cluster they belong and at what time they have been generated.

For these tests, we have reduced the monitoring and exposure time to 2 minutes (120 seconds) to show how the RabbitMQ queue is being updated, although in the joint environment these measurements would be hourly (3600s) to avoid saturation in the network and in the metrics queues.

In the process of evaluating the mNCC metrics exposure we must start with the deployment of the probes and the exposure module (Nemex for short) in each cluster, counting on a pod of probes per worker node and a Nemex per cluster. For this proof of concept, we will perform the evaluation on 2 clusters in parallel.

In Figure 42 and Figure 43 we can see a test run of the monitoring pods in which the logs of one of the probes of each cluster are obtained. These logs show the measurements taken periodically and shared with the Nemex through an internal Prometheus (Figure 44 shows the format of these logs).

```
ubuntu@alto-server:~$ kubectl logs -n nemo-net network-probe-daemonset-5grrd
* Serving Flask app 'network_probe'
* Debug mode: off
2024-12-12 16:18:20,420 - INFO - Pushing metrics to Prometheus on port 5001.
2024-12-12 16:18:20,469 - INFO - -----
2024-12-12 16:18:20,470 - INFO - Performing network measurements ...
2024-12-12 16:18:20,470 - INFO - Target host: 172.18.0.2
2024-12-12 16:18:20,470 - INFO - Target port: 80
2024-12-12 16:18:20,470 - INFO - -----
2024-12-12 16:18:22,472 - INFO - Measured Throughput: 12809.54 Mbps
2024-12-12 16:18:22,472 - INFO - Measured Energy: 1691.07 W
2024-12-12 16:18:22,472 - INFO - Time taken for Throughput measurement: 2.0016 seconds
2024-12-12 16:18:22,472 - INFO - -----
2024-12-12 16:18:24,510 - INFO - Measured Packet Loss Rate: 0.00%
2024-12-12 16:18:24,510 - INFO - Time taken for Packet Loss measurement: 2.0378 seconds
2024-12-12 16:18:24,510 - INFO - -----
2024-12-12 16:18:24,511 - INFO - Measured Latency: 0.28 ms
2024-12-12 16:18:24,511 - INFO - Time taken for Latency measurement: 0.0004 seconds
2024-12-12 16:18:24,511 - INFO - -----
2024-12-12 16:18:24,511 - INFO - Total time taken for all measurements: 4.0412 seconds
2024-12-12 16:18:24,511 - INFO - -----
2024-12-12 16:18:27,512 - INFO - -----
2024-12-12 16:18:27,513 - INFO - Performing network measurements ...
2024-12-12 16:18:27,513 - INFO - Target host: 172.18.0.4
2024-12-12 16:18:27,513 - INFO - Target port: 80
2024-12-12 16:18:27,513 - INFO - -----
2024-12-12 16:18:29,514 - INFO - Measured Throughput: 13290.12 Mbps
2024-12-12 16:18:29,514 - INFO - Measured Energy: 1754.17 W
2024-12-12 16:18:29,514 - INFO - Time taken for Throughput measurement: 2.0010 seconds
2024-12-12 16:18:29,515 - INFO - -----
2024-12-12 16:18:31,569 - INFO - Measured Packet Loss Rate: 0.00%
2024-12-12 16:18:31,569 - INFO - Time taken for Packet Loss measurement: 2.0538 seconds
2024-12-12 16:18:31,569 - INFO - -----
```

Figure 42 Logs from the network probes in the first cluster.

Document name:	D2.3 Enhancing NEMO Underlying Technology	Page:	52 of 60
Reference:	D2.3	Dissemination:	PU
	Version:	1.0	Status:
			Final



```

ubuntu@nemo-codeco:~$ kubectl logs -n nemo-net network-probe-daemonset-z45xs
* Serving Flask app 'network_probe'
* Debug mode: off
2024-12-12 17:14:02,763 - INFO - Pushing metrics to Prometheus on port 5001.
2024-12-12 17:14:02,952 - INFO - -----
2024-12-12 17:14:02,952 - INFO - Performing network measurements ...
2024-12-12 17:14:02,952 - INFO - Target host: 172.18.0.2
2024-12-12 17:14:02,952 - INFO - Target port: 80
2024-12-12 17:14:02,952 - INFO - -----
2024-12-12 17:14:04,954 - INFO - Measured Throughput: 14891.87 Mbps
2024-12-12 17:14:04,954 - INFO - Measured Energy: 1965.62 W
2024-12-12 17:14:04,954 - INFO - Time taken for Throughput measurement: 2.0011 seconds
2024-12-12 17:14:04,954 - INFO - -----
2024-12-12 17:14:06,981 - INFO - Measured Packet Loss Rate: 0.00%
2024-12-12 17:14:06,981 - INFO - Time taken for Packet Loss measurement: 2.0270 seconds
2024-12-12 17:14:06,981 - INFO - -----
2024-12-12 17:14:06,981 - INFO - Measured Latency: 0.11 ms
2024-12-12 17:14:06,981 - INFO - Time taken for Latency measurement: 0.0002 seconds
2024-12-12 17:14:06,981 - INFO - -----
2024-12-12 17:14:06,982 - INFO - Total time taken for all measurements: 4.0295 seconds
2024-12-12 17:14:06,982 - INFO - -----
2024-12-12 17:14:09,983 - INFO - -----
2024-12-12 17:14:09,983 - INFO - Performing network measurements ...
2024-12-12 17:14:09,983 - INFO - Target host: 172.18.0.4
2024-12-12 17:14:09,983 - INFO - Target port: 80
2024-12-12 17:14:09,983 - INFO - -----
2024-12-12 17:14:11,985 - INFO - Measured Throughput: 18461.14 Mbps
2024-12-12 17:14:11,985 - INFO - Measured Energy: 2437.24 W
2024-12-12 17:14:11,985 - INFO - Time taken for Throughput measurement: 2.0013 seconds
2024-12-12 17:14:11,985 - INFO - -----
2024-12-12 17:14:14,040 - INFO - Measured Packet Loss Rate: 0.00%
2024-12-12 17:14:14,040 - INFO - Time taken for Packet Loss measurement: 2.0543 seconds
2024-12-12 17:14:14,040 - INFO - -----
2024-12-12 17:14:14,040 - INFO - Measured Latency: 0.14 ms
2024-12-12 17:14:14,040 - INFO - Time taken for Latency measurement: 0.0002 seconds
2024-12-12 17:14:14,040 - INFO - -----
2024-12-12 17:14:14,040 - INFO - Total time taken for all measurements: 4.0575 seconds
2024-12-12 17:14:14,040 - INFO - -----
2024-12-12 17:14:17,043 - INFO - Waiting 120 seconds for next measures ...
2024-12-12 17:16:17,193 - INFO - -----
2024-12-12 17:16:17,193 - INFO - Performing network measurements ...
2024-12-12 17:16:17,194 - INFO - Target host: 172.18.0.2
2024-12-12 17:16:17,194 - INFO - Target port: 80
2024-12-12 17:16:17,194 - INFO - -----

```

Figure 43 Logs from the network probes in the second cluster.

```

ubuntu@alto-server:~$ curl 127.0.0.1:5001/metrics
# HELP python_gc_objects_collected_total Objects collected during gc
# TYPE python_gc_objects_collected_total counter
python_gc_objects_collected_total{generation="0"} 882.0
python_gc_objects_collected_total{generation="1"} 416.0
python_gc_objects_collected_total{generation="2"} 0.0
# HELP python_gc_objects_uncollectable_total Uncollectable objects found during GC
# TYPE python_gc_objects_uncollectable_total counter
python_gc_objects_uncollectable_total{generation="0"} 0.0
python_gc_objects_uncollectable_total{generation="1"} 0.0
python_gc_objects_uncollectable_total{generation="2"} 0.0
# HELP python_gc_collections_total Number of times this generation was collected
# TYPE python_gc_collections_total counter
python_gc_collections_total{generation="0"} 272.0
python_gc_collections_total{generation="1"} 24.0
python_gc_collections_total{generation="2"} 2.0
# HELP python_info Python platform information
# TYPE python_info gauge
python_info{implementation="CPython",major="3",minor="9",patchlevel="18",version="3.9.18"} 1.0
# HELP process_virtual_memory_bytes Virtual memory size in bytes.
# TYPE process_virtual_memory_bytes gauge
process_virtual_memory_bytes 4.87686144e+08
# HELP process_resident_memory_bytes Resident memory size in bytes.
# TYPE process_resident_memory_bytes gauge
process_resident_memory_bytes 8.7719936e+07
# HELP process_start_time_seconds Start time of the process since unix epoch in seconds.
# TYPE process_start_time_seconds gauge
process_start_time_seconds 1.75402029826e+09
# HELP process_cpu_seconds_total Total user and system CPU time spent in seconds.
# TYPE process_cpu_seconds_total counter
process_cpu_seconds_total 40.01
# HELP process_open_fds Number of open file descriptors.
# TYPE process_open_fds gauge
process_open_fds 7.0
# HELP process_max_fds Maximum number of open file descriptors.
# TYPE process_max_fds gauge
process_max_fds 1.048576e+06
# HELP network_metric_throughput Network Metric for throughput (Mbps)
# TYPE network_metric_throughput gauge
network_metric_throughput{current_node="kind-worker",remote_ip="172.18.0.2",remote_node="kind-control-plane"} 20503.9
network_metric_throughput{current_node="kind-worker",remote_ip="172.18.0.4",remote_node="kind-worker2"} 20282.44
# HELP network_metric_link_energy Energy consumption (W)
# TYPE network_metric_link_energy gauge
network_metric_link_energy{current_node="kind-worker",remote_ip="172.18.0.2",remote_node="kind-control-plane"} 2706.26
network_metric_link_energy{current_node="kind-worker",remote_ip="172.18.0.4",remote_node="kind-worker2"} 2676.93
# HELP network_metric_packet_loss Network Metric for packet loss (percentage)
# TYPE network_metric_packet_loss gauge
network_metric_packet_loss{current_node="kind-worker",remote_ip="172.18.0.2",remote_node="kind-control-plane"} 0.0
network_metric_packet_loss{current_node="kind-worker",remote_ip="172.18.0.4",remote_node="kind-worker2"} 0.0
# HELP network_metric_latency Network Metric for latency (ms)
# TYPE network_metric_latency gauge
network_metric_latency{current_node="kind-worker",remote_ip="172.18.0.2",remote_node="kind-control-plane"} 0.58
network_metric_latency{current_node="kind-worker",remote_ip="172.18.0.4",remote_node="kind-worker2"} 0.38
# HELP network_metric_link_failure Link failure (0=okay; 1=not okay)
# TYPE network_metric_link_failure gauge

```

Figure 44 Metrics saved in the internal Prometheus in the first cluster.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	53 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b> Final

In order to show this last capture, we have performed a port-forwarding to the host system, since this service is only exposed within the nemo-net network, so that other systems or users cannot access them autonomously.

In parallel, the Nemex component performs periodic queries to this Prometheus, and integrates the metrics together with the information of the cluster and the time when the exposure is performed. In this way, the components that require this information will be able to access it.

In Figure 45 we can see how from RabbitMQ we can see the pushes to the mncc\_cmv queue (mNCC Cost Map Vectorial), where at the beginning we have a series of short peaks and after a while more elongated peaks appear. This is due to the fact that, at the beginning, only one cluster was launched, but when testing with several clusters, they started to send the metrics slightly out of sync. Also, in the upper part of the figure, the logs are shown from the two Nemex pods. Reading the first lines it is possible to check that the upper-left window shows the “kind-kind” cluster and the upper-right the “kind-kind-2” cluster. These clusters were launched in two different devices in the same lab, in order to check the correct behavior of the multi-cluster monitoring and exposure functionality.

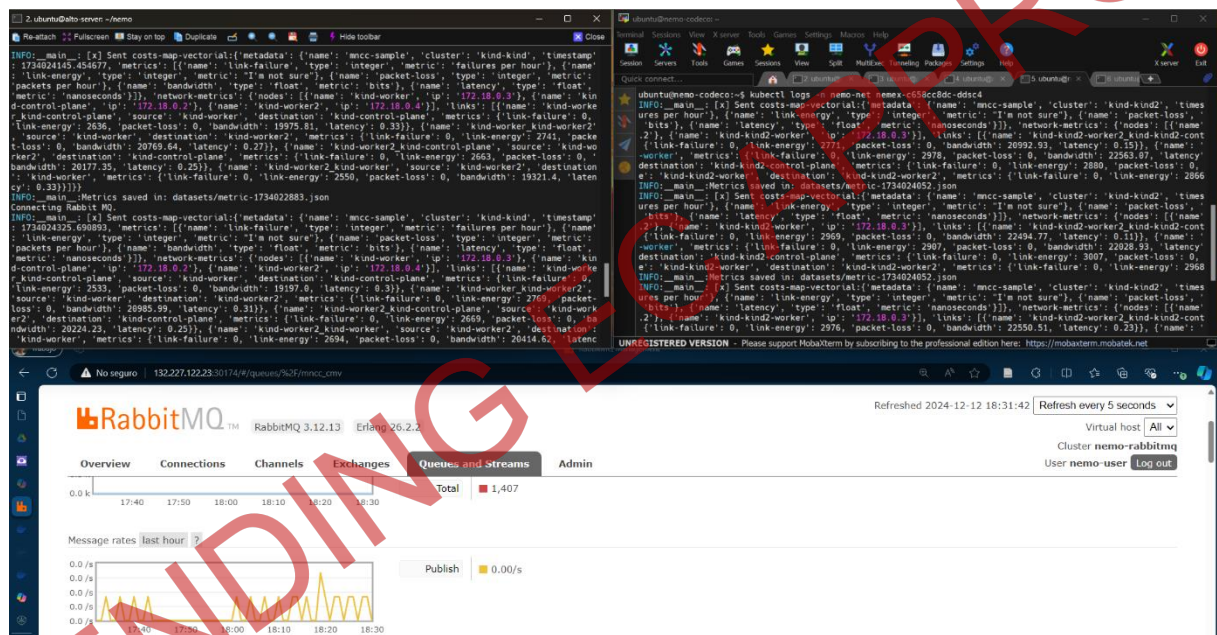


Figure 45 On the top, the logs of two Nemex components deployed in two clusters. In the bottom, the RabbitMQ updates.

These metrics in turn are accessible by the Meta Orchestrator and CF-DRL components to make decisions having knowledge of the state of the monitored network.

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	54 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
		<b>Version:</b>	1.0
		<b>Status:</b>	Final



## 7 Conclusions

---

This document reports the final outcomes of the components developed in Work Package 2, namely the CMDT, the CF-DRL and the mNCC. It also includes details on the TSN capabilities integrated on the NEMO edge-cloud continuum for satisfying specific use cases requiring time-sensitive behavior.

This deliverable summarizes the progress of the developments, the integration aspects with respect other components of the NEMO stack, as well as provides exemplary details of the functional behavior of the components.

Final integration efforts will be reported in deliverables from other work packages when required.

PENDING EC APPROVAL

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology			<b>Page:</b>	55 of 60		
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

## 8 References

---

- [1] N. w. partners, "NEMO\_D2.2-Enhancing-NEMO-Underlying-Technology\_v1.1.pdf," [Online]. Available: [https://meta-os.eu/wp-content/uploads/2024/09/NEMO\\_D2.2-Enhancing-NEMO-Underlying-Technology\\_v1.1.pdf](https://meta-os.eu/wp-content/uploads/2024/09/NEMO_D2.2-Enhancing-NEMO-Underlying-Technology_v1.1.pdf). [Accessed 15 10 2024].
- [2] "ETSI TeraFlowSDN Alignment With TIP OOPT MUST Requirements | TeraFlow," [Online]. Available: <https://www.teraflow-h2020.eu/publications/etsi-teraflowsdn-alignment-tip-oopt-must-requirements>. [Accessed 1 12 2024].
- [3] R. V. (. R. M. (. Pol Alemany (CTTC), "D4.1\_Preliminary evaluation of TeraFlow security and B5G network integration," [Online]. Available: [https://www.teraflow-h2020.eu/sites/teraflow/files/public/content-files/deliverables/D4.1\\_Preliminary%20evaluation%20of%20TeraFlow%20security%20and%20B5G%20network%20integration.pdf](https://www.teraflow-h2020.eu/sites/teraflow/files/public/content-files/deliverables/D4.1_Preliminary%20evaluation%20of%20TeraFlow%20security%20and%20B5G%20network%20integration.pdf). [Accessed 5 11 2024].
- [4] B. W. a. D. D. a. R. R. a. T. S. a. J. Mullooly, "A YANG Data Model for the RFC 9543 Network Slice Service," *IETF*, no. draft-ietf-teas-ietf-network-slice-nbi-yang-17, 2024.
- [5] J. Ramseyer, "Peering API," *IETF draft*, no. draft-ramseyer-grow-peering-api-00, 2024.
- [6] T. L. S. H. Y. Rekhter, "A Border Gateway Protocol 4 (BGP-4)," *Draft standard*, no. rfc4271, 2006.
- [7] G. D. a. C. F. a. K. T. a. M. C. a. D. B. a. B. Decraene, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing over IPv6 (SRv6)," *RFC editor*, no. RFC 9514, 2023.
- [8] "ETSI Software Development Group TeraFlowSDN - SDG TFS," [Online]. Available: [https://tfs.etsi.org/news/20230720\\_teraflowsdn\\_release\\_21/](https://tfs.etsi.org/news/20230720_teraflowsdn_release_21/). [Accessed 1 12 2024].

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	56 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
		<b>Version:</b>	1.0
		<b>Status:</b>	Final

## 9 Annexes

### 9.1 TFS interface and example intent

TeraFlowSDN controller input (python function) for descriptor file. L2VPN creation request.

```

vpn_descriptor={
  "services": [
    {"service_id": {
      "context_uuid": {"context_uuid": {"uuid": self.__params['context_uuid']}},
      "service_uuid": {"uuid": self.__params['service_uuid']} },
      "service_type": 2,
      "service_status": {"service_status": 1},
      "service_endpoint_ids": [
        {"device_id": {"device_uuid": {"uuid": self.__params['node_src']}}, "endpoint_uuid": {"uuid": "0/0/1-
"+self.__params['endpoint_src']}},
        {"device_id": {"device_uuid": {"uuid": self.__params['node_dst']}}, "endpoint_uuid": {"uuid": "0/0/1-
"+self.__params['endpoint_dst']}}
      ],
      "service_constraints": [
        {"custom": {"constraint_type": "bandwidth[gbps]", "constraint_value": self.__params['bandwidth']}},
        {"custom": {"constraint_type": "latency[ms]", "constraint_value": self.__params['latency']}}
      ],
      "service_config": {"config_rules": [
        {"action": 1, "custom": {"resource_key": "/settings", "resource_value": {
          "sub_interface_index": 0,
          "ni_name": self.__params['ni_name'],
          "vlan_id": int(self.__params['vlan_id']),
          "circuit_id": self.__params['circuit_id'],
          "remote_router": self.__params['node_dst']
        }},
        {"action": 1, "custom": {"resource_key": "/device["+self.__params['node_src']+"]/endpoint[0/0/1-
"+self.__params['endpoint_src']+"]/settings", "resource_value": {
          "sub_interface_index": 0,
          "ni_name": self.__params['ni_name'],
          "vlan_id": int(self.__params['vlan_id']),
          "circuit_id": self.__params['circuit_id'],
          "remote_router": self.__params['node_dst']
        }},
        {"action": 1, "custom": {"resource_key": "/device["+self.__params['node_dst']+"]/endpoint[0/0/1-
"+self.__params['endpoint_dst']+"]/settings", "resource_value": {
          "sub_interface_index": 0,
          "ni_name": self.__params['ni_name'],
          "vlan_id": int(self.__params['vlan_id']),
          "circuit_id": self.__params['circuit_id'],
          "remote_router": self.__params['node_src']
        }},
      ]}
    }
  ]
}

```

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	57 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
		<b>Version:</b>	1.0
		<b>Status:</b>	Final

**mNCC Intent: L2VPN request.**

```

Intent:
id: 'new_intent_2'
userLabel: 'cloud_continuum'
intentExpectations:
- expectationId: '1'
  expectationVerb: 'DELIVER'
  expectationObject:
    objectType: 'L2VPN'
    objectInstance: 'l2vpn_tfs_1'
    objectContexts:
      - contextAttribute: 'nodeSrc'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: '3.3.3.3'
      - contextAttribute: 'nodeDst'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: '5.5.5.5'
      - contextAttribute: 'endpointSrc'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: 'GigabitEthernet0/0/0/1'
      - contextAttribute: 'endpointDst'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: 'GigabitEthernet0/0/0/1'
      - contextAttribute: 'vlanId'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: '999'
      - contextAttribute: 'niName'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: 'ninametts'
  expectationTargets:
    - targetName: 'bandwidth'
      targetCondition: 'IS_EQUAL_TO_OR_GREATER_THAN'
      targetValueRange: 10.0
    - targetName: 'latency'
      targetCondition: 'IS_LESS_THAN'
      targetValueRange: 15.2
  expectationContexts:
    - contextAttribute: 'url'
      contextCondition: 'IS_EQUAL_TO'
      contextValueRange: 'http://192.168.165.10/restconf/data/ietf-l2vpn-svc:l2vpn-svc/vpn-services'
  intentContexts:
    - contextAttribute: 'NEMO_WORKLOAD'
      contextCondition: 'IS_EQUAL_TO'
      contextValueRange: 'cbb208a-d535-434b-bb35-217a64bd516b'
  intentPriority: 1
  observationPeriod: 60
  intentAdminState: 'ACTIVATED'

```

PENDING EC APPROVAL

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	58 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
		<b>Version:</b>	1.0
		<b>Status:</b>	Final

## 9.2 L2SM example intent

### mNCC Intent: L2SM network request.

Intent:

```

id: 'mncc_l2sm_1'
userLabel: 'cloud_continuum'
intentExpectations:
- expectationId: '1'
  expectationVerb: 'DELIVER'
  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: 'uc3m'
      - contextAttribute: 'domain'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: '172.18.0.2'
    expectationContexts:
      - contextAttribute: 'url'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: 'http://192.168.165.168:8080'
- expectationId: '2'
  expectationVerb: 'DELIVER'
  expectationObject:
    objectType: 'K8S_CLUSTER_CONFIG'
    objectInstance: 'spain-network'
    objectContexts:
      - contextAttribute: 'name'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: 'kind-worker-cluster-1'
      - contextAttribute: 'bearer_token'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: 'eyJhbGciOiJIUzI1NiIsImtpZCI6Iktnd1hiRk9RWGZBVkFVWU82Rk5...'
      - contextAttribute: 'api_key'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: 'https://172.18.0.3:6443'
    expectationContexts:
      - contextAttribute: 'k8s_l2_network'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: 'spain-network'
      - contextAttribute: 'url'
        contextCondition: 'IS_EQUAL_TO'
        contextValueRange: 'http://192.168.165.168:8080'
  intentContexts:
    - contextAttribute: 'NEMO_WORKLOAD'
      contextCondition: 'IS_EQUAL_TO'
      contextValueRange: 'c9cb208a-d535-434b-bb35-217a64bd516b'

```

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology	<b>Page:</b>	59 of 60
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU
		<b>Version:</b>	1.0
		<b>Status:</b>	Final

intentPriority: 1  
observationPeriod: 60  
intentAdminState: 'ACTIVATED'

PENDING EC APPROVAL

<b>Document name:</b>	D2.3 Enhancing NEMO Underlying Technology			<b>Page:</b>	60 of 60		
<b>Reference:</b>	D2.3	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final