Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or CINEA. Neither the European Union nor the granting authority can be held responsible for them.

Orchestration Data Space and initial version of scenario modelling results

Deliverable Number	D5.1
Deliverable Type	R – Document, Report
Dissemination Level	PU - Public
Deliverable Responsible	Kostas Cheliotis (FRON)
Document Version & Status	V1.0 | Final

DRAFT

Project Acronym	CARMONY
Project Title	Coordinated Action for Responsive Mixed Orchestration and Network Yield
Grant Agreement Number	101202858
Project Coordinator	Virtual Vehicle Research GmbH
Project Website	www.carmony-project.eu

The document is submitted to the European Commission for review but has not been accepted so far. As soon as the final version of this document is available, this document should be deleted.

Funded by the European Union

CARMONY has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101202858.

Author(s)

Name	Organisation
Kostas Cheliotis	FRON
Bernhard Hillbrand, Georg Nestlinger	ViF
Ana Martínez Roselló, Juan Carlos Cantos Agudo	ETRA
Pontus Svenson, Sara Janhäll, Anders Wallberg	RISE
Patrick van Egmond	LUX

Reviewers

Name	Organisation	Date
Juan Sánchez Valverde	MUR	2026-10-15
Dina Oulad Chair	HYP	2026-10-19

Change History

Version	Date	Name/Organisation	Description
0.1	2026-09-02	Kostas Cheliotis (FRON)	Added Table of Contents
0.2	2026-09-21	Bernhard Hillbrand, Georg Nestlinger (ViF)	Added Sections 4.1, 4.2, 4.3.3
0.3	2026-09-23	Ana Martínez Roselló, Juan Carlos Cantos Agudo (ETRA)	Added Section 4.3.1
0.4	2026-09-28	Patrick van Egmond (LUX)	Added Section 4.3.2
0.5	2026-10-06	Kostas Cheliotis (FRON)	Added Chapter 3
0.6	2026-10-07	Kostas Cheliotis (FRON)	Added Introduction, Conclusion, Executive Summary
0.7	2026-10-08	Kostas Cheliotis (FRON)	Proofreading, formatting
0.8	2026-10-27	All	Review revisions
1.0	2026-10-28	Katrin Walzer (ViF)	Final quality check and submission

DRAFT

Table of Contents

1. Executive Summary 5
2. Introduction 6
2.1 Objectives of this Deliverable 6
2.2 Links with WP5 objectives and other tasks 6
2.2.1 WP5 Objectives 6
2.2.2 Links with other Tasks 6
2.3 Structure of this Deliverable 7
3. Traffic Orchestration Data Space 8
3.1 Introduction 8
3.1.1 Previous work 8
3.1.2 Identification of data availability and requirements 9
3.1.3 Requirements 11
3.2 Implementation 12
3.2.1 Existing frameworks 12
3.2.1.1 Data Catalogues 12
3.2.1.2 Dataspaces 13
3.2.2 Proposed approach 14
3.2.3 Benefits 15
3.3 Next steps 16
4. Scenario Modelling 17
4.1 Simulation Scenarios 17
4.2 Regulatory Layer Scenarios 18
4.3 Simulation Setup 20
4.3.1 CARMONY API for Murcia site 20
4.3.2 CARMONY API for Luxembourg site 21
4.3.3 SUMO Simulation Setup 21
5. Conclusion 23
6. Abbreviations 24
7. References 25

DRAFT

List of Figures

Figure 1: CARMONY System architecture (from CARMONY D3.1, page 18) 9
Figure 2: Data sharing platform reference architecture 14

List of Tables

Table 1: Sample datasets, categories and utilizing components 10
Table 2: Data sharing platform requirements 11
Table 3: Use case simulation scenarios 17
Table 4: Regulatory layer scenarios 18

DRAFT

1. Executive Summary

This deliverable (D5.1) presents the progress so far for two key aspects involved in the development and implementation of the CARMONY Interlinked Orchestrator, the CARMONY data sharing platform and the simulation scenarios that the Orchestrator will be tested on. For both an initial version is presented: On the part of the data sharing platform, this deliverable presents the design and outlines the proposed solution, whereas for the simulation scenarios an initial set of scenarios covering a wide range of use cases is presented.

Regarding the data sharing platform, an initial inventory of data requirements is taken from existing CARMONY work and an analysis is performed to identify the core requirements that the data sharing platform will need to fulfil. Following that, existing widely used data management frameworks are evaluated against the core requirements established earlier, and a comprehensive solution is presented in more detail that addresses all of the requirements.

Regarding the simulation and modelling scenarios, two sets of simulation scenarios are presented, one focussing on orchestration scenarios during realistic traffic conditions and disruptions at fine scale, and another focussing on governance and regulatory policy effects at a larger scale. Additionally, a description of the simulation setup is presented, focussing on each use case as well as the traffic simulation engine.

This deliverable therefore places the initial stepping stones towards the development, implementation, and testing of the CARMONY Interlinked Orchestrator, as it establishes the reference architecture for the development of the data sharing platform, and it defines the set of scenarios that the Orchestrator will execute and evaluate.

Keywords: Dataspace, Data sharing platform, Scenario modelling and simulation

DRAFT

2. Introduction

2.1 Objectives of this Deliverable

This deliverable aims to achieve two core goals. The first is to present the initial design of the CARMONY data sharing system as it has been formulated so far. Building on work presented in previous deliverables and related activities, it captures the data needs and requirements and presents a high-level overview for the data sharing platform to be deployed in future months, acting as a guideline for its implementation. This goal links directly to the work performed so far under Task 5.1: Traffic orchestration data space.

The second core goal is to present an initial version of the scenario modelling approach that will be used by the CARMONY Orchestrator to assess the impact of traffic orchestration scenarios. This includes both an initial definition of road conditions scenarios to test the effects of traffic orchestration methods in realistic conditions, as well as scenarios relating to regulatory and governmental decisions to evaluate the effect of policy changes on road user behaviour as a whole. These initial scenarios aim to function as a roadmap for the consequent development of the CARMONY Interlinked Orchestrator, to guide the implementation of scenarios during deployment. This goal links directly to work undertaken so far in Task 5.4: Scenario modelling and simulation.

2.2 Links with WP5 objectives and other tasks

2.2.1 WP5 Objectives

This deliverable presents the initial versions and progress so far on the CARMONY data sharing platform and scenario modelling and simulation. As such it directly addresses a number of WP5 objectives as outlined in the CARMONY Grant Agreement.

First, one of the objectives of WP5 is to design and implement a data space, to be used for data sharing both internally (among CARMONY subsystems) and externally (to facilitate source data retrieval and deliver relevant outputs). This is directly fulfilled as part of Task 5.1 that designs and implements the CARMONY data sharing platform, the system that will facilitate data exchanges between the various subsystems of CARMONY. In this deliverable we present the initial design of the CARMONY data sharing platform, that will act as a roadmap and drive its development and implementation in the coming months.

A second objective of WP5 is to provide simulation models for evaluating the impact of traffic orchestration approaches. This is directly fulfilled in Task 5.4, that designs and defines the simulation models and scenarios that will be used to test the CARMONY Orchestrator, business models, and governance schemes. In this deliverable we define the initial simulation scenarios relating both to specific Use Cases as well as the Regulatory Layer that will model and assess business and governance policies.

2.2.2 Links with other Tasks

This report builds on a significant body of CARMONY work that has been undertaken and presented previously. First of all, Task 1.3: Ethics and Data Management has established the Data Management Plan for CARMONY that defines the guidelines for data management within CARMONY, and which the CARMONY data sharing platform presented here aims to follow. This will ensure that the resulting data sharing platform will adhere to FAIR principles and have provisions for secure data storage and sharing that is GDPR compliant.

DRAFT

Additionally, this report relies significantly on work undertaken as part of WP3, which defines the system architecture and designs and implements the various components for traffic management. On the one hand the system architecture defined in Task 3.1 provides an overview of the CARMONY system and highlights the overall data flows that are required for the operation of all components, and is therefore treated as the initial starting point to identify and define the requirements that the data sharing platform will need to fulfil. On the other hand, the rest of the WP3 Tasks handle the development of each of the specific components, defining their data inputs and outputs, and thus are paramount in identifying particular dataset format requirements that the data sharing platform will need to support.

This deliverable also has strong links with other WP5 Tasks, in order to achieve the overall goal of WP5 to implement and test the CARMONY Orchestrator. More specifically, Tasks 5.2 and 5.3 are expected to require data as input for the CCAM connectivity and control services as well as the user information services, which are expected to be provided via the data sharing platform. Furthermore, Task 5.5 will handle the design and implementation of the Interlinked Orchestrator, which will require both data inputs via the data sharing platform as well as use-case specific simulation scenarios for its calibration and customization. Finally, Task 5.6 will test the Orchestrator's capabilities for large-scale strategic planning, and therefore will require regulatory scenarios to be defined.

A final link is identified between this deliverable and WP6, which will demonstrate the applicability of the CARMONY Orchestrator via a number of pilot use cases. The demonstrations will require the deployment of the data sharing platform for data exchanges between all involved actors, as well as the definition of baseline and traffic orchestration scenarios in order to evaluate the Orchestrator's capabilities.

2.3 Structure of this Deliverable

This section provides an overall introduction, highlighting the objectives that this deliverable addresses, and how it links with other CARMONY Tasks. The rest of this report is structured as follows:

The following section (Section 3) will present the design of the CARMONY data sharing platform. It will review previous relevant CARMONY work that has influenced this design, provide an overview of the data needs identified so far in CARMONY, and establish the core requirements that the final data sharing platform will need to fulfil. Following that, it will review existing data management frameworks and evaluate their applicability for the CARMONY data sharing platform according to the core requirements identified earlier, present an architectural overview of the proposed implementation, and outline its benefits.

The next section (Section 4) will present the initial version of the scenario modelling approach. It defines a list of simulation scenarios based on relevant pilot Use Cases, as well as a list of simulation scenarios focussing on regulatory, governance, and business decisions. Additionally, this section presents the simulation setup that will be used for each of the two pilot sites, as well as the setup for the simulation engine that will be used to execute the traffic simulations.

Finally, Section 5 concludes the main body of the report with a summary, highlighting its main achievements and next steps.

DRAFT

3. Traffic Orchestration Data Space

In this chapter we present the data sharing platform that will facilitate the exchange of data throughout the CARMONY systems and partners. We will first review existing work from CARMONY, and based on those we will identify and define the requirements and needs that the data sharing platform will need to fulfil. Following that we will describe existing data management systems and evaluate them against the defined requirements, before presenting our proposed design for the data sharing platform.

3.1 Introduction

3.1.1 Previous work

In order to design the data sharing platform that is the Traffic Orchestration Data Space for CARMONY, we first need to identify the core requirements that it will need to fulfil. To this end we have reviewed relevant work and documentation from the CARMONY project that has been carried out and published previously. These materials include the CARMONY Grant Agreement that provides the general outline and scope of the data sharing platform, deliverable D1.2 that outlines the Data Management Plan (DMP), outputs from a workshop during the 2nd CARMONY General Assembly involving all partners, as well as ongoing work from WP3 and particularly Task T3.1 regarding the CARMONY system architecture, including both the initial work published in deliverable D3.1 and all work progress and updates that have followed.

More specifically, the CARMONY Grant Agreement outlines the scope of the work, which is to “Implement a comprehensive data space for CARMONY, compliant with IDSA standards and incorporating FIWARE components, supporting secure, scalable storage, seamless data integration, and interoperability between systems validated through pilot projects”, resulting in a “Data sharing service for the pilot sites, including protocols and data interfaces.”

Deliverable D1.2 provides the initial version of the Data Management Plan (DMP), which includes multiple points relevant to the design of the data sharing platform, that we highlight here. In addition to a summary of datasets expected to be involved in each pilot, the DMP identifies multiple data sensitivity levels along with the required access levels that will need to be implemented. Furthermore, it identifies strategies for collecting, storing, and making available the various datasets that will be handled by the data sharing platform.

During the second General Assembly of CARMONY held at Bilbao on May 12th 2026, an in-person workshop was held that included all technical and pilot demonstration partners, where each partner identified all datasets that they can either provide or that they will require for the operation of their components.

Finally the work undertaken in WP3 involves the design and development of the majority of technical components that will require data as input and/or generate data as output. As such the discussions originating from within WP3 have been taken into account here as they illuminate a lot of different facets regarding system operation. These findings have been presented in deliverable D3.1, where a diagram of the system architecture is presented outlining the main components and their relationships. We reproduce here the CARMONY System Architecture diagram from D3.1 for reference, shown in Figure 1.

DRAFT

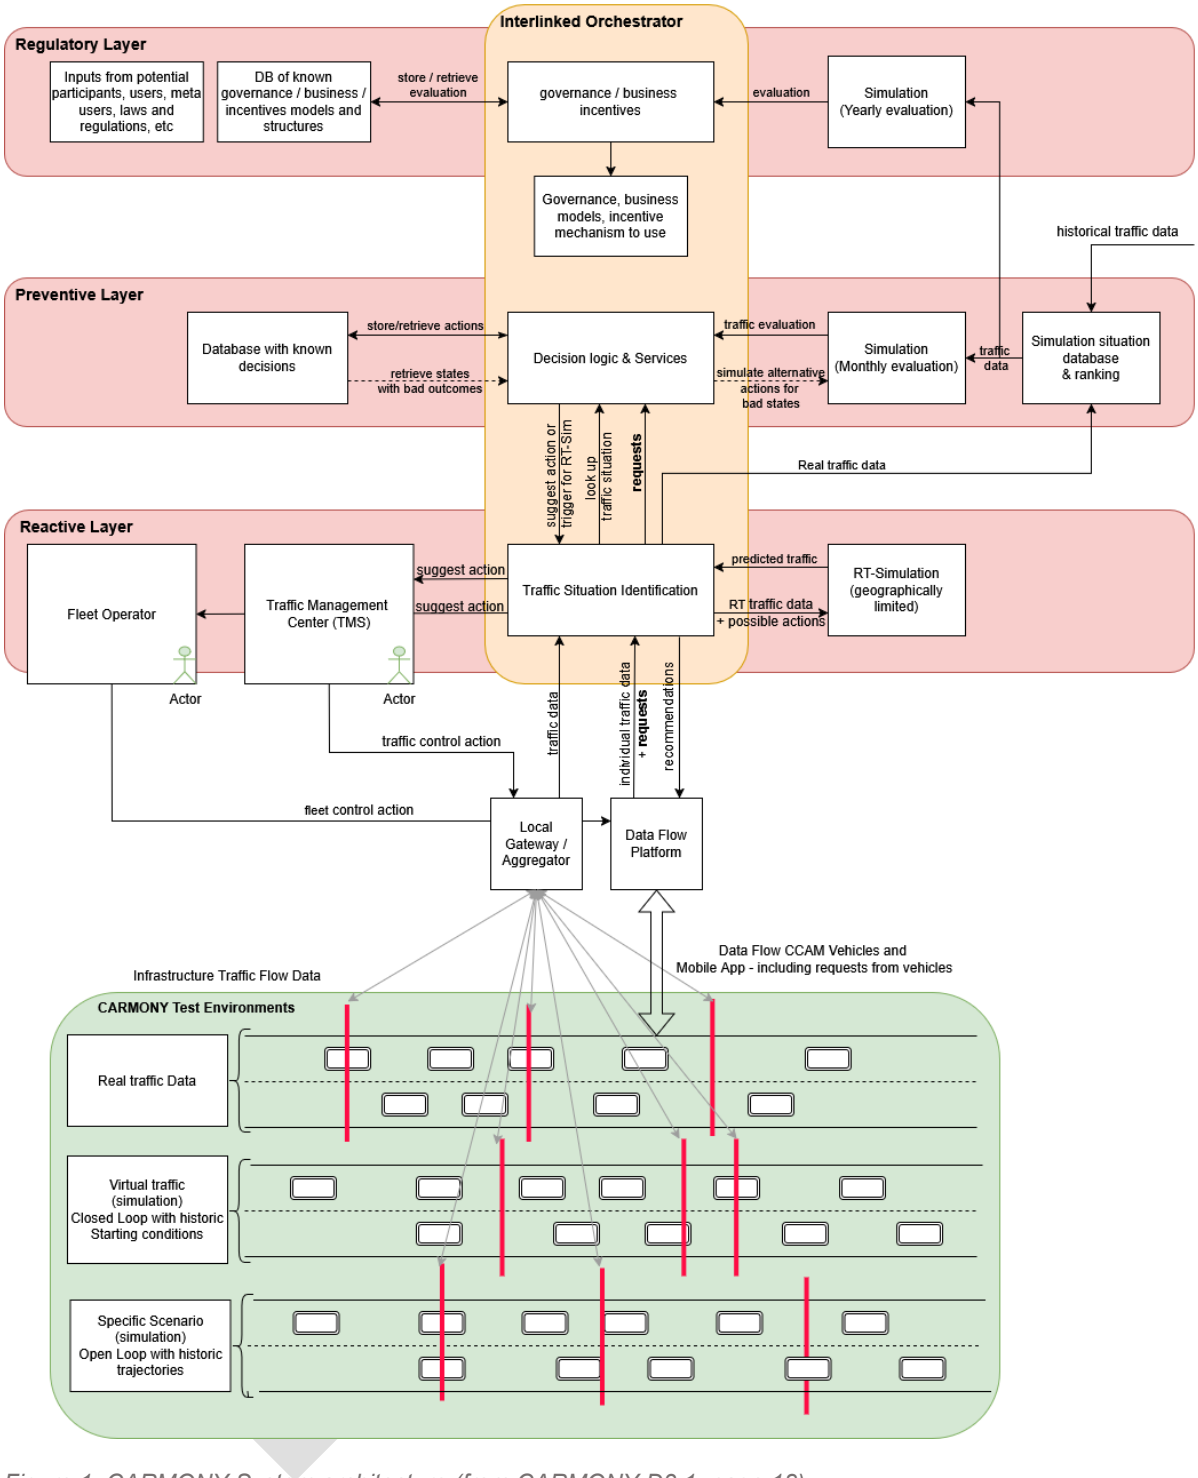


Figure 1: CARMONY System architecture (from CARMONY D3.1, page 18)

3.1.2 Identification of data availability and requirements

Following the Data Exchange Requirements as defined previously in D1.2, the CARMONY System architecture outlined in D3.1, the outcome of a partner workshop undertaken during the CARMONY 2nd General Assembly on May 12th 2026, as well as the ongoing work of WP3, at this point we have identified the majority of datasets that will be produced and consumed throughout the CARMONY System. Furthermore, after analysing the data requirements we have identified some thematic relationships between the datasets that enable the definition of high-level dataset categories, which have then been used to inform the design of the data sharing platform we will be presenting in the following sections.

We have identified four main categories for the datasets involved in the CARMONY system. These are:

1. Static datasets, that for example include historic traffic data used for training and calibrating algorithms and road network environment graphs used for executing various traffic management algorithms and simulations.
2. Real-time input data, that is expected to be available via webAPIs, and includes for example real-time traffic and incident data originating from sources external to CARMONY.
3. Operational consumable data, that includes all data originating from CARMONY components that is expected to be consumed soon after by another CARMONY component, i.e. are intermediate datasets only used during operation, and may include for example alternative action plans to evaluate via simulation or CCAM information.

- Operational databases, which include data generated during CARMONY system operation but are expected to be of long-term value, such as past scenarios and viable recommendations.

These four categories defined above capture different aspects of the dataset lifetime within CARMONY, and inform what role each dataset is expected to fulfil. More specifically, static datasets involve all data that has been generated or collected once and is not expected to ever change (or may be updated very infrequently), that is required for the operation of multiple components, and thus the most updated version of each of these datasets needs to be readily available to all partners. Real-time input data is external data required for the real-time operation of CARMONY components, whose origin is not maintained by CARMONY and thus needs applicable data retrieval processes. Operational consumable data includes datasets that will be generated and consumed within a short time span and is not expected to be made available to external recipients, but is rather only used for the generation of the final recommendations and outputs, and thus does not need to be stored afterwards. Finally operational databases include datasets that are expected to grow over the duration of the CARMONY Orchestrator operation, and we have identified it to be of potential use for future work as well, and thus needs to be stored and made available long-term.

A summary and examples of the datasets, their categories, and components requiring them are shown in Table 1. The datasets have been identified by CARMONY partners during an in-person workshop during the 2nd General Assembly on May 12th 2026, and include all datasets identified as relevant and required up to that point. While we anticipate that additional dataset requirements may come up in the future during application and implementation, the four dataset categories that we have defined here will be able to accommodate those, and therefore regarding the data sharing platform requirements we are confident that the categories defined here encompass all needs for current and future dataset needs.

Table 1: Sample datasets, categories and utilizing components

Dataset	Category	Component
historic traffic	static	decision logic
		regulatory layer
		fleet operator
road environment	static	decision logic
		route planner
baseline scenarios	operational databases	decision logic

DRAFT

recommended actions	operational databases	decision logic
		regulatory layer
route planning	operational consumable	fleet operator
		regulatory layer
recommended actions	operational consumable	simulations
real-time events	real-time input	fleet operator
real time traffic	real-time input	fleet operator
		traffic situation identification
		simulations
real-time sensors	operational consumable	HD map tiling
vehicle counts	real-time input	decision logic
historic incidents	static	decision logic
weather conditions	real-time input	decision logic

3.1.3 Requirements

Based on the findings discussed above, we define here the core requirements that the CARMONY data sharing platform will need to fulfil, along with their justification. First of all, **static datasets should be made available to all partners**. Given that static datasets have a one-to-many relationship due to different partners all using the same dataset, all partners need to have easy access to the most up-to-date version of a static dataset. Secondly, the data sharing platform **should support webAPIs and similar**, as a significant portion of real-time data is expected to be retrieved via webAPIs. Third, the data sharing platform **should be able to implement access restrictions**, as we expect some of the datasets to contain privileged information and therefore access to them should be able to be regulated. Fourth, the data sharing platform **should be compliant with standards and best practices**, to ensure security and interoperability. Moreover, the data sharing platform **should require minimal technical overhead from partners** for its implementation. Given the number of technical partners and data exchanges expected to be part of the development of the CARMONY system, as well as its deployment during pilots, we establish this minimal technical overhead requirement to facilitate smooth and straightforward implementation, and to minimize future problems during pilot deployment. Furthermore, the data sharing platform should consider the data needs after the end of the project, and to **support long-term data availability**, where needed. Finally, as identified initially in the Grant Agreement, the data sharing platform should **support secure, scalable storage, seamless data integration, and interoperability between systems**. The requirements identified here are also presented more concisely in

Table 2.

Table 2: Data sharing platform requirements

Index	Name	Description	Source
R1	static dataset availability	static datasets should be made available to all partners	D3.1
R2	RestAPI support	the data sharing platform should support RestAPIs and similar	D3.1
R3	data access control	the data sharing platform should be able to implement access restrictions	D1.2/GA

DRAFT

R4	standards and best practices compliance	the data sharing platform should be compliant with IDSA standards and incorporating FIWARE components	GA
R5	minimal technical overhead for participants	the data sharing platform should require minimal technical overhead from partners for its implementation	partner workshop
R6	long-term data availability	valuable datasets should be made available to authorized interested parties after the end of CARMONY	D3.1/partner workshop
R7	scalable	the data sharing platform should support scalable storage	GA
R8	interoperable	the data sharing platform should support seamless data integration and interoperability between systems	GA

3.2 Implementation

3.2.1 Existing frameworks

We consider here two existing candidate solutions for the data sharing platform, a data catalogue system, and a dataspace system. In this section we evaluate both against the requirements specified above, highlighting their benefits and limitations in reference to the CARMONY data sharing platform needs.

3.2.1.1 Data Catalogues

As data catalogue we define a data management system that operates as a central data hub and provides tools for easy access and publication of data, either by hosting the data in a central repository or by linking to individual contributors' repositories. A widely used system for data catalogues is the open-source CKAN system[1], used across multiple international organizations for the management of governmental open data portals. As such we consider it a representative solution of data catalogues, and we therefore will focus our evaluation of data catalogue solutions on the CKAN system.

CKAN provides a data portal with direct links to datasets and data sources, and therefore directly fulfils requirements for static and webAPI data access. Furthermore, given both its widespread use and adoption by governmental bodies, we acknowledge that it supports best practices and standards. Additionally, CKAN supports user roles and authorization, and therefore it provides tools for placing restrictions on data access where needed. Given its core approach where organizations publish datasets which are then generally available to all other participating users, we consider the CKAN system to not require extensive technical knowledge from CARMONY partners to connect and access required datasets. Finally, CKAN supports a wide extensions repository, that allows additional functionalities to be included in a particular CKAN implementation, which we expect will be able to cover additional requirements, such as the CKAN FIWARE extensions[2], that allow for the integration of CKAN within the FIWARE ecosystem.

[1] https://ckan.org/
[2] https://fiware-ckan-extensions.readthedocs.io/

One potential drawback we identify in CKAN is regarding its control and maintenance: CKAN deployments require administrator users with full privileges in order to handle users, datasets, and access. While this is beneficial during the lifetime of the CARMONY project, we identify that this might pose a potential issue in the long-term. As identified previously, we expect that some datasets generated during CARMONY might be of value even after the conclusion of CARMONY. Therefore the data sharing platform implementation should also include provisions for long-term availability of data to external partners that we are not currently aware of. A CKAN approach with its active admin role requirements might pose a limitation here, as it would require some CARMONY partners to continue operating as CKAN admins, evaluating new user access to datasets, particularly when the datasets might be of restricted access.

3.2.1.2 Dataspaces

Another approach that we consider for the data sharing platform is a dataspace, which we define here as a federated system that facilitates connections between data providers and consumers without the need for handling the datasets themselves, as data transfers are performed via a direct connection between provider and consumer (Curry, 2020). The dataspace's roles are primarily to provide tools and systems for dataspace participants to authenticate themselves, often against a trusted third identity provider, to allow participants to form data sharing agreements based on their authenticated credentials, and to provide tools that will handle the end-to-end connections for the data transfer.

Data transfers in a dataspace only take place after a data sharing agreement has been established between two participants, which in itself requires that both participants have been authenticated by a trusted third-party, and therefore by design a dataspace provides a secure system. Additionally, each data provider publishes a list of datasets available (both for static data and for real-time webAPIs), so after initial set-up all CARMONY partners would be able to have access to the required datasets. Furthermore, one of the most widely used frameworks, the Eclipse Dataspace Components (EDC)[3] implement the Dataspace Protocol (DSP) defined by the International Dataspaces Association (IDSA), and therefore are compliant with standards and best practices. Additionally, one of the major benefits of dataspaces for CARMONY is their federated nature: As all dataspace participant connections to one another are facilitated via a central handler (called the control plane in EDC), it is relatively straightforward to extend connections from one dataspace to another, allowing datasets from one provider of one dataspace to be visible to participants in other dataspaces, without significant additional work. This presents a significant benefit for the long-term support and sharing of relevant datasets to potential future users: After the conclusion of CARMONY, any datasets deemed as potentially relevant for long-term use may be published to other international mobility dataspaces (Pretzsch et al., 2022) and made available to future interested organizations, whose credentials identify them as capable of handling them. One such dataspace highly relevant for CARMONY datasets is the European Mobility Dataspace (EMDS)[4], a plan to "create a trusted environment for sharing and accessing mobility data to

[3] https://dataspace.eclipse.org/

[4] https://regions-and-cities.ec.europa.eu/cities-portal/eu-initiatives-cities/common-european-mobility-data-space_en

DRAFT

support safer, smarter, and more sustainable transport services", whose deployment is driven by the EU-funded project deployEMDS[5].

While dataspaces (and particularly the EDC framework) appear to capture most of the needs of the CARMONY data sharing platform, we note one significant drawback relating to minimizing technical overhead from CARMONY partners. Dataspaces require significant set up beforehand, with each partner that needs to access the dataspace (either as a data provider or a consumer) needing to set up specialized tools (called connectors) that will allow them to participate in and transfer data over a dataspace. While this approach would work significantly towards the longevity of CARMONY datasets, especially given the move towards dataspaces in the European and International mobility spaces, it would nevertheless impose additional workload to all CARMONY partners before system deployment and during testing.

3.2.2 Proposed approach

Given that no option fully captures all requirements for the CARMONY data sharing platform, we propose a combined solution here: A CKAN data catalogue will be implemented to handle the operational data needs of CARMONY, while a small-scale dataspace will be implemented in parallel to facilitate dataset longevity. We illustrate the proposed high-level architecture of the CARMONY data sharing platform in Figure 2.

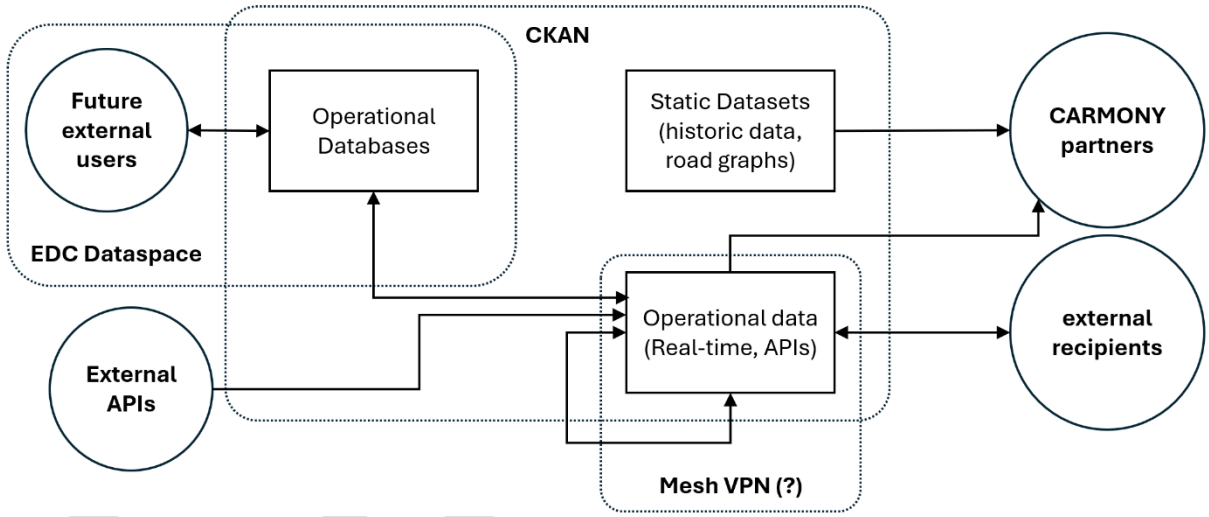


Figure 2: Data sharing platform reference architecture

More specifically, we rely on CKAN's ease of use and straightforward onboarding for all participants to set up a platform that all CARMONY partners will be able to use, share, and retrieve data from, ensuring that all operational data is readily available, with additional access restrictions where required. The majority of the set up and technical work required (setting up the CKAN instance, maintenance and admin) can be carried out by one technical partner, and all other technical and pilot demonstration partners are only required to set up a user and organizational account on the CKAN platform before receiving access to the various operational datasets, therefore minimizing technical overhead for all. Afterwards, all data

[5] https://deployemds.eu/

sharing regarding static data, real-time input data, and operational consumable data is performed via the CKAN, where all CARMONY partners have direct access.

As an additional facilitating technology we are also considering the implementation of a mesh VPN network, to handle connections between remote webAPI interfaces for internal CARMONY components. This technology is being considered for the case where each partner makes their components available via a RestAPI hosted on their own servers and the expected data pipeline during CARMONY system operation is expected to link through many of these components. If such a case is true, then we propose the implementation of such a technology (e.g. using the Tailscale[6] solution) to handle only the connections between internal CARMONY APIs. This approach is still under consideration however, and therefore we mention it here only as a potential additional layer.

In parallel to the CKAN, we propose setting up a dataspace using the EDC framework, to jump-start the long-term availability of viable CARMONY generated knowledge. This dataspace will include within its scope only the datasets deemed as relevant for long-term availability, namely the operational databases as identified earlier. This dataspace implementation will have two main aims: first, to demonstrate the viability of implementing a dataspace in cooperation to a CKAN data catalogue, and second, to bootstrap the processes for making CARMONY-generated knowledge available long-term and preparing the relevant datasets for publication in other mobility dataspaces. The majority of CARMONY partners will not be expected to participate in the dataspace, but rather the setup and testing of the dataspace will be carried out by a handful of technical partners.

3.2.3 Benefits

In previous sections we have defined the requirements that a CARMONY data sharing platform would need to fulfil, evaluated existing data management system frameworks against these requirements, and proposed a combined solution that fulfils all of the requirements. In this section we highlight here the benefits of the proposed architecture, to illustrate why we consider this the optimal approach regarding the needs of the CARMONY data sharing platform.

First of all, the majority of the day-to-day and operational data needs within CARMONY will be addressed via the CKAN data catalogue, which requires minimal set up from the majority of the partners (thus addressing R5), and should therefore enable a straightforward implementation while minimizing deployment issues. The CKAN will provide direct access to the latest versions of all static datasets, fulfilling R1. Additionally, the CKAN will also provide direct links and documentation regarding the use of internal component endpoints enabling all CARMONY development teams to have access to any required APIs (alternatively if we opt for the mesh VPN approach for the APIs, CKAN will still provide documentation, but connections will be routed through the VPN). This set up will ensure that all recipients, both internal (e.g. CARMONY components) and external (e.g. CCAM and CARMONY system users and stakeholders) have direct access to outputs as needed, fulfilling R2.

At the same time a small team of CARMONY technical partners will undertake the development and testing of an EDC dataspace focussing on making selected operational databases

[6] https://tailscale.com/

available within a dataspace. This will prepare the selected datasets for inclusion in larger federated mobility dataspaces during and after CARMONY project conclusion (thus fulfilling R6), taking into account any restrictions that may be imposed on access (as dictated by R3). Furthermore, by implementing this flow we will have an established process for additional datasets considered viable for long-term availability, if such are identified later during the CARMONY project. As this EDC dataspace will be implemented in parallel to the CKAN, it will not interfere with normal CARMONY operations, nor it will require additional effort from the majority of partners.

This proposed combined setup of a CKAN data catalogue and an EDC dataspace also fulfils the R4 requirement identified earlier, regarding compliance with IDSA standards and incorporating FIWARE components, with the CKAN data catalogue incorporating FIWARE-compliant extensions, while the EDC connector is fully compliant with IDSA standards. Furthermore, both the CKAN and EDC approaches constitute modern frameworks aiming to address current needs in data management and sharing, and are therefore fully capable of enabling scalable and interoperable solutions, therefore addressing requirements R7 and R8.

3.3 Next steps

At this point we have presented an overview of the CARMONY data sharing platform, taking into account all requirements as identified from previous CARMONY work, and have established the benefits of the proposed approach and how they address the requirements. We present here a work plan for the implementation of the above mentioned platform, and the steps needed. Development for the implementation of the data sharing platform will commence on M19 and will continue until the end of Task 5.1 on M33. The final version of the CARMONY data sharing platform will be reported in deliverable D5.4, due in M33.

The primary focus will be on developing and deploying the CKAN data catalogue, which will be performed in two broad phases: First a CKAN instance will be deployed and set up with a minimal subset of datasets from each category, requiring the input and contribution from only one or two additional partners, to verify that the approach is sound. Following that, an onboarding guide will be prepared and circulated to all CARMONY partners for them to join the platform and make their data available, and the CKAN data catalogue will be implemented in full. At this stage, the CKAN platform will be made available to be validated via the Murcia and Luxembourg pilots.

At the same time an EDC dataspace will be deployed with the participation of a handful of technical partners (we estimate two to three partners to be involved in this), to test and validate the dataspace implementation. All throughout the CARMONY project we will be evaluating datasets for their viability as long-term datasets, adding them to the EDC dataspace when required. Finally, at a later stage in the project we will identify the partners that will be responsible for hosting and facilitating the availability of CARMONY data in the long term via connection with other mobility dataspaces, allowing for CARMONY-generated knowledge to be preserved after the end of the CARMONY project.

DRAFT

4. Scenario Modelling

This chapter presents the scenario modelling approach adopted in CARMONY to assess the potential impact of traffic orchestration services and regulatory interventions across different mobility contexts. The selected scenarios represent the CARMONY use cases from urban and cross-border traffic management (UC1, UC2, UC4 and UC6), focusing on how orchestration mechanisms can improve traffic efficiency, support policy objectives, and enhance multimodal mobility services in simulation. UC3 and UC5 are scheduled for real-world testing only and will not be covered in this chapter.

First, a set of simulation scenarios is defined to evaluate the operational performance of the CARMONY orchestrator under realistic traffic conditions and network disturbances. Building on these baseline scenarios, additional regulatory layer scenarios are introduced to investigate the influence of policy measures, governance decisions, infrastructure configurations, and user behaviour on the effectiveness of orchestration services. These variations provide a broader understanding of the interactions between technical solutions and regulatory frameworks.

To ensure a consistent and reproducible evaluation process, the chapter also outlines the simulation environment, including the modelling tools, interfaces, data exchange mechanisms, and integration of the CARMONY orchestration components. Together, these elements establish the foundation for the subsequent assessment of project use cases and key performance indicators.

4.1 Simulation Scenarios

Table 3: Use case simulation scenarios

Scenario ID	Description	Related KPIs
UC1_S1	City of Murcia: Traffic with • Origin and Destination of some connected vehicles (CVs) are outside the LEZ: the orchestrator provides a route avoiding the LEZ. • Origin and/or Destination of some other CVs are inside the LEZ and the CVs are allowed: the orchestrator suggests the best route. • Origin and/or Destination of some CVs are inside the LEZ and the CV is not allowed: the orchestrator suggests alternative transport options and guides CVs to a convenient transfer point. Note: Variations with traffic situations that are representative for different times of the day.	KPI1, KPI2, KPI3, KPI4, KPI5, KPI8
UC2_S1	City of Murcia: Road blockage at an intersection due to a bus incident. CVs need to be redirected. Note: Variations with different intersections and/or traffic situations.	KPI11, KPI12, KPI14
UC4_S1	A3 highway between France and Luxembourg City with a dedicated carpool and bus lane on the left side: Non-Carpool Vehicle Access Optimization, dynamically authorizing commuters to use the carpool lane to balance traffic loads.	KPI25, KPI26, KPI27, KPI28,

		KPI29, KPI30, KPI31
UC4_S2	A3 highway between France and Luxembourg City with a dedicated carpool and bus lane on the left side: Freight Slot Allocation, where drivers benefit from dedicated access to improve cross-border freight flow.	KPI25, KPI26, KPI27, KPI28, KPI29, KPI30, KPI31
UC6_S1	A3 highway between France and Luxembourg City: Dynamic Speed Limit approach, variable speed signs display adaptive recommendations to all drivers, stabilizing the flow.	KPI37, KPI38, KPI39, KPI40, KPI41, KPI42

The baseline is always the same scenario without the orchestrator's intervention.

The KPIs refer to the CARMONY deliverable D2.4, in which the KPI have been defined and described in detail.

4.2 Regulatory Layer Scenarios

The scenarios will test the orchestrator, business models, and governance schemes. The terminology related to business models have been further investigated in CARMONY and is based on the conclusions from D4.2. Thus, the focus of the scenario testing is the business model system that is built on the business models of all different actors together. The business model system is further divided into the business logic describing how each actor makes decisions, and the business model related to each actor describing the way each actor earns return and experience costs. In this context the business model system is tested in the regulatory layer scenarios.

The regulatory layer scenarios are formed to test the viability of the viability of the orchestrator approach, that in D4.2 is referred to as the business model system. This system includes aspects of decision making, mechanism design and incentive models that will be further described in D4.5. The scenarios described in Table 3 will be used to test how the orchestrator wants to solve the challenges in the transport system of both Murcia and Luxembourg, together with a description of the resulting actions and the new distribution of traffic.

Table 4: Regulatory layer scenarios

Scenario ID	Description
UC1_RL_S1	City of Murcia: Extension of UC1_S1: LEZ extent. Two or three candidate boundaries; the output is the traffic redistribution and the loading on boundary roads, which shows how the orchestrator's routing shifts flow rather than simply removing it.
UC1_RL_S2	City of Murcia: Extension of UC1_S1: Time-dependent restrictions. Entry restrictions active only at peak hours; the output is the change in peak entries and how the orchestrator's advice (re-timing, re-routing, park-and-switch) is distributed across trip types.

DRAFT

UC1_RL_S3	City of Murcia: Extension of UC1_S1: Instrument comparison. The same emission target inside the zone is approached by (i) orchestrator advice to avoid the zone, (ii) a parking-price change, (iii) a signal strategy favouring the ring road. The output is which instrument reaches the target and at what cost in travel time, illustrating that the orchestrator coordinates several instruments rather than one.
UC1_RL_S4	City of Murcia: Extension of UC1_S1: CV/CAV penetration rate applied to all of the above.
UC1_RL_S5	City of Murcia: Extension of UC1_S1: Compliance rate. How many of the vehicles follow the advice they are given by the orchestrator? (This is more relevant for the orchestration than the connected vehicle penetration rate!)
UC2_RL_S1	City of Murcia: Extension of UC2_S1: Priority ordering during an incident. One blockage; the orchestrator applies alternative priority orders among public transport, emergency and recovery vehicles, and general traffic. The output is delay by actor class (ART1 agents in traffic, ART2 fleet operators, public transport), showing that the orchestrator's recommendations serve several actors with different objectives.
UC2_RL_S2	City of Murcia: Extension of UC2_S1: Alternative bus routes. For the same blockage, two or three pre-defined diversion routes for public transport; the output is passenger delay and the effect on general traffic along the diversion.
UC2_RL_S3	City of Murcia: Extension of UC2_S1: Responsibility allocation. The same blockage handled under two contract variants for who decides what (traffic management system, public transport operator, mobility service provider). Single-period output is response time and delay; this shows the multi-actor coordination the orchestrator is designed for.
UC2_RL_S4	City of Murcia: Extension of UC2_S1: CV/CAV penetration rate. Applied to all of the above, since compliance with the detour recommendation is the mechanism by which the orchestrator relieves the blockage.
UC2_RL_S5	City of Murcia: Extension of UC2_S1: Compliance rate. How many of the vehicles follow the advice they are given by the orchestrator? (This is more relevant for the orchestration than the connected vehicle penetration rate!)
UC4_RL_S1	A3 highway between France and Luxembourg City: Lane capacity and length. Vary the lane's length and capacity against a given demand profile (a representative peak period); the output is lane throughput, relief in the general lanes and queue formation at the merge.
UC4_RL_S2	A3 highway between France and Luxembourg City: Lane opening policy. Compare a permanently open lane with time-of-day opening (peak only) and with opening on the orchestrator's recommendation when observed loading exceeds a threshold; the output is throughput and delay for carpool and non-carpool traffic.
UC4_RL_S3	A3 highway between France and Luxembourg City: Single-occupancy access through booking. Allow single-occupancy vehicles into the lane via the orchestrator's booking service, with a fixed daily quota; vary the quota. The output is lane throughput, general-lane relief and the utilisation of the booking capacity. Booking is a central orchestrator service and this scenario demonstrates it in its simplest form.

UC4_RL_S4	A3 highway between France and Luxembourg City: Compliance and enforcement. Vary the share of unbooked or non-eligible entries and the enforcement rate; the output is the sensitivity of lane throughput to non-compliance.
UC6_RL_S1	A3 highway between France and Luxembourg City: Activation threshold. Vary the traffic-state threshold (flow, density or speed) at which the orchestrator activates speed advice; the output is whether the jam forms, and total delay.
UC6_RL_S2	A3 highway between France and Luxembourg City: Advised speed. Vary the advised speed at a fixed activation threshold; the output is flow, delay and the number of stop-and-go events.
UC6_RL_S3	A3 highway between France and Luxembourg City: Advisory versus legal limit. Compare the orchestrator's speed recommendation, with a compliance rate assumed for the peak period, against a variable legal speed limit with near-full compliance. This is the recommend-versus-command comparison in a single-period form and shows what compliance level the advisory approach needs to match a mandatory one.
UC6_RL_S4	A3 highway between France and Luxembourg City: CV/CAV penetration rate. Advice reaches only connected vehicles; vary the share and observe at what penetration the advised speed propagates to the whole flow.
UC6_RL_S5	A3 highway between France and Luxembourg City: Compliance rate. How many of the vehicles follow the advice they are given by the orchestrator? (this is more relevant for the orchestration than the connected vehicle penetration rate!)

4.3 Simulation Setup

4.3.1 CARMONY API for Murcia site

The main workflow between the CARMONY API and the rest of services in the Murcia pilot, including the simulation, covers route retrieval, journey creation, and the establishment of real-time communications to receive traffic-light, incident and low emission zones information.

The data provided includes routing information (route geometries, route identifiers, distances, and alternative routes), journey-related data (journey IDs, stop nodes, intersections, and associated traffic lights), and real-time traffic-light information exchanged through a Gateway WebSocket Service. The platform also distributes traffic incidents and Low Emission Zone (LEZ) information through a dedicated Events WebSocket. Data is primarily exchanged in JSON format, while route geometries will potentially be represented using GeoJSON structures (LineString with coordinate arrays in [longitude, latitude] format), final format will be defined in the upcoming months. Incident information follows the CIFS standard and includes attributes such as identifiers, timestamps, location, type, description, validity period, and affected road segments. LEZ information includes polygon geometries, schedules, descriptions, and time-zone information.

Additional points of interest include support for authentication and authorization, detailed error handling, journey recovery and reconnection mechanisms. For integration purposes, a publicly available OpenAPI specification and AsyncAPI specification are provided. In addition, a versioned document containing the definitions of all available endpoints has been shared by

email with the partners involved in the integration. The document also outlines planned future enhancements, including authenticated incident creation and confirmation, as well as the publication and dissemination of emergency vehicle position data through the platform.

4.3.2 CARMONY API for Luxembourg site

For the Luxembourg pilot, the data exchange supporting the CARMONY simulation and orchestration activities focuses on two motorway corridors, the A3 and A13. For both corridors, historical and real-time traffic information is available, including vehicle flow and average speed, with real-time updates typically provided at intervals of approximately 30 seconds to one minute. Historical data are available from 2018 onwards and use ISO 8601 timestamps, while vehicle classification currently distinguishes cars and trucks. Incident and traffic degradation information is also available. On the A3, particular attention is given to the carpooling lane, for which historical traffic flow and an estimated theoretical capacity are available, while more detailed information on effective capacity, eligible vehicle categories and actual carpool use remains to be further developed. On the A13, the available information additionally covers roadworks, lane or road closures and traffic incidents, with location and temporal information allowing events to be associated with specific road segments and directions. Data are provided through the Luxembourg CITA traffic information services, with DATEX II used for available traffic-status information and JSON identified as the preferred exchange format for CARMONY integration.

4.3.3 SUMO Simulation Setup

As already outlined in D3.3, the simulation setup relies on three main data sources:

- The API as provided by ETRA and HYP individually, provides the topology as well as specific snapshots of counting data.
- Historical traffic data for the duration of several months
- Locations of induction loops

The essential ingredients for a SUMO simulation are

- A SUMO road network file (.net.xml), which defines the geometric layout of roads, lanes, junctions, and traffic lights.
- Additional Files (.add.xml), which are used for extra features like bus stops, induction loops, variable speed signs, or specific traffic light logic.
- Route/Demand File (.rou.xml), which specifies the vehicles, vehicle types (car, bus, truck), their characteristics (speed, acceleration), and the precise paths (edges) they take from origin to destination.

Step 1: For the road network generation, we query the topology from the corresponding API. This topology definition is then converted to the OSM format with the help of Python code. Doing so, we can make use of the OSM to SUMO converter script that comes with each SUMO installation to eventually obtain a SUMO road network.

Step 2: To create traffic demand (see Step 3), SUMO needs to be able to map given counting data to the network. This is done by running the SUMO tool mapDetectors.py, which uses the

DRAFT

network file and a file of counting locations to create a new file that defines induction loops with regards to the edges of the network file.

Step 3: Finally, to create traffic demand, the following procedure is executed:

- Prepare Count Data: Traffic counts are converted into the format required by SUMO, specifying the measured vehicle volumes at individual detector locations and time intervals.
- Create Vehicle Flows: Measured traffic counts were converted into SUMO vehicle flows by using traffic counts and detector locations.
- Generate Candidate Routes: Based on the network topology, a set of possible routes is created.
- Estimate Routes with routeSampler.py: The SUMO tool routeSampler.py was used to generate vehicle demand that matches the observed detector counts. The script iteratively samples from the candidate route set and adjusts the number of vehicles assigned to each route until the simulated detector flows closely reproduce the measured traffic counts. This allows the creation of realistic route files without requiring explicit origin/destination matrices.
- Create SUMO Demand Files: The routes selected by routeSampler.py were converted into SUMO route files (*.rou.xml), defining vehicle departures, vehicle types, and route assignments for the simulation period.

The current implementation relies on SUMO 1.27.0. As SUMO is distributed free of charge under the Eclipse Public License v 2.0, migration to a more recent release might be considered. All use-case scenarios will be simulated in an open-loop and closed-loop configuration. The open-loop simulations establish the baseline, i.e., without an orchestrator in the loop, while the closed-loop simulations incorporate the orchestrator into the feedback loop.

DRAFT

5. Conclusion

This deliverable has presented the initial progress on two key aspects of the CARMONY Interlinked Orchestrator related to the work of Tasks 5.1 and 5.4, namely the description and overview of the data sharing platform that will support the data exchange pipeline throughout CARMONY, and the initial scenarios and modelling approach that will be used for testing the Orchestrator. Each of these aspects plays a vital role in the implementation and subsequent operation of the Interlinked Orchestrator, and this deliverable establishes the groundwork and provides the roadmap for the future development of both.

More specifically, regarding the data sharing platform related to Task 5.1, related CARMONY work was reviewed to identify the data requirements that the platform would need to support, establishing 5 key requirements: To make static datasets readily and securely available to all CARMONY partners; to provide support for webAPIs for real-time data transfers; to be able to implement data access restrictions for sensitive data; to be compliant with standards and best practices; and to require minimal technical knowledge from partners in order to participate. Two widely used data management system frameworks were evaluated against these requirements, with each found to satisfy most but not all requirements. Based on this evaluation the proposed solution was presented, consisting of a CKAN system to handle the majority of CARMONY's operational data needs, and a small-scale dataspace to be deployed in parallel incorporating only a subset of CARMONY's datasets to prepare them for long-term availability.

Regarding the scenario modelling and simulation, two sets of simulation scenarios were presented: one list consisting of scenarios aimed at evaluating the Orchestrator under road conditions and during incidents at a fine scale, and another containing scenarios focussing on governance and regulatory policies and their impact at a larger scale. Following the definition of the scenarios the initial setup and preparation for the simulation implementations was presented, for each pilot site and for the simulation engine that will handle the execution of all simulations.

This deliverable will therefore provide a roadmap in the development of the CARMONY Interlinked Orchestrator and is expected to be used as a reference in this endeavour, on the one hand by providing an architectural overview to be used during the development of the data sharing platform, and on the other by establishing a set of scenarios that the Orchestrator will need to execute and evaluate.

DRAFT

6. Abbreviations

Abbreviations of terms used within the deliverable, in alphabetical order.

Term	Definition
API	Application Programming Interface
CARMONY	Coordinated Action for Responsive Mixed Orchestration and Network Yield
CCAM	Cooperative Connected and Automated Mobility
CV	Connected Vehicle
CAV	Connected Automated Vehicle
DMP	Data Management Plan
DSP	DataSpace Protocol
EDC	Eclipse Dataspace Components
EMDS	European Mobility DataSpace
IDSA	International Dataspaces Association
KPI	Key Performance Indicator
LEZ	Low Emission Zone
OSM	OpenStreetMap
PU	Public
SUMO	Simulation of Urban MObility (traffic simulation engine)
UC	Use Case
VPN	Virtual Private Network
WP	Work Package

DRAFT

7. References

[1] “CARMONY Grant Agreement Number AMD-101202858-6”, 2026-05-22

[2] "CARMONY D1.2: Initial Ethics Manual & Data Management Plan, V2.0 Final", 2025-10-25

[3] “CARMONY D3.1: Specifications and conceptual architecture - Initial Version, V1.0 Final”, 2026-04-28

[4] “CARMONY D3.3: Agent-based simulation controllers and simulation platform – Initial Version”, 2026

[5] Curry, E. (2020). Dataspaces: Fundamentals, Principles, and Techniques. In: Real-time Linked Dataspaces. Springer, Cham. https://doi.org/10.1007/978-3-030-29665-0_3

[6] Pretzsch, S., Drees, H., Rittershaus, L. (2022). Mobility Data Space. In: Otto, B., ten Hompel, M., Wrobel, S. (eds) Designing Data Spaces . Springer, Cham. https://doi.org/10.1007/978-3-030-93975-5_21

DRAFT