A Novel Logistics Optimisation Module for Improving Response Deployment of Critical Resources

Evangelos Filippas
ASTRIAL GmbH
e.filippas@astrial.de

Akhil Kumar
ASTRIAL GmbH
a.kumar@astrial.de

Muhammad Ali
ASTRIAL GmbH
m.ali@astrial.de

Sabina Ziemian
ASTRIAL GmbH
s.ziemian@astrial.de

Antonios Kostaridis
ASTRIAL GmbH
a.kostaridis@astrial.de

Evangelos Sdongos
ASTRIAL GmbH
e.sdongos@astrial.de

ABSTRACT

Effective crisis response requires both situational awareness and efficient coordination of logistics resources across multiple agencies. While Command, Control, Communications and Intelligence / Incident Management Systems (C3I/IMS) provide a Common Operational Picture (COP) and vehicle dispatching, they lack embedded optimisation capabilities for dynamic routing and allocation. This work operationalises logistics optimisation within a real C3I/IMS by introducing a module that formulates response logistics as a Multi-Depot Capacitated Vehicle Routing Problem with heterogeneous fleet characteristics, inventory coupling, compatibility constraints, and explicit time components. The module functions as an event-driven intelligence layer, automatically generating allocation plans, geospatial routes, Estimated Time of Arrival (ETA) predictions, and diagnostic outputs. A synthetic civil protection scenario demonstrates system integration and workflow coherence. The contribution lies in embedding optimisation-driven logistics intelligence within a multi-agency command-and-control environment, improving response efficiency, transparency, and interoperability, while also supporting preparedness through integrated simulation capabilities.

Keywords

Crisis management; Disaster Response; Humanitarian logistics; Logistics optimisation; C3I systems; Decision support systems, Training platform

INTRODUCTION

Crisis and disaster management environments are characterised by high uncertainty, time pressure, degraded infrastructure, and multi-agency/multi-team coordination requirements. While substantial advances have been made in enhancing situational awareness through sensor networks, geospatial platforms, and information fusion systems, logistics coordination and deployment remains a persistent operational bottleneck. The effective distribution of civil protection and disaster relief resources - including heavy equipment, emergency vehicles, medical supplies, temporary shelters, and trained personnel - continues to be impeded by structural limitations that become particularly visible during large-scale or complex incidents. Humanitarian and relief logistics has long been recognised as a critical determinant of response effectiveness, yet operational challenges persist when decisions must be taken under severe time and information constraints (Van Wassenhove, 2006; Kovács & Spens 2007; Rottkemper & Fischer, 2013; Besiou et al., 2021).

Recent events illustrate recurring failure modes such as delayed dispatch, fragmented visibility of inventory and transport capacity, unclear responsibility chains, and transportation disruptions caused by infrastructure damage. Such patterns have been repeatedly documented in humanitarian supply chains, where capacity limitations, coordination breakdowns, and insufficient preparedness constrain response performance (Jahre & Jensen, 2010; Besiou & Van Wassenhove, 2020). From an operational research perspective, such conditions can be formalised as dynamic and stochastic vehicle routing problems characterised by evolving and uncertain information (Pillac et al., 2013), combined with risk-aware facility location and inventory prepositioning decisions in disaster environments (Campbell & Jones, 2011). In addition, crisis response requires urgency- and equity-aware decision making that is not easily captured by conventional cost-driven objective functions (Holguín-Veras et al., 2013).

Over the past two decades, information systems for emergency management have improved coordination and information sharing. Early work by Turoff (2002) highlighted the need for computer-supported emergency response systems capable of aggregating and disseminating actionable information. Subsequent work within the information systems for crisis response and management (ISCRAM) community advanced the concept of the Command-and-Control and COP, enabling multi-agency collaboration, real-time situational awareness, and improved communication across hierarchical command structures (Van de Walle & Turoff, 2006; Van de Walle et al., 2010). Modern C3I/IMS platforms integrate sensor data, mapping services, incident reports, and resource tracking into unified dashboards, enhancing information transparency and coordination efficiency. Decision-support research further stresses the importance of embedding analytics within operational workflows to improve coordination under uncertainty (Comes et al, 2011).

However, only a few systems attempt to integrate logistics and transport planning capabilities (Widera et al., 2017), and most crisis information systems primarily act as aggregation and coordination layers (Comes et al., 2011; Ransikarbum & Mason, 2016): they provide visibility into incidents and resource status, but fail to effectively embed optimisation capabilities that can dynamically compute dispatch alternatives, balance competing demands, or simulate allocation outcomes under transport disruption. As a result, resource planning and transportation decisions are often manually performed or are supported by external tools disconnected from the COP, creating a persistent disconnect between information visibility and decision computation.

In parallel, humanitarian logistics research has developed a substantial body of optimisation and planning approaches, including pre-positioning models for emergency supplies (Rawls & Turnquist, 2010), dynamic vehicle routing under evolving conditions (Pillac et al., 2013), facility location and capacity planning (Campbell & Jones, 2011). These models demonstrate analytical potential, although remain theoretical frameworks or standalone tools not integrated into the operational command-and-control environments used by civil protection and disaster relief agencies. In contrast, commercial supply chain management platforms embed route optimisation, inventory balancing, and transportation planning within integrated enterprise environments (Simchi-Levi et al., 2008; Chopra & Meindl, 2016). Nevertheless, direct transfer to crisis operations is non-trivial due to disrupted infrastructure, intermittent communications, multi-level decision hierarchies, and non-economic objectives such as urgency and life-saving prioritisation (Holguín-Veras et al., 2013), as well as inter-agency interoperability requirements with legacy systems.

These gaps collectively motivate integrated solutions that embed optimisation-driven logistics intelligence directly within a COP to support both headquarters-level coordination and field-level situational awareness in real time. Present work prototype leverages OpenRouteService and VROOM for network-based travel-time estimation and capacitated routing. However, such tools do not inherently address inventory coupling, crisis workflows, or COP integration.

The present work addresses this gap by presenting the design and development of a Logistics Optimisation Module embedded within a C3I/IMS, developed in the context of SYNERGISE (https://www.synergise-project.eu/) and FORESIGHT (https://foresight-project.eu/) Horizon European projects. The module connects distributed warehouse inventories, heterogeneous fleet characteristics, structured logistics requests, and geospatial routing data to generate optimised dispatch recommendations, route calculations, and transparent ETA predictions. A synthetic civil protection and disaster relief scenario is used to demonstrate architectural coherence, allocation feasibility, and routing transparency.

The contribution lies in operationalising optimisation within a multi-agency crisis information system by embedding decision-support analytics into an event-driven COP workflow. The novelty does not lie in a new optimisation methodology, but in the system-level integration of established optimisation techniques within operational command-and-control environments. Unlike standalone or manually triggered approaches, the proposed event-driven architecture automatically recomputes allocation and routing decisions in response to changes such as updated requests, resource availability, or transport conditions. This reduces operator workload, minimises delays, and ensures that the COP reflects up-to-date optimisation results in real time. By coupling optimisation outputs with situational awareness layers, the module supports both headquarters-level planning and field-level coordination, bridging the gap between information visibility and actionable decision support. This enables more responsive, transparent, and adaptive logistics coordination in rapidly evolving crisis scenarios.

The remainder of this paper is structured as follows. First, the system architecture and module positioning is presented. Next a description of the optimisation model and integration logic is provided. Finally, a synthetic validation scenario is introduced and present work concludes with contributions and future directions.

SYSTEM ARCHITECTURE

The Logistics Optimisation Module is an analytical engine embedded within a C3I/IMS, functioning as a decision-support component integrated with the COP rather than a standalone tool.

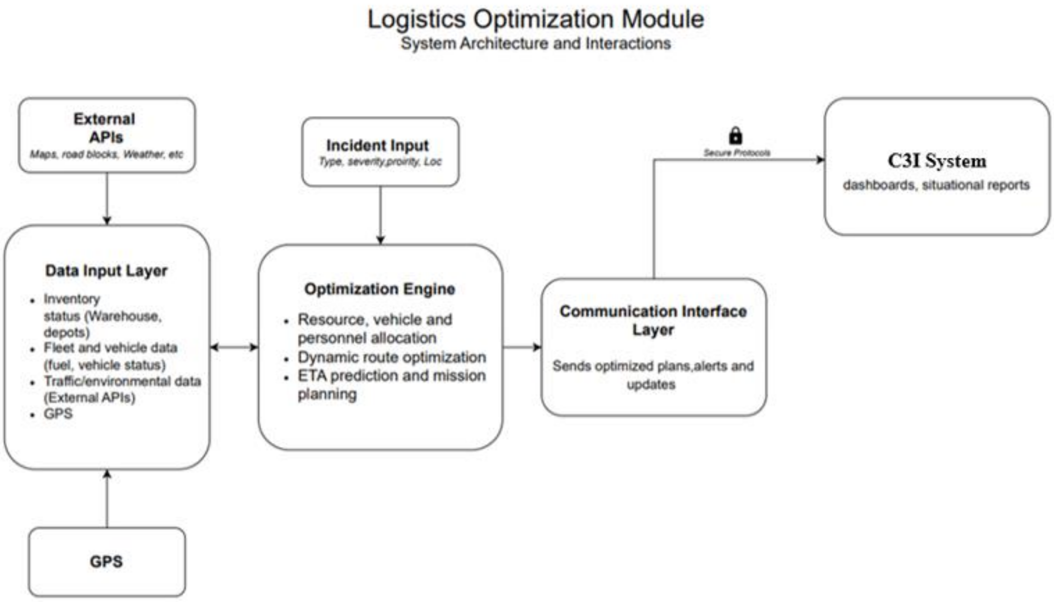

Figure 1. Logistics Optimisation Module architecture and interactions within the C3I/IMS ecosystem

Role within the C3I/IMS environment

The COP provides a shared, map-based view of incidents, resources, and operations. The module enhances this by transforming resource visibility into dispatch recommendations, routing plans, and ETA predictions. It links inventory, fleet capabilities, and demand within a unified operational interface.

Logistics decisions in crisis response span field assessment, coordination centres, and execution. The module aligns with these workflows by supporting request handling, feasibility assessment, dispatch planning, and monitoring.

Positioning of the Logistics Optimisation Module in the system

The module operates between (a) resource and inventory data sources, (b) the C3I/IMS workflow and COP interface, and (c) external contextual data services (e.g., road network state, traffic, incident constraints). The C3I/IMS provides operational context, while the module performs allocation, routing, and ETA computation. Results (vehicle assignments, routes, ETAs) are reintegrated into the COP. The architecture supports two operational modes:

1. Response mode for real-time dispatch and routing under current constraints
2. Simulation mode for evaluation of alternative strategies for planning and preparedness.

Architectural layers and components

The system follows a layered architecture:

- Data layer: manages warehouses, inventory, fleet assets, and operational events
- Integration layer: enables event-driven communication and data exchange

- Optimisation layer: performs allocation, routing, and feasibility checks
- Application layer: provides COP visualisation and user interaction

The input–process–output structure is, depicted also in Figure 2 is as follows:

- Inputs: inventory, vehicles, requests, and contextual constraints
- Process: feasibility validation, allocation, routing, ETA estimation
- Outputs: warehouse activation, vehicle assignments, routes, ETAs, diagnostics

This structure enables continuous recalculation under changing conditions. Finally, the module integrates via standardised data formats and event-driven communication, supporting real-time interaction with the COP and extensibility to multiple transport modalities (e.g., UAVs, rail, maritime).

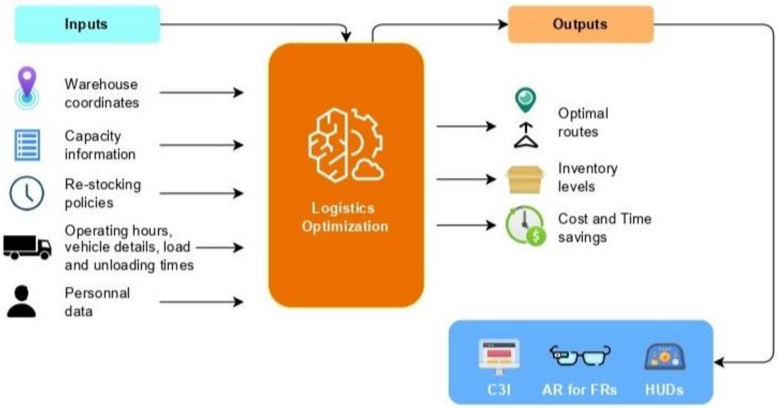

Figure 2. Input-Output interaction of the Logistics Optimisation Module within the C3I/IMS workflow

LOGISTICS OPTIMISATION MODEL

Under operational constraints, the Logistics Optimisation Module supports dynamic allocation and routing of heterogeneous resources (vehicles and relief items) from multiple geographically distributed depots (warehouses) to one or more incident destinations, considering also heterogeneous fleet characteristics, capacity constraints, and routing decisions under operational limitations. The problem is formulated as a Multi-Depot Capacitated Vehicle Routing Problem (MD-CVRP), a well-established extension of the classical Vehicle Routing Problem (VRP) in operations research literature (Laporte, 2009; Toth & Vigo, 2014; Montoya-Torres et al., 2015; Braekers et al., 2016). In the present implementation, the formulation is further extended to include item-vehicle compatibility constraints and time-related components (travel time plus handling times), reflecting the operational realities of crisis logistics environments. The optimisation aims to: (i) satisfy requested material quantities, (ii) assign feasible vehicles given capacity and compatibility constraints, (iii) generate operationally feasible routes across depots and destinations, and (iv) provide transparent ETA and diagnostic outputs suitable for decision support within the C3I/IMS environment.

Data model (input/output)

The Logistics Optimisation Module relies on a structured data model formalising supply (warehouses and fleet), demand (logistics requests), and decision-support outputs. The optimisation process consumes three primary inputs: (i) inventory data, (ii) vehicle data, and (iii) structured logistics request events.

The inventory dataset represents distributed warehouse infrastructure. Each depot is geospatially referenced and associated with available resource types and quantities. Optional attributes include weight class, handling requirements, or compatibility flags. This dataset defines supply-side constraints and feasible dispatch origins. The vehicle dataset describes fleet capabilities through vehicle identifiers, type/category, starting location, availability status, and capacity attributes (e.g., weight or volume). Compatibility constraints determine which resource types a vehicle may transport, enabling heterogeneous fleet modelling under capacity restrictions. Together, these inputs define a parameterised MD-CVRP instance under inventory, capacity, and compatibility constraints.

The module produces both human-interpretable and machine-readable outputs, including:

- Warehouse-level allocation decisions
- Vehicle assignment and load configuration
- Ordered route plans with geospatial visualization
- Segment-level and cumulative ETA estimates
- Operational diagnostics identifying shortages or infeasibilities

ETA calculations incorporate travel time and operational handling durations where applicable. Diagnostic outputs report inventory feasibility, capacity utilisation, and unmet demand conditions. All results are exported in structured format to support integration with COP visualisation, storage, and reporting services. The data model separates supply, transport, and demand structures while preserving semantic links between them. This structure enables modular optimisation, event-driven recalculation, and incremental system integration.

Optimisation approach

The Logistics Optimisation Module implements a variant of the MD-CVRP extended to incorporate heterogeneous fleets, item–vehicle compatibility constraints, and temporal components. The underlying Vehicle Routing Problem (VRP) and its extensions are NP-hard, and the inclusion of these additional constraints further increases complexity, making exact optimisation impractical for real-time applications. Therefore, heuristic methods such as VROOM are employed to provide efficient and timely solutions.

Routing and travel-time estimation are supported by OpenRouteService based on OpenStreetMap data, while capacitated routing optimisation is performed using VROOM. The engine combines a constructive heuristic with local search based on neighbourhood exploration, making it suitable for time-sensitive crisis scenarios. This approach, however, implies a trade-off, as global optimality is not guaranteed.

The optimisation logic is described at a conceptual level, capturing allocation of resources from multiple depots to vehicles, combined with routing and time estimation under capacity, compatibility, and inventory constraints. The objective is to minimise total routing cost, currently represented by estimated travel time derived from network speed limits. The formulation remains extensible to incorporate disruption-aware routing and composite operational metrics such as traffic conditions, hazards, or prioritisation policies.

In future versions, the cost function may incorporate composite operational metrics including dynamic traffic information, road blockages, hazard zones, infrastructure damage, prioritisation policies, or risk exposure levels. It could also account for specific operational privileges of emergency vehicles, such as the ability to exceed standard speed limits, bypass traffic signals, or use contra-flow movements under authorised emergency conditions. Such extensions would require integration with real-time traffic feeds, disruption-aware routing models, and context-sensitive speed profiling, as explored in dynamic vehicle routing literature (Pillac et al., 2013) and time-dependent shortest-path modelling (Ichoua et al., 2003).

Decision variables include transported quantities per warehouse–vehicle pair, binary vehicle–warehouse assignments, and vehicle route sequences. Constraints ensure demand satisfaction (where feasible), inventory limits, vehicle capacity compliance, and compatibility requirements. Compatibility is currently modelled through simplified binary indicators, however, emergency logistics often require more granular modelling of vehicle capabilities, including volumetric limitations, equipment sensitivity (e.g., vibration, temperature, or hazardous materials handling), securing mechanisms, and specialised transport requirements. While these aspects are currently abstracted into simplified compatibility flags to ensure computational tractability and prototype validation, future extensions could refine this representation by incorporating multi-dimensional loading constraints, heterogeneous compartment modelling, or resource-specific transport attributes (Toth & Vigo, 2014); Braekers et al., 2016).

In addition to allocation and capacity feasibility, the formulation enforces structural route consistency conditions derived from classical CVRP flow conservation principles (Toth & Vigo, 2014). For each vehicle, the computed route must originate from its designated starting depot location and terminate at the designated delivery destination associated with the logistics request. Intermediate visits are restricted to selected warehouse nodes determined by the allocation decision, and each such warehouse is visited at most once per vehicle route to avoid redundant traversal.

Beyond physical feasibility and route consistency, the model emphasises temporal transparency through a component-resolved Estimated Time of Arrival (ETA) formulation. The ETA is not limited to routing travel time but incorporates segment distances and corresponding travel durations together with operational parameters such as loading and unloading activities. Each resource type includes metadata describing handling requirements, and these operational handling factors are integrated into the cumulative ETA calculation to provide a more realistic representation of end-to-end dispatch time per vehicle, enabling logistics officers to understand not only the “best plan” but also the key contributors to delay. Future work may extend the time estimation process to incorporate dynamic personnel availability, multi-stage handling operations, and equipment-specific loading constraints. Such refinements would enable a more granular and context-sensitive modelling of operational service times beyond the current prototype implementation.

Operational workflow within the C3I/IMS ecosystem

The model is executed in an event-driven workflow aligned with crisis and disaster relief decision-making processes. At implementation level, the operational logic follows a structured processing pipeline, illustrated in Figure 3. A logistics request is published via the platform’s messaging infrastructure implemented using Kafka (Auradkar et al., 2012), consumed by the Logistics Optimisation core, passed to the routing engine, processed by the ETA module, and subsequently visualised within the COP while being exported in structured JSON format (Bray, T., 2017) for interoperability with other C3I/IMS components.

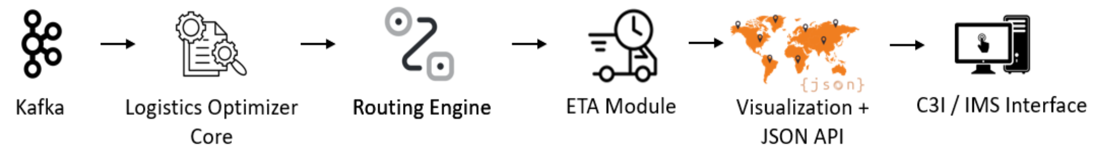

Figure 3. Operational processing pipeline of the Logistics Optimisation Module within the C3I/IMS ecosystem.

Rules engine and feasibility layer

In the current prototype, a rule-based feasibility layer is implemented to ensure that valid allocation decisions are passed to the routing engine. This layer is embedded within the pre-processing stage of the optimisation pipeline, where operational constraints are evaluated prior to route computation. That layer acts as a policy and constraint enforcement mechanism, verifying that allocation decisions satisfy capacity limitations, vehicle-resource compatibility requirements, and prioritisation policies before invoking the optimisation engine. In practice, this is achieved through sequential filtering and validation steps, including inventory availability checks, vehicle capability matching, and capacity feasibility verification.

Representative operational rules include:

- Capacity constraint: A vehicle can only be assigned resources if the total load remains within its capacity limits (e.g., a rescue van cannot transport multiple generators if their combined weight exceeds allowable capacity).
- Compatibility constraint: Certain resources require specific vehicle types or handling capabilities (e.g., crane components are assigned exclusively to crane trucks and not to standard vans).
- Prioritisation logic: When total demand cannot be fully satisfied, allocation prioritises critical resources (e.g., medical supplies are delivered before non-critical equipment under constrained capacity).

By embedding these rules within the optimisation pipeline, the system ensures that only feasible and operationally consistent solutions are generated. Future work will extend this capability through a dedicated rule-definition interface, enabling end-users to configure and adapt operational policies via a Graphical User Interface (GUI) within the C3I/IMS environment.

Upon receiving a logistics request event, the module automatically:

1. Parses demand and destination
2. Evaluates inventory availability
3. Filters feasible vehicles
4. Computes allocation and routing
5. Generates routes and ETAs
6. Publishes results to the COP

This pipeline-based architecture ensures modularity, scalability, and separation of concerns between data ingestion, optimisation logic, routing computation, ETA estimation, and visualisation. By relying on standardised JSON templates and event streaming, the system maintains loose coupling with other C3I/IMS services while enabling real-time recalculation when operational conditions change.

From a functional perspective, the module provides a set of core operational capabilities including:

1. Real-time request ingestion without manual triggering.
2. Automatic inventory matching across distributed warehouses.
3. Vehicle capability filtering under heterogeneous fleet conditions.
4. Route generation minimising travel time while respecting capacity and route-consistency constraints.
5. Detailed ETA transparency including segment travel and handling durations.
6. Operational diagnostics identifying shortages, infeasible assignments, and capacity violations.
7. Machine-readable export of allocation summaries, vehicle assignments, ETA breakdowns, and diagnostic indicators.

In practical civil protection and disaster relief operations, the system supports a continuous decision cycle. Requests generated from field assessments are processed automatically, producing dispatch plans and ETA predictions. As conditions evolve (e.g., demand, inventory, or accessibility), the system dynamically recalculates solutions, enabling adaptive decision-making.

PRELIMINARY VALIDATION - CONCEPT OF OPERATIONS

To demonstrate the operational behaviour of the proposed Logistics Optimisation Module, a synthetic civil protection scenario was constructed reflecting a typical multi-resource dispatch request under heterogeneous fleet and inventory constraints. The scenario represents a field-level material demand communicated to a coordination centre through the event-driven C3I/IMS infrastructure.

The logistics request required the delivery of 2 Personal Protective Equipment (PPE) Boots, 2 Generators, 2 Crane Components, and 1 First Aid Kit to a geospatially defined destination located in Berlin, Germany, at latitude 52.428388 ° N and longitude 13,523065 °E (WGS84). In addition, WGS84-based coordinates can be further mapped to human-readable semantic locations (e.g., landmarks or place names) to enhance usability for end-users, which is considered as part of future work. The request was submitted through the messaging layer as a structured "LOGISTICS_REQUEST_CREATED" event in JSON format, automatically triggering the optimisation workflow without manual intervention.

The scenario was intentionally defined at a limited scale to isolate and evaluate core functionalities of the system, including distributed inventory matching, heterogeneous vehicle filtering, route generation using geospatial network data, explicit ETA computation including handling times, and operational diagnostics reporting. While simplified, the instance captures essential structural characteristics of crisis logistics allocation problems and enables transparent validation of the optimisation and integration pipeline.

The synthetic scenario is based on a distributed warehouse network representing multiple logistics depots within the Berlin metropolitan area. Each warehouse is geospatially referenced and integrated into the optimisation framework as a potential supply origin. Table 1 summarises the geographical configuration of the warehouse infrastructure that defines the spatial supply network for the optimisation process.

Table 1. Geospatial Configuration of the Warehouse Network

Warehouse ID	Warehouse Name	Latitude	Longitude
W1	Depot_Berlin_Mitte	52.498690	13.494721
W2	Depot_Berlin_Nord	52.551360	13.385024
W3	Depot_Berlin_Süd	52.503013	13.276789
W4	Depot_Berlin_Ost	52.528145	13.419503
W5	Depot_Berlin_West	52.457334	13.423820
W6	Depot_Berlin_Pankow	52.477604	13.256366
W7	Depot_Berlin_Treptow	52.530273	13.389612
W8	Depot_Berlin_Spandau	52.506153	13.290281
W9	Depot_Berlin_Tempelhof	52.501982	13.214466
W10	Depot_Berlin_Kreuzberg	52.509289	13.443405

Following event ingestion, the MD-CVRP-based optimisation process evaluated distributed warehouse inventories and determined feasible supply allocations. Table 2 summarises the warehouse-level allocation outcome derived from the prototype interface.

Table 2. Warehouse Allocation Summary

warehouse_id	warehouse_name	items_supplied
W3	Depot_Berlin_Süd	crane_component:2, ppe_boots:2
W5	Depot_Berlin_West	first_aid_kit:1, generator:2

The optimiser selected Depot_Berlin_Süd (W3) as the primary supply depot for crane components and PPE boots, while generators and the first aid kit were allocated based on stock availability and compatibility constraints from an alternative depot, i.e. Depot_Berlin_West (W5). This confirms correct enforcement of inventory feasibility constraints and proper coupling between demand satisfaction and depot-level stock availability.

Table 3. Vehicle Assignments and Load Configuration

vehicle_id	vehicle_type	items_assigned
VEH4	crane truck	crane_component:1
VEH11	crane truck	crane_component:1
VEH22	rescue van	ppe_boots:2, first_aid_kit:1
VEH_H4	truck (heavy)	generator:2

The optimisation process automatically selected four vehicles to satisfy capacity and compatibility requirements, as detailed in Table 2. The resulting allocation demonstrates correct enforcement of heterogeneous fleet constraints, with equipment assigned in accordance with vehicle type, load-bearing capacity, and compatibility rules, thereby confirming that both capacity limitations and item–vehicle constraints are consistently respected within the optimisation workflow. For each vehicle, the system generated ordered route segments derived from OpenStreetMap network data via the routing engine. Travel times were computed based on legally permitted speed limits, and synthetic operational handling times were integrated into the cumulative ETA calculation. Table 3 summarises the total time per vehicle.

Table 4. Total Time per Vehicle

Vehicle	Total Time (min)
VEH22	53.92
VEH11	70.52
VEH4	70.95
VEH_H4	143.04

The heavy truck (VEH_H4) exhibits the longest delivery time (143.04 minutes), primarily due to extended loading (60 minutes) and unloading (50 minutes) durations rather than excessive travel distance. Therefore, ETA estimate incorporates operational handling times in addition to pure travel time, providing temporal transparency suitable for decision support. Breakdown of travel and handling components is presented in Table 5.

The optimisation results were visualised using a lightweight development interface designed to validate functional correctness prior to full C3I/IMS embedding. The interface displays matched warehouse resources, vehicle assignments, and load distribution. This view confirms correct integration between inventory data and routing allocation outputs. The timeline view presented in Figure 4. provides an intuitive representation of staggered vehicle departures and segment transitions. It visually differentiates route segments and highlights cumulative arrival times, supporting temporal reasoning and coordination between operational roles. Moreover, the geospatial interface in Figure 5. displays multi-vehicle routes with distinct colour coding, connecting depots to the crisis destination. This integration demonstrates how optimisation outputs are translated into actionable map-based recommendations suitable for inclusion within a Common Operational Picture.

Finally, diagnostics module provides structured overview of resource feasibility and fleet utilization. Table 6 compares requested and available quantities, indicating number of vehicles required and assigned, identifying unfulfilled demand or additional vehicle requirements. In present scenario, all requested items were successfully matched with available stock and appropriately assigned vehicles, with no unallocated resources or unmet demand. This confirms full demand satisfaction under the prevailing inventory and capacity constraints and demonstrates the system's ability to transparently report allocation feasibility within the operational workflow.

Table 5. Detailed ETA Breakdown by Vehicle and Route Segments

Vehicle	Segment	Distance (km)	Travel Time (min)	Loading Time (min)	Unloading Time (min)	Cumulative ETA (min)
VEH11 (crane truck)	Start → Stop 1	8.78	10.49	25	0	37.49
VEH11 (crane truck)	Stop 1 → Stop 2	20.84	21.03	0	12	70.52
VEH22 (rescue van)	Start → Stop 1	4.89	8.90	2	0	12.90
VEH22 (rescue van)	Stop 1 → Stop 2	12.63	14.93	5	0	32.84
VEH22 (rescue van)	Stop 2 → Stop 3	9.69	14.08	0	7	53.92
VEH4 (crane truck)	Start → Stop 1	9.77	10.92	25	0	37.92
VEH4 (crane truck)	Stop 1 → Stop 2	20.84	21.03	0	12	70.95
VEH_H4 (truck heavy)	Start → Stop 1	9.34	16.96	60	0	78.96
VEH_H4 (truck heavy)	Stop 1 → Stop 2	9.69	14.08	0	50	143.04

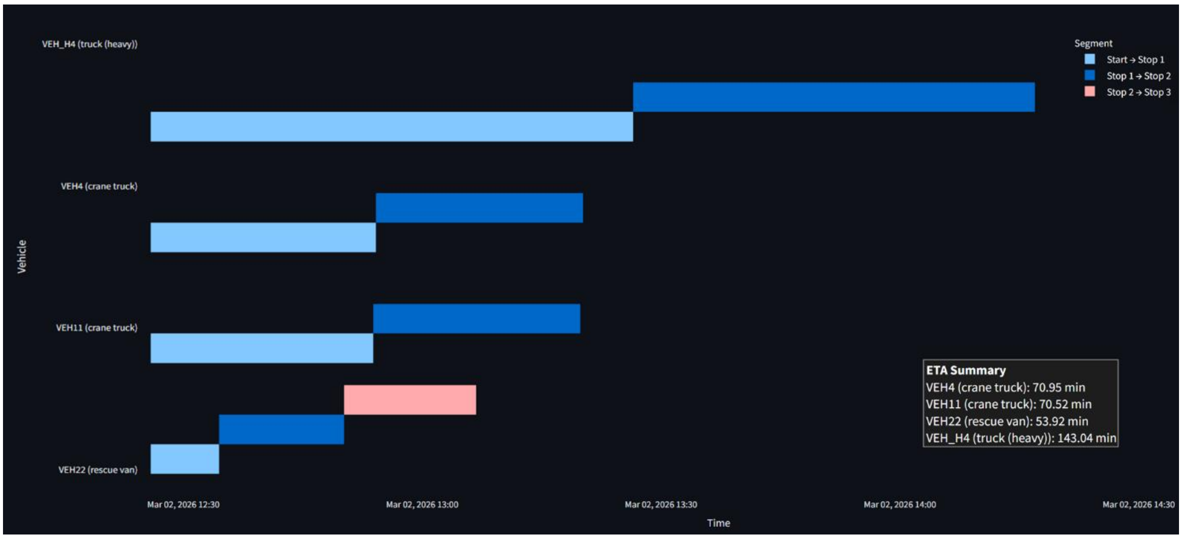

Figure 4. Timeline Visualisation of Vehicle Routes and Cumulative ETA Estimation.

Handling corner cases such as partial fulfilment and infeasible demand is supported through the diagnostics component of the module. When it is not possible to satisfy all requested quantities due to limited inventory, insufficient vehicle capacity, or compatibility constraints, the system generates the best feasible partial allocation while maintaining operational consistency. In such cases, the diagnostics output provides a structured comparison between requested and allocated quantities, including information on available resources, unassigned items, and vehicle utilisation. This enables decision-makers to immediately identify shortages, capacity limitations, and infeasible assignments within the operational workflow. When multiple partial allocation solutions are possible, selection is guided by feasibility constraints, prioritisation of critical resources, and minimisation of routing cost, ensuring that the most operationally relevant solution is selected under constrained conditions. To illustrate this behaviour, consider a scenario where requested quantities exceed available resources and vehicle capacity. In this case, the system produces a partial fulfilment solution together with diagnostic indicators highlighting unmet demand and allocation gaps.

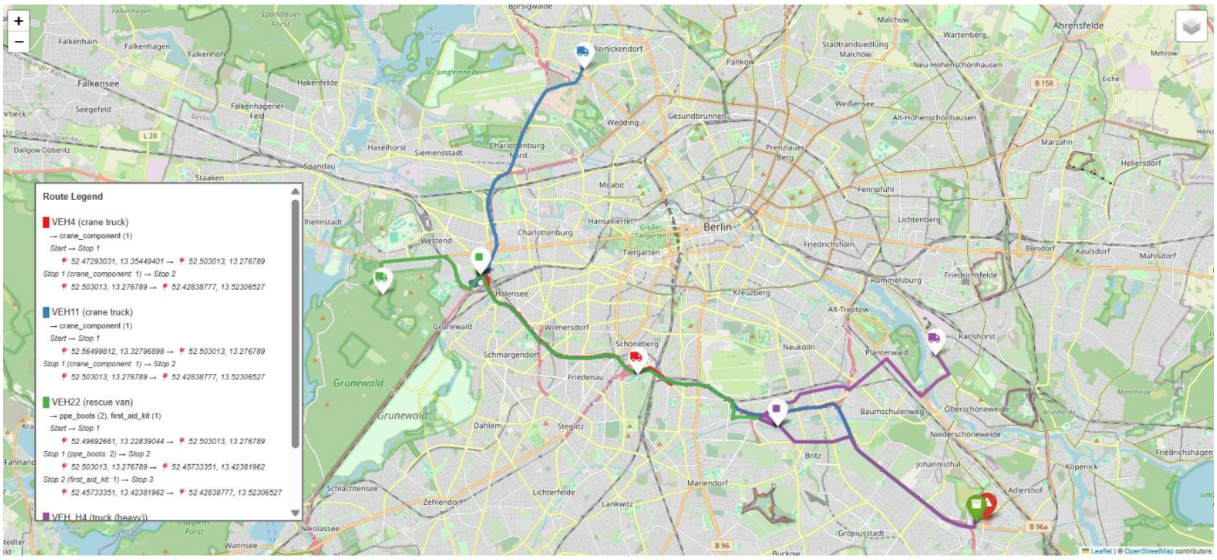

Figure 5. Map-based Route Visualisation with Warehouse and Destination Nodes.

Table 6. Operational Diagnostics Report: Resource Feasibility and Vehicle Utilisation

Resource Type	Requested	Available	Not Available	Vehicle Type	Vehicles Free	Vehicles Needed	Assigned Vehicles	Unassigned Items	Additional Vehicles
Crane_Component	2	2	0	crane truck	6	2	2	0	0
First_Aid_Kit	1	1	0	truck	2	1	1	0	0
Generator	2	2	0	truck (heavy)	4	1	1	0	0
PPE_Boots	2	2	0	truck	2	1	1	0	0

Table 7: Example Logistics Request (Input Event)

Parameter	**Value**
Event Type	LOGISTICS_REQUEST_CREATED
Destination Latitude	52.428388
Destination Longitude	13.523065

Resource Type	Quantity
Generator	5
Crane Component	3

This structured feedback supports transparent decision-making by explicitly exposing infeasibilities and enabling operators to adjust requests, reallocate resources, or trigger additional logistics actions. Future work will further refine prioritisation strategies and explore alternative allocation policies under constrained conditions.

Table 8: Diagnostics Output under Partial Fulfilment Conditions

Resource Type	Requested	Available	Unassigned Items	Vehicles Needed	Vehicles Assigned
Generator	5	2	3	2	1
Crane Component	3	3	2	3	1

Although the present validation scenario is synthetic and limited in scale, the results demonstrate correct enforcement of inventory feasibility constraints, heterogeneous vehicle compatibility filtering, and capacity-respecting allocation decisions across the fleet. The system produces geospatially grounded route generation and explicit ETA transparency that incorporates both travel and handling times, while simultaneously ensuring machine-readable JSON interoperability for integration within the broader C3I/IMS ecosystem. The prototype therefore confirms that the module successfully operationalises an MD-CVRP-based formulation within an event-driven crisis management workflow.

From a computational performance perspective, the routing component of the module relies on the VROOM optimisation engine, which is specifically designed for efficient vehicle routing and is capable of solving medium-sized problem instances within seconds under standard configurations. Public benchmarks provided by the VROOM project indicate execution times typically ranging from tens to several hundreds of milliseconds depending on problem size and search intensity. This demonstrates a favourable trade-off between computational efficiency and solution quality, which is essential for time-sensitive crisis response scenarios. It is important to note that above performance discussion refer primarily to the routing component. The overall system performance additionally depends on upstream and downstream processes, including inventory matching, rule-based feasibility filtering, ETA computation, and system integration within the C3I/IMS environment.

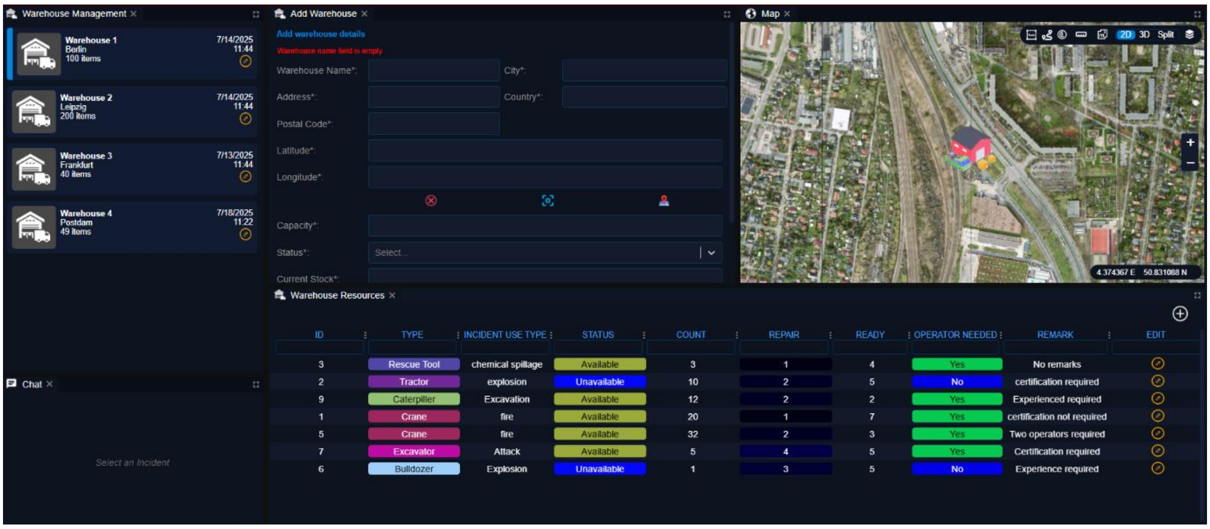

Figure 6. C3I/IMS Warehouse Management and COP Integration (Work in Progress).

The module is being currently integrated within the broader C3I/IMS environment. Figure 6 illustrates the emerging warehouse management interface within the COP. In parallel with this progressive system integration, further development validation and performance evaluation will extend beyond algorithmic refinement to include close collaboration with end-user organisations. Through planned field demonstrations, user-training sessions, and structured evaluation activities, the module will be assessed under realistic operational conditions to validate routing logic, time estimations, allocation strategies, and interoperability assumptions using real-world data. These activities will support iterative refinement of service-time modelling, compatibility constraints, and decision-support outputs, while simultaneously consolidating the module's embedding within the C3I/IMS architecture in alignment with real civil protection workflows, preserving scalability and interoperability within the broader ecosystem.

CONCLUSION

Present work operationalises logistics optimisation within a real C3I/IMS, developed in the context of the SYNERGISE and FORESIGHT projects. The optimization module operationalises a MD-CVRP formulation within an event-driven crisis management workflow, integrating distributed warehouse inventories, heterogeneous fleet characteristics, compatibility constraints, geospatial routing, and explicit time modelling into a unified decision-support capability for dispatching and optimal routing of vehicles. By coupling inventory feasibility, capacity-respecting allocation, geospatial route computation, explicit ETA transparency, and structured diagnostics reporting, the system bridges the gap between optimisation research and operational crisis coordination platforms.

Future work will proceed along several complementary directions. First, algorithmic refinement will incorporate disruption-aware routing, time-dependent travel conditions, and richer multi-dimensional loading constraints, including volumetric and equipment-specific transport requirements. Second, time modelling will be extended to account for dynamic personnel availability, multi-stage handling processes, and context-sensitive operational delays. In addition, the scope of optimisation will be expanded beyond ground vehicles to include Unmanned Aerial Vehicles (UAVs) enabled logistics, enabling hybrid ground-air dispatch strategies for time-critical deliveries and access-constrained environments. Furthermore, staff logistics and personnel dispatch planning will be incorporated, linking first responders and specialised operators with assigned vehicles, UAV platforms, and equipment handling requirements to support integrated asset–personnel coordination. Moreover, the current prototype minimises estimated travel time as a proxy for rapid response, future development will involve a more holistic re-evaluation of the objective function in collaboration with civil protection and disaster relief user organisations. In particular, optimisation criteria will be expanded beyond time or distance minimisation to incorporate urgency levels, life-saving priorities, protection of critical infrastructure, and mission-critical asset allocation policies. This co-designed objective formulation will aim to better reflect the multi-criteria nature of crisis decision-making, balancing efficiency with equity and operational criticality. Finally, and most critically, system integration will be validated within the context of SYNERGISE and FORESIGHT projects through field demonstrations and structured operational and performance evaluation, with civil protection end-user organisations to ensure operational realism, interoperability, and alignment with real-world crisis workflows.

ACKNOWLEDGMENTS

The present work has been supported by the SYNERGISE project that is jointly funded from the European Union's Horizon Europe research and innovation programme, State Secretariat for Education, Research, and Innovation from Switzerland, the Japan Science and Technology Agency, the Korea Ministry of Science and ICT, and the Korea Electronics and Telecommunications Research Institute under the Grant Agreement no. 101121321. Present work has been also supported by the FORESIGHT Horizon Europe Project that is jointly funded from the European Union's Horizon Europe research and innovation programme (GA No. 101168521), and the State Secretariat for Education, Research, and Innovation from Switzerland.

REFERENCES

Auradkar, A. at al. (2012). Data infrastructure at LinkedIn. In Proceedings of the 28th IEEE International Conference on Data Engineering (ICDE 2012), Arlington, VA, USA, 1370–1381. https://doi.org/10.1109/ICDE.2012.147

Besiou, M., Pedraza-Martinez, A. J., & Van Wassenhove, L. N. (2021). Humanitarian operations and the UN Sustainable Development Goals. Production and Operations Management, 30(12), 4343–4355. https://doi.org/10.1111/poms.13579

Besiou, M., & Van Wassenhove, L. N. (2020). Humanitarian operations: A world of opportunity for relevant and impactful research. Manufacturing & service operations management, 22(1), 135-145. https://doi.org/10.1287/msom.2019.0799

Braekers, K., Ramaekers, K., & Van Nieuwenhuyse, I. (2016). The vehicle routing problem: State of the art classification and review. Computers & Industrial Engineering, 99, 300–313. https://doi.org/10.1016/j.cie.2015.12.007

Bray, T. (2017). The JavaScript Object Notation (JSON) data interchange format (STD 90, RFC 8259). Internet Engineering Task Force (IETF). https://doi.org/10.17487/RFC8259

Campbell, A. M., & Jones, P. C. (2011). Prepositioning supplies in preparation for disasters. European Journal of Operational Research, 209(2), 156-165. https://doi.org/10.1016/j.ejor.2010.05.003

Chopra, S., & Meindl, P. (2016). Supply chain management: Strategy, planning, and operation (6th ed.).

Pearson Education.

Comes, T., Hiete, M., Wijngaards, N., & Schultmann, F. (2011). Decision maps: A framework for multi-criteria decision support under severe uncertainty. Decision Support Systems, 52(1), 108-118. https://doi.org/10.1016/j.dss.2011.05.008

Coupey, J., Nicod, J. M., & Varnier, C. (2024). VROOM v1. 14, Vehicle routing open-source optimization machine. Verso. http://vroom-project.org

Holguín-Veras, J., Pérez, N., Jaller, M., Van Wassenhove, L.N. and Aros-Vera, F. (2013), On the appropriate objective function for post-disaster humanitarian logistics models. Journal of Operations Management, 31: 262-280. https://doi.org/10.1016/j.jom.2013.06.002

Ichoua, S., Gendreau, M., & Potvin, J. Y. (2003). Vehicle dispatching with time-dependent travel times. European Journal of Operational Research, 144(2), 379–396. https://doi.org/10.1016/S0377-2217(02)00147-9

Jahre, M., & Jensen, L.-M. (2010). Coordination in humanitarian logistics through clusters. International Journal of Physical Distribution & Logistics Management, 40(8/9), 657–674. https://doi.org/10.1108/09600031011079319

Kovács, G., & Spens, K. M. (2007). Humanitarian logistics in disaster relief operations. International Journal of Physical Distribution & Logistics Management, 37(2), 99–114. https://doi.org/10.1108/09600030710734820

Laporte, G. (2009). Fifty years of vehicle routing. Transportation Science, 43(4), 408–416. https://doi.org/10.1287/trsc.1090.0301

Montoya-Torres, J. R., Franco, J. L., Isaza, S. N., Jiménez, H. F., & Herazo-Padilla, N. (2015). A literature review on the vehicle routing problem with multiple depots. Computers & Industrial Engineering, 79, 115-129. https://doi.org/10.1016/j.cie.2014.10.029

OpenRouteService. (2024). OpenRouteService API documentation. https://openrouteservice.org/

OpenStreetMap. (2024). OpenStreetMap. https://www.openstreetmap.org

Pillac, V., Gendreau, M., Guéret, C., & Medaglia, A. L. (2013). A review of dynamic vehicle routing problems. European Journal of Operational Research, 225(1), 1–11. https://doi.org/10.1016/j.ejor.2012.08.015

Ransikarbum, K., & Mason, S. J. (2016). Goal programming-based post-disaster decision making for integrated relief distribution and early-stage network restoration. International Journal of Production Economics, 182, 324–341. https://doi.org/10.1016/j.ijpe.2016.08.030

Rawls, C. G., & Turnquist, M. A. (2010). Pre-positioning of emergency supplies for disaster response. Transportation Research Part B: Methodological, 44(4), 521–534. https://doi.org/10.1016/j.trb.2009.08.003

Simchi-Levi, D., Kaminsky, P., & Simchi-Levi, E. (2008). Designing and managing the supply chain: Concepts, strategies, and case studies (3rd ed.). McGraw-Hill.

Rottkemper, B., & Fischer, K. (2013). Decision making in humanitarian logistics – A multi-objective optimization model for relocating relief goods during disaster recovery operations. In T. Comes, F. Fiedrich, S. Fortier, J. Geldermann & T. Müller (Eds.), Proceedings of the 10th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2013), Baden-Baden, Germany, May 2013, pp. 647–656.

Toth, P., & Vigo, D. (2014). Vehicle routing: Problems, methods, and applications (2nd ed.). SIAM. https://doi.org/10.1137/1.9781611973594

Turoff, M. (2002). Past and future emergency response information systems. Communications of the ACM, 45(4), 29–32. https://doi.org/10.1145/505248.505265

Van de Walle, B., & Turoff, M. (2006). ISCRAM: Growing a global R&D community on information systems for crisis response and management. International Journal of Emergency Management, 3(4), 364–369. https://doi.org/10.1504/IJEM.2006.011302

Van De Walle, B., Turoff, M., & Hiltz, S.R. (2010). Information Systems for Emergency Management (1st ed.). Routledge. https://doi.org/10.4324/9781315703473

Van Wassenhove, L. N. (2006). Humanitarian aid logistics: Supply chain management in high gear. Journal of the Operational Research Society, 57(5), 475–489. https://doi.org/10.1057/palgrave.jors.2602125

Widera, A., Lechtenberg, S., Gurczik, G., Bähr, S., & Hellingrath, B. (2017). Integrated logistics and transport planning in disaster relief operations. In T. Comes, F. Bénaben, C. Hanachi, M. Lauras, & A. Montarnal (Eds.), Proceedings of the 14th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2017) (WiPe Paper – Logistics and Supply-Chain, pp. 752–764). Albi, France.