University of St. Gallen

School of Management

to obtain the title of

Bachelor of Arts in Business Administration

Bachelor Thesis on

Beyond Dots on a Map

Designing and Evaluating a Truck-First Visibility Prototype for B2B Road Freight Logistics

---

Submitted by:

Louis Hollinger

21-619-853

Approved on the application by:

Prof. Dr. Bruno Rodrigues

Date of Submission:

May 19, 2026

Acknowledgements

First and foremost I would like to express my deepest gratitude to my supervisor, Prof. Dr. Bruno Rodrigues. The month-long journey of iterating on the prototype would not have been possible without his weekly guidance and support. Our meetings gave the work structure and helped me refine the research direction. They also kept me motivated and in many cases his deep topic knowledge helped in unblocking my progress. His feedback encouraged me to connect the technical prototype more clearly to the underlying business problem and to keep the thesis grounded in the practical realities of logistics visibility.

I am also grateful to the interview partners who provided their valuable time to share their experience and perspectives with me. Their input helped me better understand the operational relevance of the problem and the conditions under which a visibility system could become useful in practice. The feedback and tips I got from many of the interview partners reached beyond what I could have imagined or expected.

Finally, I would like to thank my parents for their support throughout my studies and during the writing of this thesis. Their encouragement meant a great deal to me.

St. Gallen, May 19, 2026 Louis Hollinger

Abstract

This thesis investigates how visibility gaps in B2B road-freight handoffs can be reduced through a low-cost end-to-end visibility prototype. In third-party logistics settings, shipment information often deteriorates when goods leave the focal firm's direct operational control and enter partner-controlled transport legs. Existing tracking frequently relies on milestone scans and delayed status updates, leaving blind intervals that limit receiving preparation, exception handling, and status clarification. Following a design science research approach, the thesis derives requirements from practitioner input and logistics literature, translates them into a truck-first visibility architecture, implements a functional prototype, and evaluates it through technical tests, stakeholder feedback, and business-feasibility reasoning. The artifact combines lightweight shipment-side identification through BLE, RFID, or software-based association with reusable truck-side sensing. Then, backend interpretation of raw events into parcel-related state enables dashboard-based access with specified sharing across organizational boundaries. The evaluation shows that the prototype can capture, buffer, upload, and interpret BLE, RFID, GNSS, and liveness events at prototype level. The business analysis indicates that passive RFID and grouped-identifier scenarios are more plausible for broader deployment than disposable BLE identifiers attached to every parcel. The thesis contributes a practically relevant architecture for improving handoff visibility under current technical, economic, and organizational constraints. Its contribution lies in demonstrating a feasible architectural path toward better logistics visibility, while the system remains a prototype rather than a production-ready solution for universal parcel-level real-time tracking.

Table of Contents

List of Abbreviations

3PL	Third-party logistics
AI	Artificial intelligence
API	Application programming interface
B2B	Business-to-business
BLE	Bluetooth Low Energy
CSV	Comma-separated values
DSR	Design science research
EPC	Electronic Product Code
EPCIS	Electronic Product Code Information Services
ERP	Enterprise resource planning
ETA	Estimated time of arrival
FMCG	Fast-moving consumer goods
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HAT	Hardware Attached on Top
ICT	Information and communication technology
ITS	Intelligent transportation systems
IoT	Internet of Things
KPI	Key performance indicator
LTE	Long-Term Evolution
OEM	Original equipment manufacturer

PPK2	Power Profiler Kit II
QR	Quick Response
RFID	Radio-frequency identification
RQ	Research question
RSSI	Received Signal Strength Indicator
SQ	Subquestion
TDABC	Time-driven activity-based costing
TMS	Transportation management system
UHF	Ultra-high frequency
USB	Universal Serial Bus
UUID	Universally unique identifier
WISMO	Where is my order
WMS	Warehouse management system

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Problem Statement and Motivation

"The information about the package is just as important as the package itself" [24]. This statement is attributed to Frederick W. Smith, the founder of FedEx, and captures a long-standing problem in logistics in a simple form. Moving the goods is only one part of transport performance, while knowing where those goods are, what state they are in, and who can act on that information is another. In practice, however, logistics visibility still often breaks down when goods move across organizational boundaries [29] and when the actor who needs the information is different from the actor who currently controls the transport process [90, 93].

Modern supply chains rely heavily on networks of partners [85, 90]. Transport legs are outsourced, carriers subcontract to one another, and receiving operations depend on information produced by parties outside the focal logistics operator's direct operational control. This creates fragmented provider landscapes in which consistent end-to-end information availability is difficult to achieve in practice [90].

In such ecosystems, visibility often remains milestone-based. Shipments are observed at discrete scan points, which leaves blind intervals that become most consequential when goods move into partner-controlled transport legs [29, 90]. At these handoffs, incomplete capture and delayed or non-standardized data exchange reduce information quality [66, 90] (see also Appendices A.2 and A.4). The operational consequence is concrete. Receiving sites plan with weaker arrival evidence, exceptions are noticed later, and status clarification falls back to calls, emails, and partner-side checks [38, 39] (see also Appendices A.1, A.6, and A.7).

This motivates the development of a visibility concept that unobtrusively expands and integrates with preceding tracing systems by producing additional real-time data, while keeping the cost-sensitive and fragmented logistics setting in mind [90].

1.2 Research Objectives

This thesis designs and evaluates a low-cost visibility prototype for reducing information gaps in business-to-business (B2B) third-party logistics (3PL) handoffs. Its contribution is to translate this practical visibility problem into requirements, an artifact architecture, a working prototype, and a bounded feasibility evaluation. This design science research (DSR) path leads to the following main research question:

Main Research Question (RQ).

> How can a low-cost end-to-end visibility prototype be designed to reduce visibility gaps in B2B third-party logistics handoffs?

The main research question is answered through three supporting subquestions. Together, these subquestions structure the design, development, and evaluation of the artifact.

Subquestion 1 (SQ1). Which requirements and architectural choices enable a low-cost visibility prototype to address B2B third-party logistics handoff gaps under practical operational constraints?

Subquestion 2 (SQ2). How can a visibility prototype turn low-level tracking signals into decision-relevant information for B2B logistics handoffs?

Subquestion 3 (SQ3). Under what conditions can the operational value of visibility data justify the cost of implementing a low-cost tracking prototype?

1.3 Scope and Structure

These questions define the analytical route of the thesis, but they are answered within a deliberately limited scope. The thesis focuses on B2B road-freight processes in which visibility gaps emerge at inter-organizational handoffs and during third-party transport

execution [29, 90, 93]. Its scope is the design and feasibility-oriented evaluation of a low-cost visibility prototype under realistic operational constraints. The thesis therefore examines whether such an artifact can be made technically and operationally plausible, rather than trying to answer whether it is ready for full-scale deployment.

The empirical grounding comes from a small, purposive set of expert inputs from companies active in the German-speaking logistics context. These inputs are used to derive requirements and assess plausibility, while the more detailed evaluation combines prototype evidence, stakeholder feedback, and business-feasibility reasoning. The resulting claims are therefore analytical and context-sensitive rather than statistically generalizable across all logistics settings [75].

The thesis follows a DSR framework because it addresses a practical visibility problem through the design and evaluation of an artifact [36]. Chapter 2 develops the conceptual and related-work foundation, Chapter 3 explains the research design, Chapter 4 derives and implements the artifact, Chapter 5 evaluates it, and Chapter 6 summarizes contributions, limitations, and future work.

Chapter 2

Fundamentals and Related Work

Road-freight visibility is not created by movement data alone. A vehicle position becomes operationally meaningful only when it can be related to the right object, the right process moment, the right organizational context, and be relayed to the right party [25, 28]. A coherent account of road-freight visibility, therefore, has to separate visibility from tracking, freight objects from transport resources, raw observations from logistics events, and information access from unrestricted transparency.

2.1 Fundamentals

2.1.1 Visibility, Tracking, and Traceability

Supply-chain visibility is sometimes used as a synonym for live tracking, but the literature treats it as a wider information capability. Francis defines supply-chain visibility as knowledge about the identity, location, and status of entities moving through a supply chain, captured through timely event messages that relate planned and actual times [25]. The definition is important because it already contains more than location. Visibility concerns an entity, a status, a time relation, and a process expectation. A map point can show where a vehicle is, but it does not by itself say which goods are affected, whether a transfer has occurred, or whether an operational decision should change.

The value of visibility also depends on information quality. Shared information must be accurate and available in a usable form before it can improve supply-chain performance [9]. Swink et al. describe visibility quality in similar terms, specifically underlining access to timely and complete formatted data from internal or external sources [80]. Data abundance does not by itself produce visibility, however. Information has to be interpretable and decision-relevant at the moment when intervention is still feasible.

This distinction helps separate visibility from neighboring terms. Tracking usually refers to following the current position or status of an object. Traceability and visibility are often used interchangeably, even though they emphasize different aspects of supply-chain information sharing [74]. Visibility primarily concerns whether relevant information about a flow is accessible for monitoring and decision-making in real or near real time, while traceability primarily concerns the reconstruction of a past path or chain of events.

Transparency reaches further. It implies disclosure across organizational boundaries and can include information beyond what is strictly necessary for a specific operational decision. Interorganizational information transparency is shaped by dependence between firms, by the access privileges, and by the compatibility between different information systems [15].

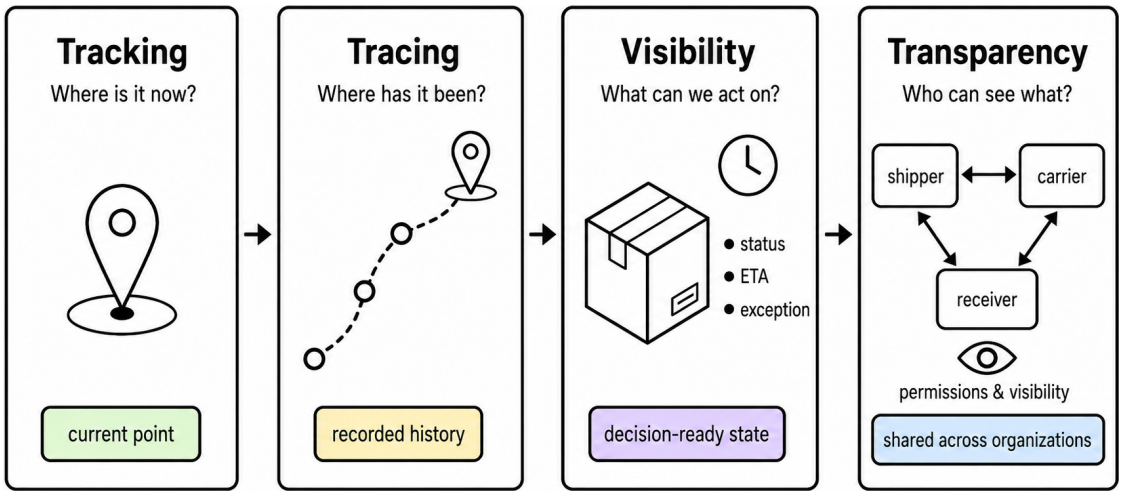

Figure 2.1: Conceptual distinction between visibility terms.
Source: Original representation generated with OpenAI ImageGen-2.

In logistics contexts, these terms often overlap in practice, but this analytical separation is important because it prevents inflated claims. A system which gives a receiver earlier arrival information does not automatically provide full traceability, for example. The visibility level depends on the object being observed, the quality and timeliness of the information, and the organizational boundary across which it is shared.

2.1.2 Freight Objects, Identity, and Association

Visibility is always visibility of something. Road-freight operations involve several object levels at the same time: vehicle, trailer, tour, transport leg, consignment, shipment,

handling unit, parcel, and item [25, 28, 35]. These levels are operationally connected, but they are not identical. A truck may be directly located through vehicle telematics, while the relevant business question may concern whether a particular parcel has passed a specific scan-gate.

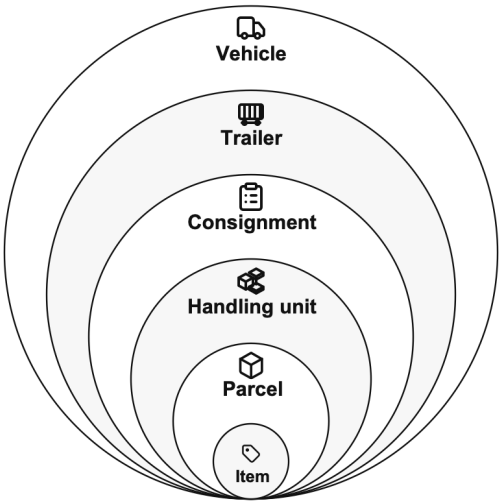

Figure 2.2: Nested visualisation of selected road-freight objects. Lower-level freight objects are interpreted within higher-level transport contexts.

Source: Original representation based on [25, 28, 35].

Freight-ITS literature makes this object-level distinction explicit. Giannopoulos describes the development of freight information systems as a movement from information at the level of transport means toward information at the level of load units and individual items [28]. Smart tracking devices have historically been associated with larger transport assets such as trucks and containers, with less coverage of pallets or individual packages [35]. Taken together, this indicates a gradual shift from vehicle-centric toward freight-centric visibility [28, 35]. Vehicle-centric visibility observes the transport resources, while freight-centric visibility concerns lower-level freight objects whose state matters in the logistics process. Vehicle-centric information becomes freight-centric only when there is a credible connection between the vehicle and the freight object. Through manual assignment, an identifier read, or an explicit event association, for example [31, 35].

2.1.3 Sensing and Event-Based Logistics Information

Road-freight visibility often begins with observations produced by sensing or identification technologies. Freight-ITS research has long treated mobile data transmission and the Global Positioning System (GPS) as important drivers of automatic vehicle location and dispatch applications [28]. Road-freight Internet of Things (IoT) research similarly discusses location sensing, asset tracking, and fleet-management data as capabilities through which firms can monitor commodities, drivers, vehicles, and traffic information [23]. Mobile telematics extends this perspective by treating trucks, containers, and loading units as IoT-enabled transport assets with localization and sensing functions [17, 34].

Different identifier technologies create separate forms of object evidence. Table 2.1 summarizes the main trade-offs between optical identifiers, passive radio-frequency identification (RFID), and active radio identifiers such as Bluetooth Low Energy (BLE) beacons for parcel-level visibility.

Table 2.1: Comparison of selected freight-object identification technologies, synthesized from logistics-technology literature [11, 17, 35, 95].

Technology	Typical evidence contribution	Capture or infrastructure burden	Main limitation
Barcode and QR code	Low-cost object identity at a defined process point	Manual scan or scan-gate, vision-based capture	Evidence exists only when the code is scanned and remains dependent on process discipline
Passive RFID	Object identity without line of sight, often through gate or reader-zone capture	Dedicated readers, antenna placement, and controlled read zones	Read performance depends on orientation, materials, distance, and reader configuration
Active RFID and BLE beacon	Repeated radio signal from an active object-side device	Nearby receivers, gateways, phones, or vehicle-mounted base stations	Better recurring presence evidence, but with battery and cost trade-offs

The difference between observation and event is central. A location fix, RFID read, BLE advertisement, Barcode scan, or manual confirmation is an observation. A departure,

arrival, exception, or handoff status is an interpreted logistics event or state. Supply-chain event-management literature adds a second lens. Exceptions can be understood as deviations arising from the gap between planning and execution, and event analysis is needed to support real-time decision-making [46]. Logistics exception handling also requires more than low-level event data. They need state information because systems need to support detection, interpretation, and response [38]. This literature helps explain why raw sensing data rarely carries operational value by itself. It becomes useful when it is interpreted into something that can be acted upon.

Estimated time of arrival (ETA) prediction illustrates this distinction further. ETA can be defined as reliable, up-to-date arrival-time prediction. Existing track-and-trace systems, however, often show current location information without explaining how the remaining transport is likely to unfold [8]. Heinbach et al. include ETA in their freight intelligence platform [34]. Arrival prediction and exception detection are therefore downstream uses of visibility data.

2.1.4 Interorganizational Visibility and Handoffs

Visibility becomes harder when transport information crosses organizational boundaries. In 3PL settings and subcontracted road freight, information fragmentation is the normal state of affairs, and the involved parties may therefore each hold different pieces of the operational picture.

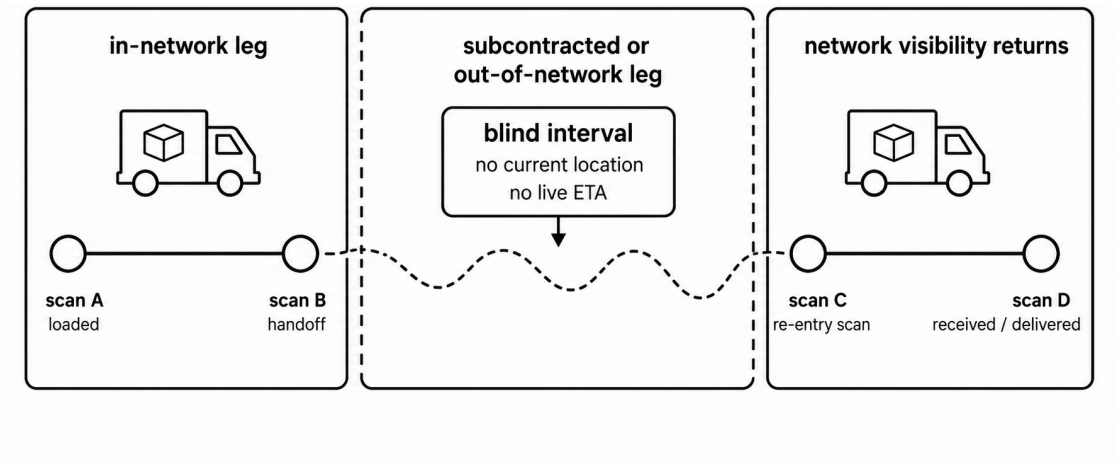

Figure 2.3: Representation of the visibility gap between scan points during a subcontracted transport leg.

Source: Original representation generated with OpenAI ImageGen 2.

Interoperability literature frames this as a structural problem. Pan et al. define interoperability as the ability of independent logistics and supply networks to conduct operations and business with one another, and they connect visibility, traceability, and transparency problems to interoperability gaps [66]. Standards such as GS1 and Electronic Product Code Information Services (EPCIS), alongside application programming interfaces (APIs) and platform architectures, are presented there as operational mechanisms for exchanging logistics data. But interoperability requires more than the technical layers to be compatible with one another.

Data sharing across firms is constrained by data governance and data sovereignty. These topics have become central concerns, since actors need ways to collaborate without releasing confidential information [15, 27]. Information sharing and process coordination are therefore important elements of logistics integration that are limited by contractual and relational governance between 3PL users and providers [88].

Work on information sharing with 3PLs summarizes the interorganizational barriers. Outsourced logistics processes depend on trust and the technological capacity of participating firms [85]. The sources point to a basic tension. Visibility is valuable because information moves across firm boundaries, but exactly those boundaries make unrestricted disclosure unrealistic [15, 64].

2.2 Related Work

2.2.1 Visibility Gaps in Road Freight

Inbound and road-freight studies show that visibility gaps are operationally consequential. Kalaiarasan et al. study inbound logistics visibility and motivate their work through the difficulty of managing deviations, especially changes in arrival time [39]. Their companion study identifies blind spots in inbound flows and emphasizes the need to capture key process steps [40]. Their main contribution lies in showing how missing or delayed status evidence affects the ability to react to deviations before they affect downstream operations.

Outbound logistics provides a concrete example of this limitation in Bolte and Goll's retailer case: an unclear handover point and missing loading manifest create a black-hole interval between the retailer and postal service provider [29]. The case is relevant here

because it shows that checkpoint scans alone do not resolve visibility gaps when the transition between actors is not represented by reliable handover or loading evidence.

Real-time visibility is examined by Wycislak in a complex transportation network, where subcontracting is identified as a strong factor affecting real-time visibility compliance [93]. The study links visibility issues to missing enterprise resource planning (ERP), transportation management system (TMS), and warehouse management system (WMS) connectivity, inadequate driver resources, privacy concerns, and the lack in willingness of carriers to share commercially sensitive information. A later deployment study makes these obstacles more concrete by showing how manual truck data collection and asymmetric costs and benefits between the parties can slow real-time visibility platform adoption [92].

Road-freight digitization research adds a broader adoption context. Road freight is fragmented and strongly oriented toward cost minimization, which can slow the adoption of digital technologies [69]. Molero et al. make this implementation problem more concrete. Transport information and communication technology (ICT) projects can be slowed by equipment costs, ICT-training effort, and uncertainty over returns [51]. These studies help explain why visibility concepts that appear technically mature may still be difficult to realize in everyday road-freight settings. Whether a tracking technology exists is one part of the problem. Whether it fits existing environments is another.

2.2.2 Standards, Platforms, and Visibility Services

Event standards address the representation side of visibility. EPCIS and GS1 work provides structured concepts for how logistics events can be represented so that different actors and systems can interpret them without complex translation layers [2, 31, 73]. EPCIS-based visibility services can be organized in a variety of architectural setups, and standardized visibility data can support reuse across distributed supply chains [82].

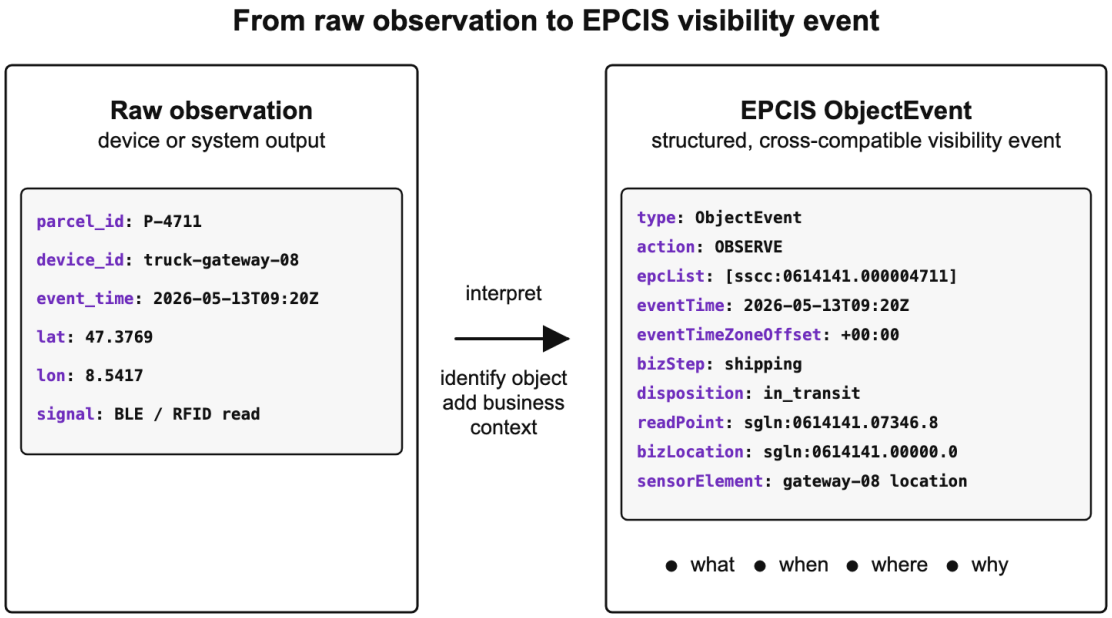

Figure 2.4: Translation of a raw parcel and location observation into an EPCIS-style ObjectEvent.

Source: Original representation based on the EPCIS standard [31]

Visibility platforms and control towers address the coordination and decision-support side. Supply-chain control towers can be understood as multi-tier visibility and exception management systems [67] and have also been described as socio-technical hubs that combine technology, organization, and processes to capture and use supply-chain data for decision-making [86]. Road-freight transport management platforms can include services for shippers, carriers, freight forwarders, and subcontractors [33]. These works lay out a blueprint of how visibility data can become part of broader ecosystems.

Real-time transportation visibility platforms occupy a similar space. Such platforms depend on carrier participation, telematics or enterprise-system integration, and the handling of privacy and commercial-sensitivity concerns [93]. Heinbach et al.'s freight intelligence platform similarly displays telematics-enabled data as a variety of services [34]. These approaches give insight toward what kind of information can be derived from collected visibility data.

Technology-oriented work contributes component-level evidence about RFID infrastructures, BLE proximity signals, and mobile telematics [11, 34, 95]. These streams are especially useful for the design of the artifact's hardware, but they often lack the

combined, integrated evaluation of the components in the specific setting of a visibility prototype.

Direct technical precursor work is represented by Coita's low-cost parcel-tracking prototype, which evaluates off-the-shelf RF technologies under a target component budget of three to five euros per parcel module [17]. Coita's study mainly asks whether simple parcel modules can communicate reliably with a mobile base station in a delivery-vehicle setting, and whether signal continuity or signal loss can serve as evidence of parcel presence or removal. It is therefore useful here as an empirical reference for low-cost parcel-side identification and for the trade-offs between cost, power consumption, and communication reliability. The present thesis builds from a different level of analysis. RF-based parcel evidence is treated as one input to a wider visibility platform. The different layers have to work together to create better third-party handoff visibility.

Interorganizational data-sharing research adds the governance boundary. Standards and APIs help independent networks exchange information, but visibility services must also preserve data sovereignty, access privileges, appropriate granularity, and trust [15, 27, 66].

2.3 Synthesis of the Literature

The literature establishes a coherent but incomplete picture. It shows that road-freight visibility requires a chain that connects observations, object identity, event interpretation, and controlled information sharing across organizational boundaries.

The remaining difficulty is at the intersection of these streams. Vehicle tracking is mature, although vehicle visibility does not produce freight visibility by itself. Event standards provide strong semantics, but they do not guarantee that the necessary association evidence exists in a given transport context. Platforms can distribute visibility data, but often assume partner participation, system integration, or data availability. Identifier technologies can strengthen object evidence, although each introduces its own particular constraints that must be smoothly navigated. Lastly, administrative frameworks explain why sharing must be bounded while still allowing adequate data to be transferred.

For road-freight handoffs across organizational boundaries, the specific question is how reliable freight-object visibility can be maintained once the object leaves the focal firm's

direct control and evidence becomes distributed across a variety of vehicles, partners, and information systems.

Chapter 3

Methodology

3.1 Research Design

This thesis follows a DSR logic. In the information systems literature, Hevner et al. describe DSR as the creation and evaluation of artifacts intended to solve identified organizational problems [36]. Peffers et al. [68] describe the DSR methodology as a six-step process that is illustrated in Figure 3.1.

This sequence fits the present thesis because the work develops and evaluates a prototype response to the identified logistics visibility problem. The prototype is therefore treated as the DSR artifact through which the problem is translated into requirements that form the basis for an artifact design that then later gets evaluated against the objectives.

The research design is therefore neither a pure analytical study nor a pure implementation report. *Problem identification and motivation* are addressed through the framing of visibility gaps in B2B 3PL handoffs as a relevant and practical research problem. The *objectives for a solution* are defined by combining expert interviews and targeted literature review into a requirement base that translates problem evidence into design goals and constraints. *Design and development* are where the translation of these requirements into a coherent prototype concept happens. *Demonstration* and *Evaluation* are finally addressed through the assessment of the resulting prototype, including technical feasibility checks, stakeholder feedback, and requirement-oriented evaluation.

An important implication of this application of DSR is that the process is not fully linear. In artifact-oriented work, some design questions only arise once an initial concept is translated into a working prototype and confronted with constraints. For that reason, this thesis treats development as an iterative process with refinement loops. These refinements

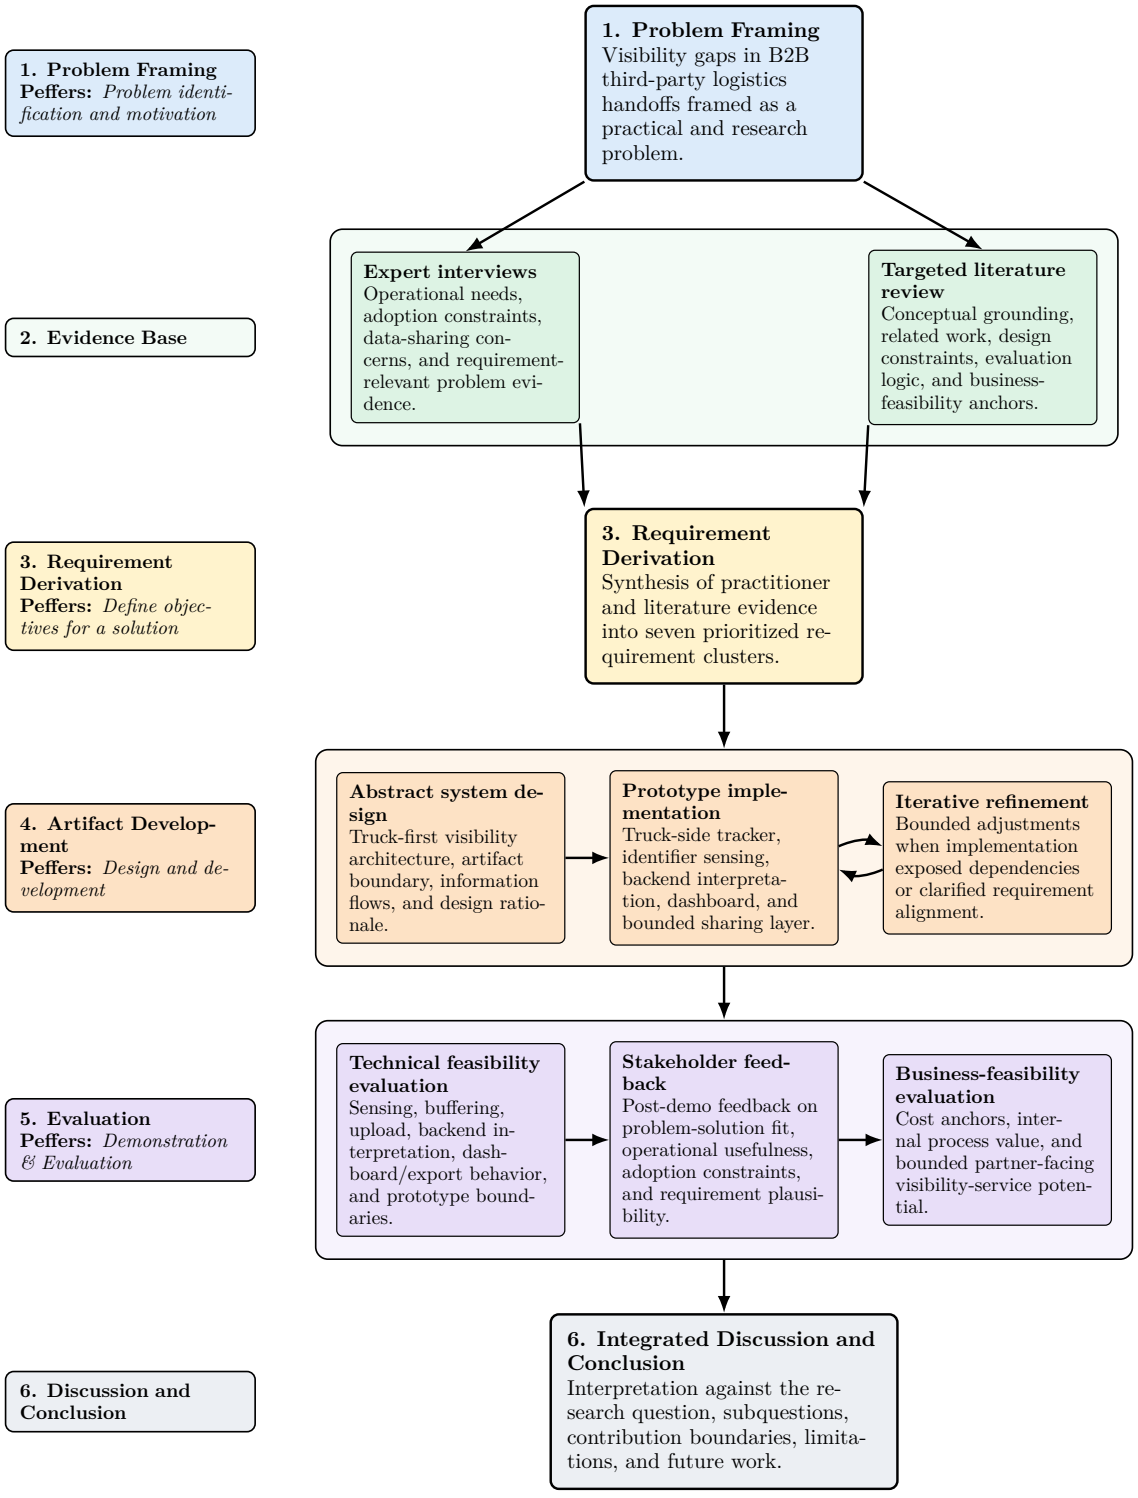

Figure 3.1: The main research phases and their alignment with the core design science steps adapted from Peffers et al.

Source: Original representation based on Peffers et al. [68].

were accepted only when they improved alignment with the derived requirements without changing the research problem after the fact.

3.2 Empirical and Literature Inputs

3.2.1 Expert Interviews

The empirical foundation of the thesis consists of six practitioner sources from logistics or logistics-adjacent operational settings. Five of these sources were collected through semi-structured expert interviews, while one was collected via a structured written questionnaire. The participant selection process began with outreach to approximately 30 logistics-related organizations in the German-speaking area, but only six accepted. The resulting sample is therefore purposive and access-based and not statistical.

This sample includes perspectives spanning from receiving and inbound operations to freight organizers to systems-integration contexts. This variety is important because visibility gaps in third-party transport emerge across organizational boundaries and therefore benefit from input that reflects several sides of the transport chain.

The semi-structured interviews were conducted with the help of a prepared, but not shared, question guide. Participants were encouraged to elaborate by providing examples and to explain how visibility problems appear in their practice. Methodologically, the interviews were therefore expert-oriented and problem-discovery oriented.

The spoken interviews were recorded after participant consent and then transcribed, redacted, and cleaned into analysis-ready transcript files. The written questionnaire response was retained to provide different perspectives on the process and problem setting, despite its detached collection format.

To move from raw interview findings to evidence, the material was then analyzed using qualitative content analysis by assigning relevant transcript excerpts to data-near categories. These initial categories were then consolidated into higher-level findings and requirement-relevant clusters for the artifact design.

For the demonstration and the evaluation, practitioners were later recontacted to give feedback on the prototype concept direction. These practitioners assessed whether the

designed response remained plausible and useful to solve the operational concerns they had helped surface earlier.

3.2.2 Literature Review and Source Selection

Hevner et al. emphasize that DSR has to connect a relevant problem environment with established foundations and methods [36]. The interviews, therefore, supplied evidence of practical relevance, while the literature helped determine how far these observations aligned with existing concepts and comparable artifact classes. It further revealed known constraints in logistics visibility and information systems research.

A source was relevant when it clarified the visibility problem or elaborated on a plausible solution. Within Peffers et al.'s [68] process, the targeted literature review therefore mainly supports the movement from *problem identification and motivation* to *objectives for a solution*. It gives the thesis reference points for judging feasible solution directions and criteria for assessing whether the proposed artifact is a meaningful response.

3.3 Requirement Derivation and Traceability

Requirement derivation is where the DSR *define objectives for a solution* step becomes thesis-specific. For this thesis, the requirement process translated interview evidence and literature findings into objectives that could guide a designed artifact.

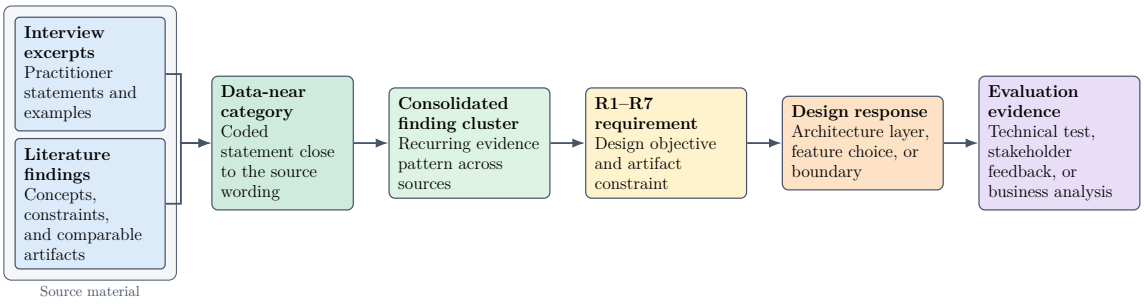

Figure 3.2: Traceability chain linking empirical and literature evidence to requirements, design responses, and evaluation evidence.
Source: Original representation.

The consolidated requirement-relevant clusters from the interviews were further combined with the literature base. This resulted into seven prioritized requirements. This

structure gives the artifact a defensible target for design and later evaluation [36, 70]. Traceability turns this translation into an inspectable chain of reasoning. By linking evidence, requirement statements, and design responses, the thesis can show why a design choice belongs in the artifact and what kind of evidence supports it.

3.4 Artifact Development and Iterative Refinement

Artifact development is where the seven requirements are addressed by an end-to-end visibility prototype. Its methodological role is to make the proposed response concrete enough to be demonstrated, refined, and evaluated.

The development journey can be described in a three-part process, starting at an abstract system design, then moving towards a concrete prototype implementation that is then iteratively refined throughout the process. This build process functions methodologically as the point at which design assumptions, integration dependencies, and feasibility constraints became inspectable and evaluatable.

3.5 Evaluation Design

Evaluation design follows directly from the artifact boundary and the claims the thesis can defend. DSR requires the utility of an artifact to be demonstrated through appropriate evaluation methods [36]. The evaluation is therefore organized into three parts. The technical evaluation tests hardware and software behavior. The practitioner feedback probes for the problem-solution fit. The business analysis explores plausible value mechanisms.

The business-feasibility perspective is included in the evaluation because practical relevance is part of the DSR claim. It asks whether the visibility information generated by the prototype could support operational value. This part can be looked at as a literature-informed analytical extension to the evaluation. It remains qualitative and mechanism-based in identifying plausible economic value paths.

3.6 Limitations and Ethical Considerations

On an empirical level, the thesis favors depth of problem grounding and traceability over pure statistical breadth. The interview participant base is small and only covers a

German-speaking operational context. The length of the interviews and the open format, however, allowed the answers to be thorough and detailed.

The artifact evaluation is constrained by the thesis scope. Because this thesis only created a prototype, claims from the evaluation cannot be transferred to real-world settings. The performed evaluation only applies to the narrow scenarios tested, and further deployment studies would have to be performed before the prototype could be turned into a product.

The same researcher built and evaluated the artifact. This can cause researcher proximity bias. The thesis addresses the resulting risks through explicit evidence traceability. Claims require literature grounding. Requirements are assessed openly and are based on feedback from practitioners. Additionally, safeguards provide several points at which the chain from evidence to claim can be inspected.

From an ethical viewpoint, considerations about consent, confidentiality, and sensitive company information need to be taken into account. Spoken interviews are only recorded after having obtained participants' explicit consent. We identified interview participants through the organization only. The interview transcripts were redacted where needed. Additionally, to avoid exposing company-specific details that might reveal sensitive information about the participating companies' operations, the transcript generalized or completely omitted this kind of information.

Chapter 4

Design and Implementation

This chapter first synthesizes the design logic that shapes the prototype, then presents the end-to-end concept and architecture, and finally documents the implementation layer by layer. The purpose of the chapter is to show how the thesis problem was turned into an evaluatable prototype rather than to present a production-ready logistics platform. This chapter mainly answers **SQ1** by specifying the artifact design and **SQ2** by showing how event and state data goes from captured raw data to interpretable state, while **SQ3** is evaluated more directly in Chapter 5.

4.1 Requirement Synthesis with Literature Backing

This section condenses the empirical and literature evidence into seven requirement clusters. The requirements are derived by combining interview evidence with literature findings into clusters that are empirically recurring and relevant for design decisions within the prototype scope. The final requirement clusters will finally be subdivided into functional and non-functional requirements in table 4.1.

4.1.1 R1: Visibility Must Reduce Blind Intervals Between Milestones

Several interviews indicate that visibility today still usually relies on scan or milestone status events. The interview partner from the large international logistics company described current tracking as a milestone chain created through scans at operational touchpoints (e.g., pickup, dispatch, final delivery, etc.). The same interview characterized the outcome of this process as evidence that a parcel was scanned at a specific time and place, rather than as continuous insight into the transport leg itself (Appendix A.2). The Quantifor interview partner pointed to a similar data gap in handoff processes between organizations.

When a small e-commerce retailer hands a package to a carrier, the first recorded tracking point may only appear later in the carrier's sorting network. If the parcel is lost before that point, the preceding transfer remains invisible (Appendix A.4). The Brack Alltron AG interview partner described the same problem for outsourced distribution, where the commissioning firm may lack visibility over what an external logistics provider will actually deliver the next day (Appendix A.1). These demands, therefore, highlight an operational problem created by the time gap between scan milestones.

The literature supports this requirement by treating visibility as timely, decision-relevant event information rather than as a historical scan archive. Francis, Somapa et al., and Callefi et al. all mention that visibility depends on accessible, useful information about the moving goods and the road-transport processes [12, 25, 79]. Wycislak adds that real-time visibility becomes necessary where fragmented carrier systems and weak intercompany integration prevent shippers from observing the current transport situation [90]. **R1** therefore treats structured visibility events as a method in which the artifact can fill the intervals between formal scans.

4.1.2 R2: Arrival and Load Information Must Support Receiving Preparation

The interviews repeatedly link visibility to preparation before physical arrival. The Valora Gruppe interview partner explained that a store may know a two-hour delivery window but still lack precise knowledge of when a shipment will actually arrive. The same interview illustrates the practical limitation of short-notice alerts: notifications by phone or email shortly before arrival can be helpful, but staff and store processes often have to be planned earlier than the final half hour (Appendix A.5). The Galliker written response shows that inbound work already depends on structured booking data, including supplier, customer, expected delivery date and time, and ramp information (Appendix A.6). The Brack Alltron AG interview revealed that when inbound goods arrive without a reliable forecast of volume and composition, temporary staffing cannot be adjusted for sorting those goods into inventory (Appendix A.1). The Manor interview partner added that knowing what arrives and when matters for store replenishment, while even short delays in fresh food can create far worse operational issues (Appendix A.3).

Inbound-visibility literature supports this requirement directly. Kalaiarasan et al. derive requirements around data collection, data handling, and decision-oriented data use, and

later show that insufficient visibility particularly weakens deviation management when expected arrival times change [39, 40]. Wycislak and Pourhejazy reach a similar conclusion in a dock-booking and control-tower context, where inbound efficiency depends on real-time data integration, operational preconditions, dynamic responsiveness, and intercompany integration [94].

The requirement, therefore, demands more than arrival timing as it becomes only truly actionable when it is connected to information about what is arriving, and whether a deviation requires replanning.

4.1.3 R3: Operational Fit and Retrofitability Must Constrain the Design

The interviews show that visibility must accommodate fragmented B2B logistics environments. The interview partner from the large international logistics company described heterogeneous subcontractor practices: some use handheld, connected scanning devices while others provide reports by email on request. There is no common standard across the partner network. Physical adoption has its own boundary. A tracker is only viable if it is sufficiently affordable and disposable, or if a closed-loop process allows it to be recovered. Otherwise, the device and its recovery process become part of the cost and adoption problem (Appendix A.2). Manor's CTO added an organizational limit by explaining that supplier onboarding is difficult across a base ranging from global corporations to local firms without mature ERP systems (Appendix A.3). The Valora Gruppe interview partner similarly emphasized that visibility data would not sit in an isolated transport process, but between the logistics provider and other partners along the supply chain (Appendix A.5).

The adoption literature supports this as an assimilation constraint rather than only a technical design choice. Tracking technologies create benefits within organizational and interorganizational conditions rather than independently of them [10]. Distribution-logistics and road-freight studies further show that visibility initiatives face persistent barriers such as resource limits, limited expertise, resistance to change, and broader IoT implementation challenges [5, 26, 32].

The artifact must therefore fit existing routines and partner heterogeneity while keeping operational burden low enough and integration simple enough for adoption in the fragmented B2B shipping industry.

4.1.4 R4: Cost Must Be Justified by Operational Use

Cost arises throughout the interviews as a boundary for viable visibility features. The Quantifor interview partner noted that even inexpensive trackers become costly when deployed at large scale. Relevant cost drivers include infrastructure, communication, computation, and human resources. The same interview expresses an important limitation: the use case must be defined first, after which tracking can be assessed for economic justification against cheaper alternatives (Appendix A.4). The interview partner from the large international logistics company similarly noted that customer willingness to pay for visibility add-ons is limited (Appendix A.2). Galliker's written response provides the positive side of the same logic: better preparation of personnel resources can support internal process improvement and cost optimization (Appendix A.6). Cost is therefore not an isolated purchasing issue. It must be assessed in relation to the operational use that visibility data enables.

Road freight and tracking technology literature supports this economic constraint. Road freight is fragmented and cost-sensitive, which can slow digital adoption and makes additional visibility services dependent on profitability and competitiveness [33, 69]. Asset-visibility and RFID studies add that implementation value must be assessed against installation cost, unit economics, break-even points, and return on investment [23, 44, 65]. Interorganizational tracking projects can also distribute costs and benefits unevenly across actors, especially in 3PL contexts [87].

These findings justify treating visibility granularity as an economic design decision rather than collecting the technical maximum amount of data. The requirement is therefore to match sensing granularity to operational decisions rather than to maximize data capture for its own sake.

4.1.5 R5: Raw Events Must Become Usable Operational Information

Multiple interviews suggest that simply collecting more data does not solve the visibility problem. The Quantifor interview partner described logistics data as heterogeneous, incomplete, and distributed across disconnected systems. The challenge is to bring that information into a state in which it can be analyzed. The same interview argued against extracting everything from isolated systems and instead emphasized standardized interfaces for the most important information (Appendix A.4). The Valora Gruppe

interview partner gave a practical example: if GS1-compliant electronic messages are not returned on time, manual booking may be required, and downstream store goods-receipt or sales processes may be blocked (Appendix A.5). The Brack Alltron AG interview partner framed a similar structuring problem in receiving, where delivery notes, carrier information, and internal actors must be reconciled (Appendix A.1). The Manor interview showed the integration dimension: status data can flow from carrier systems into ERP systems, but only where the corresponding interfaces exist (Appendix A.3). Galliker's written response similarly stresses the need for complete documents and good data quality, because missing or incomplete information can block goods until clarification work is completed (Appendix A.6). These points show that the problem extends beyond having the right interface. Data must become intelligible enough for people and systems to act on it.

The literature supports this as a capture-to-information requirement. EPCIS and visibility-service studies show why event capture, access, standard data formats, and service interfaces matter for interpretable supply-chain transparency [2, 82]. Interoperability and event-processing research add that data exchange must be reliable across systems and organizations, and that heterogeneous supply-chain events have to be converted into process-relevant information before they can guide action [41, 66].

4.1.6 R6: Visibility Information Must Be Reliable and Trustworthy Enough for Use

The interviews indicate that visibility data becomes more useful when users actually can rely on it. Galliker's written response links efficient flow of parcels and supplier evaluation to good data quality, while also noting that missing data can delay entire processes because goods remain blocked until clarification is resolved (Appendix A.6). The Valora Gruppe interview partner reinforces this dependency in system terms: missing or incorrectly booked electronic messages can trigger manual workarounds and hinder downstream processes (Appendix A.5). The Quantifor interview partner describes the data foundation as decisive, because data must be sufficiently complete, related across systems, and placed on a reliable basis before it can support analysis. Automation is also out of scope until the data is proven to be reliable (Appendix A.4). The interview partner from the large international logistics company adds a practical capture concern: manual scans may be forgotten, performed incorrectly, or delayed through partner processes

(Appendix A.2). Reliability is therefore just as much about eliminating errors as it is about ensuring that information remains dependable and understandable when capture conditions are imperfect.

The literature supports this requirement by treating reliability as a central part of visibility quality. Swink et al. and Kalaiarasan et al. emphasize accuracy and completeness as conditions for visibility to support decisions [39, 80]. Pan et al. add the interoperability perspective by stressing secure, reliable, and traceable data exchange, while Wycislak shows operationally that incomplete timestamps, delayed updates, inaccurate addresses, and wrong planning data can prevent real-time visibility benefits from materializing [66, 90]. The artifact should therefore present reliable information without turning imperfect capture evidence into false certainty.

4.1.7 R7: Data Must Be Shared Selectively Across Ownership Boundaries

The interviews reveal a tension between the need for cross-company visibility and the need to protect internal operational information. The Brack Alltron AG interview partner describes external freight forwarders as serving several firms at the same time, while the commissioning firm may not receive complete information about incoming vehicles, package volumes, or delivery composition (Appendix A.1). The interview partner from the large international logistics company describes customers who move shipments through multiple logistics service providers and prefer to view information in their own software environment rather than logging into each provider portal separately. At the same time, the interview identifies the difficulty of getting subcontractors to release data to external systems or allow tracking devices on their vehicles (Appendix A.2). The Manor interview partner makes this boundary concrete through the supplier-portal example: delivery notes, invoices, and shipment-status updates could be standardized, but only if external suppliers and forwarders participate in a shared access surface or expose the required information through interfaces (Appendix A.3). The Quantifor interview partner adds the strategic reason for this reluctance. Companies create competitive advantage through efficient internal processes and do not want to expose those processes beyond the firm (Appendix A.4).

The literature supports treating selective sharing as a boundary-design problem. In 3PL contexts, information sharing can improve collaboration, competitive advantage, and customer service, but visibility mechanisms still depend on trust and commitment [9,

85]. ICT-enabled sharing is further shaped by both trust and distrust, while collaborative information sharing faces barriers around data security and compatibility [43, 50]. Data-centric service models may create additional value from logistics data, but only when this data is shared in a controlled way [98]. The requirement is therefore controlled cross-company usefulness: shipment-related status, exception signals, proof-relevant history, and bounded export access should be shareable without dissolving data ownership or exposing proprietary process knowledge.

Table 4.1: Design-oriented translation of requirement clusters into functional and non-functional requirements

Req.	Derived functional requirement	Derived non-functional requirement
R1	The artifact shall capture transport- and handoff-relevant events between formal scan points and delayed partner milestones.	The artifact shall remain useful even when capture is intermittent, delayed, or partly manual.
R2	The artifact shall provide arrival, load, and deviation information in a form that supports receiving preparation and pre-unload decisions.	The information shall be timely and clear enough for staffing, space, ramp, and escalation planning.
R3	The artifact shall support visibility generation within existing workflows and heterogeneous partner inputs.	The artifact shall be retrofittable without extensive fixed infrastructure, full system replacement, or major partner-side process changes.
R4	The artifact shall provide decision-relevant visibility at a bounded level of granularity.	Total system cost and operating burden shall remain proportional to the value created in low-margin, high-volume settings.
R5	The artifact shall transform raw sensing and handoff observations into structured, decision-relevant visibility information.	The information logic shall remain interpretable, interoperable, and manageable across incomplete and heterogeneous source data.
R6	The artifact shall preserve evidence, provenance, and uncertainty cues so that users can distinguish observed events from derived visibility states.	The resulting information shall remain sufficiently reliable, understandable, and robust under imperfect capture conditions.
R7	The artifact shall provide role-aware access and bounded sharing across organizational boundaries.	Data ownership, confidentiality, and partner autonomy shall be protected.

4.2 Abstract System Design

R1 and **R2** describe the need for more than just additional tracking data. They require visibility to become parcel- and load-relevant before receipt. **R5** to **R7** add further demands: captured observations have to become reliable, selectively sharable, and structured, operational state information. **R3** and **R4** set the decisive boundary conditions to make the artifact plausible in real-world operations. The design problem is therefore one of locating sensing, interpretation, and access in a way that creates additional certainty without multiplying infrastructure and process burden.

An autonomous tracker-on-every-parcel design would have to place components for sensing, communication, and power on each shipment unit. The maintenance and recovery burden would similarly grow parallel to the shipment units. Such an arrangement conflicts directly with the retrofit and cost requirements and would make reliable operation harder to sustain as the number of tracked units grows. A more defensible arrangement is therefore to keep the shipment side minimal and to concentrate the heavier sensing antennas, communication protocols, and compute and energy burden on the truck-side. One reusable installation could serve many shipment units across repeated tours.

At the abstract level, several design directions could be imagined. However, these options either remain dependent on external or human capture discipline, leave the moving transport leg under-observed, or reintroduce shipment-side cost and recovery constraints. The remaining direction is to distribute responsibilities: keep shipment-side identification lightweight, place reusable and widely compatible observation capability near the transport context, and treat the generated observations as material that needs further interpretation before it becomes operational information.

A truck-based design offers itself due to the truck being the medium in which the parcel traverses current visibility gaps. The moving nature of the truck and the long duration of time any given parcel stays in the truck are advantageous for reporting multiple data points to the backend. The truck can also provide power and presents a place where a fixed installation of hardware is feasible.

From the seven requirements follows a four-part system model that can make this truck-based system a reality. The first part is the **Parcel-Side Identifier Layer**, which identifies parcels or grouped cargo such as handling units and pallets. The second part is the **Truck-Side Tracker Layer**, which generates observations between formal milestones and

relays the captured parcel identities. The third part is the **Backend Interpretation and Derived State Layer**, because vehicle movement alone does not reveal load composition. Raw observations do not yet support planning. It has to relate shipment identity with transport context and other handoff-relevant observations over time, without treating every observation as final truth. The fourth part is the **Dashboard and Bounded Sharing Layer**, because the resulting information must be usable internally and selectively exposable externally.

Visibility in this model progresses from observation to relation to state. Captured observations, therefore, have to preserve time and place and have to deliver information that facilitates relation computation so that state information, such as current status and exception signals, can be derived through the combination of that data. What reaches the user is not an isolated reading, but an interpreted transport and load state.

This allocation extends visibility beyond milestone traces while remaining aligned with the retrofit, cost, reliability, and governance conditions established by the requirement base. It closes the response to **SQ1** by specifying the architecture that makes a third-party handoff visibility artifact operationally plausible.

4.3 Implementation by Layer

The implementation section follows the built prototype system layer by layer, from parcel-side identification to dashboard and sharing surfaces. It provides the main evidence for **SQ2** by showing how field observations become structured event data, how those events are interpreted into derived state, and how the resulting visibility is presented to users through bounded sharing surfaces.

4.3.1 Parcel-Side Identifier Layer

The parcel-side identifier layer provides the signal that enables the linkage between a parcel or handling unit to the truck. Without such a join point, the artifact would only enable vehicle tracking rather than parcel-relevant visibility. The layer is intentionally limited by only supplying identity or presence evidence, while other layers must provide the rest [25, 28, 35].

The first implemented route is an active BLE beacon. The prototype used a *Seeed Studio XIAO ESP32-C3* as shown in Figure 4.1. This exact board choice should not be

understood as a production recommendation. It was useful because it provided BLE capability, is widely documented, has a small board footprint, and low component cost. The component is a 21 mm by 17.5 mm module with 2.4 GHz Wi-Fi and Bluetooth support, and it is readily available for purchase in the low single-digit USD range [76, 77]. That made it preferable to a consumer BLE tag for the experiment because the advertising interval and payload could be changed during prototyping. The BLE variant can be understood as a small powered identifier attached to a parcel or handling unit. Its task is very narrow: it broadcasts a unique BLE advertising payload from which the same objects identity can be recovered [95].

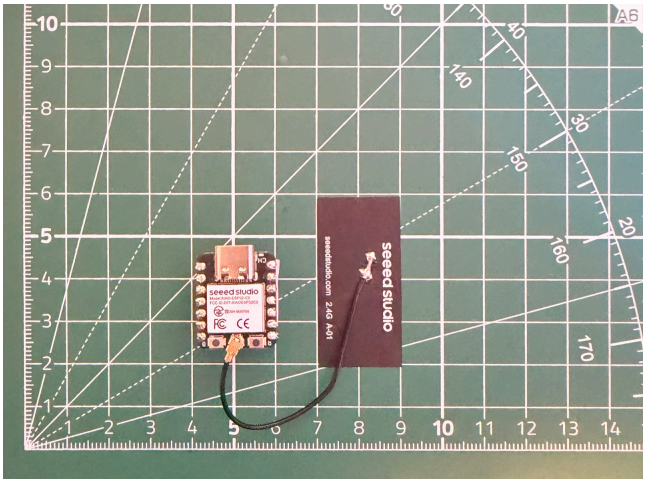

Figure 4.1: Image of the *Seeed Studio XIAO ESP32-C3* with its 2.4Ghz antenna used for the prototype BLE identifier.

Source: Own photograph.

The payload was deliberately kept minimal. The implemented firmware emits a universally unique identifier (UUID)-like beacon token inside the BLE manufacturer-data advertisement. The configured prototype payload only contains a 128-bit beacon token that gets transmitted without two-way connection being necessary. Richer parcel information is not written into the beacon. Keeping the parcel-side payload this small avoids turning every identifier into a data-maintenance object that would have to be updated whenever shipment information changes. The broader shipment metadata can be retained within the backend and programmatically referenced via the identifier at a later stage.

The experience with the BLE build highlighted the importance of simplicity on the parcel side. Both battery life testing using a CR2032 setup and development under Universal Serial Bus Type-C (USB-C) power revealed that, beyond average current consumption,

the most challenging aspects are the power spikes experienced during startup and BLE advertisement. The system uses a duty-cycled approach: the beacon wakes up briefly, sends its advertisement, and then goes back to sleep to minimize its battery draw. Even with this logic, the need for reliable power supply showed already in the development phase that a powered BLE beacon is difficult to justify as the default identifier for all parcels.

The second approach implemented involved passive ultra-high-frequency (UHF) RFID. In this scenario, the parcel-side object is outfitted with a simple passive tag. Both sticker labels and hard card-style tags were tested, with stickers proving to be the more practical choice for attachment to parcels and handling units. The sticker used in the prototype is therefore the passive UHF RAIN RFID sticker label using an *EM Microelectronic em*|*echo-V / EM4425* dual-frequency RFID IC. The chip supports UHF EPC Gen2 / ISO/IEC 18000-63 style operation, while the prototype evaluation focuses only on the UHF read path because that is what was used with the truck-side reader [22].

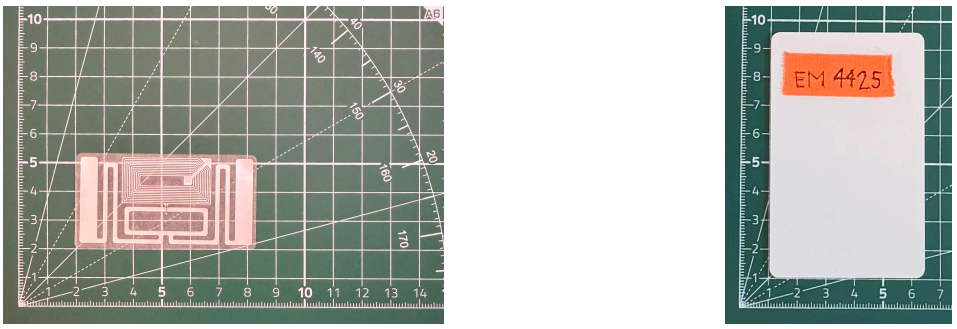

(a) Sticker label **(b)** Hard card-style tag

Figure 4.2: Passive UHF RFID sticker label and hard card-style tag tested as parcel-side identifiers, with the sticker selected as the more practical attachment option.
Source: Own photograph.

Similarly to the BLE identifier, the RFID tag is used only as a source of a unique identifier. The parcel-side element remains cheap, battery-free, and mechanically simple, while the more demanding sensing work is moved away from the individual parcel [11, 42, 45]. RFID is the technology that most closely represents the concept of strictly upholding the minimal-identifier principle [42, 62]. Compared to NFC, UHF RFID offers a more practical read range for the logistics industry [11, 42, 62]. In contrast to UWB, LoRa, or other active radio technologies, it eliminates the need to add yet another powered device to the parcel side [6, 11, 42].

RFID is not free of constraints either. Readability still depends on surrounding material and antenna geometry. The thesis, therefore, does not treat an RFID tag as continuous parcel autonomy. It is a passive object identity that can be detected when the surrounding infrastructure is able to read it [11, 45, 49, 73]. Chapter 5 provides a more in-depth discussion of the tag and reader behavior as part of the technical evaluation.

The prototype also supports software-side linkage as a third association route. This route is different from BLE and RFID because it does not create a new radio observation. Instead, an existing process or upstream system provides the relation between a parcel or handling unit and the truck directly to the backend. In a deployment, this could come from an existing barcode or QR code scan during loading, from a shipper-side loading confirmation, or from another system that already knows which parcels were placed onto which truck. The route was retained because it makes the design more compatible with heterogeneous operations. A firm can use the same visibility backend even when it does not want every association to depend on additional parcel-side hardware. Its limitation is also clear. The association is only as good as the loading information that is supplied, and it would still depend on process discipline, where the relation is created through human scanning. The backend layer below returns to this linkage route when discussing parcel-beacon assignments and grouped handling units.

Consequently, these three routes should not be viewed as mutually exclusive definitive solutions but as configurable evidence and linkage pathways. A deployment can choose an all-BLE approach, an all-RFID setup, or a combination of both. Or even software-side linkage where external partner systems already furnish the necessary loading relationships. Table 4.2 summarizes the resulting decision logic.

The importance of this flexibility became more apparent during the implementation phase than was initially anticipated with the primary focus on BLE. By supporting multiple association routes, the artifact becomes more aligned with the constraints of **R3** and **R4**, as it allows different firms to strike their own balance between cost, reusability, operational discipline, and desired level of visibility.

Another technique for reducing the amount of hardware required on the parcel side is the use of grouped handling units. One identifier does not necessarily need to correspond to a single parcel. Instead, it can represent a pallet, a master box, a roll container, or another stable collection of parcels if the operational workflow does not separate that

Table 4.2: Decision logic for parcel-side association routes, informed by logistics-technology and interoperability literature [11, 35, 45, 66, 95]

Route	Passive parcel side	Automatic observation	Reusable identifier	Process-data reliance	Best-fit decision logic
BLE beacon	–	✓	✓	Low	Use when a powered, reusable identifier is acceptable and repeated proximity evidence is useful, especially for handling units or higher-value flows.
UHF RFID tag	✓	✓	Partial	Low	Use when parcel-side cost and battery-free operation dominate, while accepting that reads require truck-side antenna placement and backend smoothing.
Software-side linkage	✓	–	✓	High	Use when a shipper or upstream system already knows the loading relation and can create, verify, and correct the association as operations change.

union delivery. This is important because the cost of visibility is not determined solely by the price of individual tags or beacons but by the number of identifiers required to track a complete shipment flow [3, 44, 65].

The main output of this layer is the `beaconId` join key. It allows the backend to link two physical items, the parcel and the truck it is in, with another.

4.3.2 Truck-Side Tracker Layer

The truck-side tracker is the reusable physical visibility node. Mounted in the vehicle, it observes parcel-side identifiers, captures truck location and runtime status, and sends normalized events to the backend. This reduces transport blindness between scan points while keeping sensing and transmission in the vehicle instead of across hundreds of parcel-side tags [23, 28, 34].

The prototype uses a *Raspberry Pi 5* as a prototyping host because it offered interface flexibility, local storage, BLE support, USB expansion, Hardware Attached on Top (HAT) compatibility, and enough compute headroom for rapid iteration [58]. The relevant design point is the allocation of powered, connected capability to the vehicle side itself.

This tracker functionality differentiates the device from a basic GPS tracker. If the device would have only used a GPS unit, it would only know the truck's current latitude and longitude. As such, the tracker was made as a sensing node and a cellular uplink device

all in one. To fulfill global navigation satellite system (GNSS) positioning and the LTE cellular link, the *SIM7670G LTE Cat-1/GNSS HAT* was used [89].

Meter-level truck positioning is sufficient for the target visibility problem, because the artifact needs to locate a truck-level transport context rather than perform centimeter-level indoor localization. That said, GNSS antenna placement remains a practical installation issue. In practical terms, an inside-cabin or cargo-area placement is viable only where it preserves sufficient sky visibility; windows, roof structures, and other obstructions can reduce the GNSS signal, so some vehicle setups require a more exposed placement with direct line of sight to the sky [83].

Cellular communication was chosen because the transport leg cannot rely on site Wi-Fi, controlled gateway coverage, partner-owned infrastructure, or uniform OEM telematics access. This setup keeps the prototype self-contained from a retrofit perspective while still requiring practical mounting, power, and radio placement [12, 32, 69].

BLE scanning uses the Pi's onboard capability and therefore adds little hardware burden to the tracker [58]. RFID support was implemented with a *YANZEO SR791 UHF RFID Reader*, an integrated long-range UHF reader with a 10 dBi antenna, a USB interface, 860–960 MHz operation, and UHF EPC Gen2 support [97]. This path becomes more installation-sensitive because the reader has to be positioned inside the cargo area where parcel-side tags can be detected [30, 62]. RFID support is optional to the specific deployment. The implementation, therefore, supports configurability at both ends of the association chain: the parcel side can remain minimal, while the truck-side node can support the desired technologies variably.

Figure 4.3: Truck-side tracker prototype with Raspberry Pi 5, the SIM7670G LTE/GNSS HAT, the RFID reader antenna, the GNSS antenna, and the LTE antenna attached.
Source: Own photograph.

A separate tracker battery was considered but rejected in favor of vehicle power. Providing a continuous GNSS location feed with Cellular uplinks simply becomes easier when it can rely on the truck for power, rather than depending on a battery that has to be charged or replaced. The current prototype is a USB-C device powered from a standard 12 V vehicle accessory socket via an adapter. In a brownfield installation, such an accessory-socket power path can itself be a realistic long-term retrofit option when cabling and driver ergonomics are acceptable. A more permanent 12 V vehicle-power connection converted to the required device voltage is a second option for installations where concealed cabling or operational robustness matter more.

The tracker has to gather location evidence, observe parcel-side identifiers, maintain its own runtime health, and deliver events under imperfect connectivity. In the implemented prototype, these responsibilities are split across runtime services such as the location uplink, BLE uplink, RFID uplink, and resource monitoring scripts. In implementation, it became apparent that these divisions were important so the device could handle reboots and system errors without human intervention. System services were needed so that the device could restart its sensing and upload paths reliably after erroring out during field-like use.

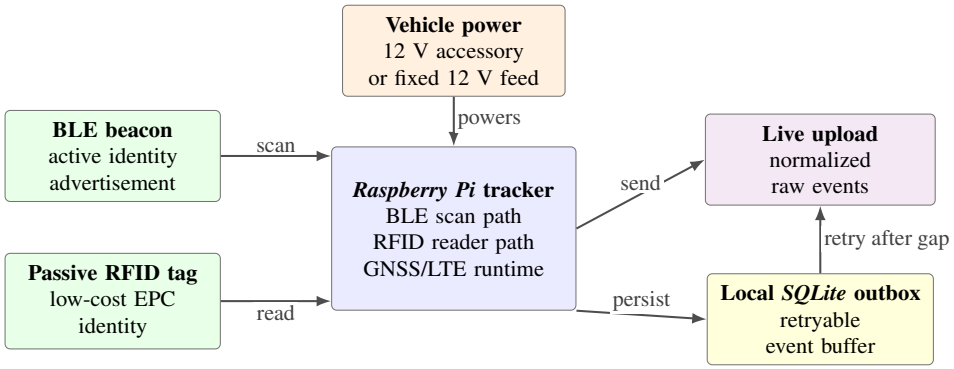

Figure 4.4: Truck-side hardware setup and the main sensing and connectivity paths.
Source: Original representation.

The runtime can be understood as a short pipeline across three observation families. For BLE, the input is a manufacturer-data advertisement; for RFID, it is a raw reader frame from the UHF reader; for location, it is a modem response to GNSS AT commands. In each case, the tracker performs only a bounded transformation: it extracts the relevant identifier or location state, wraps the result in the shared raw-event envelope, and queues the event before upload. The code listings below (Listings 1 and 2) are selective implementation evidence for this pipeline. They show the input-to-output boundary for BLE, RFID, and location events, then show how queued events survive connectivity gaps.

There is another difference in location tracking that only became apparent during testing and provides direct support for **R6**. Even when the tracker is awake and processing, it may not have a fix available from GNSS at a given time. For that reason, heartbeat messages and *no-fix* events are also preserved in the device's run time. This provides useful information to the backend so that "device online but no fix" can easily be distinguished from "device offline". The interval settings used in the prototype, including a 30-second location and heartbeat cycle, should be understood as prototype settings rather than an optimized production recommendation. They were frequent enough to exercise the end-to-end chain and reveal reliability issues, but the final reporting interval would depend on operational value, data cost, and power constraints in a deployed system.

The most important implementation refinement was local buffering. During field-like use, mobile connectivity gaps were not hypothetical. If the tracker had only transmitted live events, the loss of connectivity would have created exactly the kind of blind interval the artifact is supposed to reduce. The tracker, therefore, uses local *SQLite* outboxes and retry logic so that events can survive temporary uplink failures. Listing 2 shows how a

Listing 1 Tracker-side normalization of BLE, RFID, and GNSS observations into raw-event envelopes.

```
# BLE input: advertisement_data.manufacturer_data
identity = parse_beacon_identity(advertisement_data.manufacturer_data)
event_seq += 1
payload = {
    "eventId": new_event_id(),
    "contractVersion": CONTRACT_VERSION,
    "trackerId": TRACKER_ID,
    "eventType": "ble_scan",
    "trackerTsMs": ts_ms,
    "seq": event_seq,
    "beaconId": identity.beacon_id,   # e.g. ble:<uuid>
    "signal": {"rssi": rssi} if rssi is not None else {},
    "sourcePayload": identity.source_payload,  # for debug
}

# RFID input: raw 64-byte reader frame
epc = extract_epc(frame)
event_seq += 1
payload = {
    "eventId": new_event_id(),
    "contractVersion": CONTRACT_VERSION,
    "trackerId": TRACKER_ID,
    "eventType": "rfid_scan",
    "trackerTsMs": ts_ms,
    "seq": event_seq,
    "beaconId": f"rfid:{epc}",
    "sourcePayload": {"epc": epc, "source": "rfid"},   # for debug
}

# GNSS input: SIM7670 AT-command response
gnss_location, cgps_raw = parse_cgpsinfo(at_cmd(ser, "AT+CGPSINFO"))
event = build_event(
    conn,
    "gps_fix",
    location=clean_location,
    source_payload={"cgpsinfoRaw": cgps_raw, "cgnssinfoRaw": cgnss_raw} # for debug
)
```

backoff delay postpones the next attempt after a failure that prevents the tracker from repeatedly querying the backend during poor connectivity. A dead-letter path is used for permanent client-side failures, such as malformed requests, so that one unrecoverable event does not block the rest of the queue.

Listing 2 Focused excerpt of durable truck-side outbox handling for retries and dead-lettered events.

```
rows = conn.execute(    # get items in outbox
    """
    SELECT id, event_id, payload, attempts
    FROM outbox
    WHERE next_attempt_ms <= ?
    ORDER BY id ASC
    LIMIT ?
    """,
    (now_ms(), MAX_FLUSH_BATCH),
).fetchall()

for row_id, event_id, payload, attempts in rows:
    status, body = post_payload(payload)
    if 200 <= status < 300:     # success case: remove from outbox
        conn.execute("DELETE FROM outbox WHERE id = ?", (row_id,))
    elif is_permanent_http_failure(status): # permanent fail case: store for debug
        move_to_dead_letter(conn, row_id=row_id, event_id=event_id,
                            payload=payload, status_code=status,
                            error_text=body)
    else:   # retry case: retry in ? ms
        next_attempt = now_ms() + retry_backoff_ms(attempts + 1)
        conn.execute(
            """
            UPDATE outbox
            SET attempts = ?, next_attempt_ms = ?
            WHERE id = ?
            """,
            (attempts + 1, next_attempt, row_id),
        )
```

The tracker layer intentionally returns raw normalized events. Hence, the tracker's output focuses on what it observed and when it observed it. It is the backend's role to determine what those recorded observations signify about the operational context [38, 41, 46].

4.3.3 Backend Interpretation and Derived State Layer

The third layer combines version adherence with contract, keeps a raw history, computes association, and derives state. These responsibilities are discussed together because they form one continuous transformation. The tracker layer ends by producing normalized observations from the field. The backend layer begins when those observations have to become interpretable enough to support operational decisions. This makes the layer central to **SQ2** because it explains how event and state data is structured and how noisy observations are interpreted.

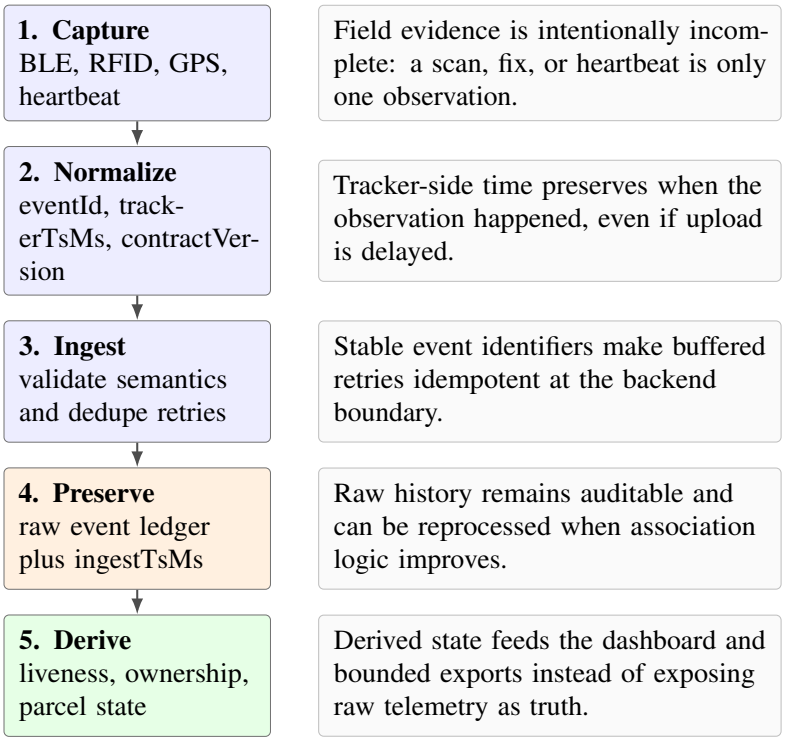

Figure 4.5: Event flow from truck-side capture through contract validation and raw storage to derived read-side state.

Source: Original representation.

This layer therefore implements **R5** and **R6** by turning raw events into reliable, queryable state. A raw GNSS fix is useful evidence, but it is not yet parcel visibility. A BLE or

RFID observation is useful evidence, but it is not yet proof that a parcel is loaded on a truck.

The first design decision in this flow is the shared event contract. The contract ensures that different data sources send the data in such a format that the backend supports. While working on the prototype, it became apparent that ingested data can change with identifier iteration or tracker firmware. To make it easy for the backend to know how to interpret what data. A contract ensuring that the data format received is the data format expected was added as a first check before the data gets entered into the database. Listing 3 shows how the code acts as a guard for incoming events.

Listing 3 Focused excerpt of the backend raw-event contract and semantic validation.

```
// general checks
export const ingestRawEvent = mutation({
  args: {
    eventId: v.string(), // required fields
    contractVersion: v.string(),
    trackerId: v.string(),
    eventType: eventTypeValidator,
    trackerTsMs: v.number(),
    seq: v.number(),
    beaconId: v.optional(v.string()), // optional fields
    location: v.optional(locationValidator),
    signal: v.optional(v.object({ rssi: v.optional(v.number()) })),
    status: v.optional(v.object({ locationMode: v.optional(locationModeValidator) })),
    sourcePayload: v.optional(v.any()), // optional field without any verification
  },
  handler: async (ctx, args) => {
    assertCompatibleContractVersion(args.contractVersion);
    assertEventSemantics(args);
  },
});

// specific checks (if this, then that)
function assertEventSemantics(args: RawEventArgs) {
  if ((args.eventType === "ble_scan" || args.eventType === "rfid_scan") && !args.beaconId) {
    throw new Error(`${args.eventType} requires beaconId.`);
  }
  if (args.eventType === "gps_fix" && args.location?.method !== "gnss") {
    throw new Error("gps_fix requires location.method='gnss'.");
  }
}
```

The second decision is to keep tracker-side time separate from backend ingestion time. Tracker-side time records when the field observation occurred. Backend ingest time

records when that observation reached the backend. In road transport, these two moments can diverge because of mobile connectivity gaps and buffered uploads. The backend, therefore, stores both and updates tracker liveness only when the incoming event is appropriate for the latest tracker-time snapshot. This has the practical effect of preventing the replay problem, whereby an old observation transmitted late makes the associated truck or parcel appear older than it actually is. Delayed uploads thus can enrich the history without contaminating the operational live view.

Deduplication is the third choice in this same sequence toward data reliability and directly relates to **R6**. Because a deliberate decision in the tracker software was to always try retransmitting event data from its outbox after failure, it is critical that the `eventId` is treated as an idempotency key on the backend. If the backend has already received an event with this ID, those retries receive a success response (to stop further retries), but the event does not get stored a second time in the backend.

Listing 4 Focused excerpt of retry-safe backend deduplication and raw-event preservation.

```
// has this eventId already been seen?
const existing = await ctx.db
  .query("ingest_dedupe")
  .withIndex("by_event_id", (q) => q.eq("eventId", args.eventId))
  .first();

if (existing) { // duplicate does not get inserted again into raw events
  return { ok: true, duplicate: true, rawEventId: existing.rawEventId };
}

// new event gets inserted into raw events
const ingestTsMs = Date.now();
const rawEventId = await ctx.db.insert("raw_events", { ...args, ingestTsMs });

// eventId gets stored in dedupe database
await ctx.db.insert("ingest_dedupe", {
  eventId: args.eventId,
  trackerId: args.trackerId,
  rawEventId,
  createdTsMs: ingestTsMs,
});
```

The storage of accepted events as raw history is treated as separate from the later derivation of state. Raw events are essentially the individual pieces of evidence observed in the field by the tracker. Derived state is the operational interpretation placed on top of those raw data points. By keeping the raw event history, a deployment can do further analysis on the

same data after the fact and can reliably trace back association and state information to specific events, thereby catching repeated errors and being able to ameliorate the system over time [2, 41, 73].

Association is where that preserved raw evidence is put into the context of the parcel. This is the layer that can make or break the end-to-end prototype architecture. The previous layers successfully collect data on a variety of things, but without the association of this data to one another, little information can be pulled from the raw ledger. Because the workload was previously separated from the identifier to a mix of identifier, tracker, and backend. The backend now has the task to recombine this data.

The first association task is to compute, with the available evidence, what parcel is on what truck, because identifiers are linked to parcels or handling units already. It is now the task of the backend to compute what identifier is located on which truck. The sensing information from the trucks' trackers supplies a good evidence base. However, identifier reads can be ambiguous, as for example when a BLE beacon gets read by two trucks that are parked next to each other [13, 71, 95].

In such a scenario, the state would intentionally be set as ambiguous until the beacon is only seen by a single truck. If this beacon was previously also associated with the truck, the ambiguous gap gets backfilled with association records for this truck. The opposite case is where this change from one truck to another actually takes place. A beacon gets detected by Truck A, later gets detected by both Truck A and B, which results in an ambiguous state, and lastly, only gets seen by Truck B. In such a scenario, the system would keep the ambiguous state during the transfer time. After it is no longer detected by Truck A, the state changes to show that this beacon and therefore also the related parcels are now on Truck B.

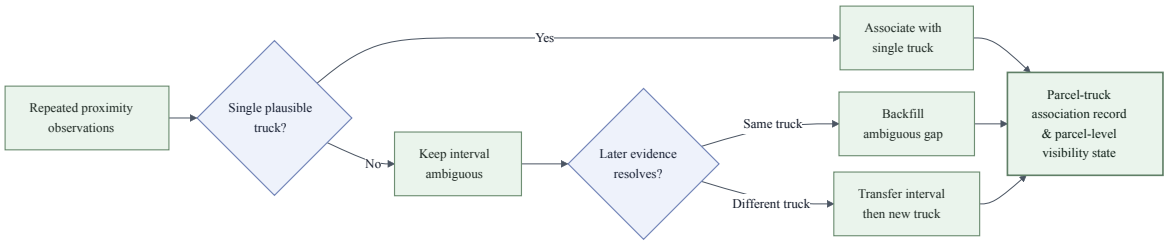

Figure 4.6: BLE-based association path from ambiguous proximity evidence to a parcel-level visibility state.

Source: Original representation.

RFID, on the other hand, provides a different type of association evidence. Because RFID can only be sensed by the tracker when the tag is getting powered by the RFID antenna, cases exist where an RFID identifier only gets scanned a couple of times by the tracker [42, 62]. This tracker can therefore only deliver limited association evidence to the backend. An identifier is therefore treated as on the truck as soon as it gets detected by the tracker, and only gets disassociated from the tracker if the parcel or the identifier is detected by another tracker. A more advanced association mechanism for RFID could include vehicle speed and stop times to more accurately compute the start and end times of an RFID identifier's presence in a truck.

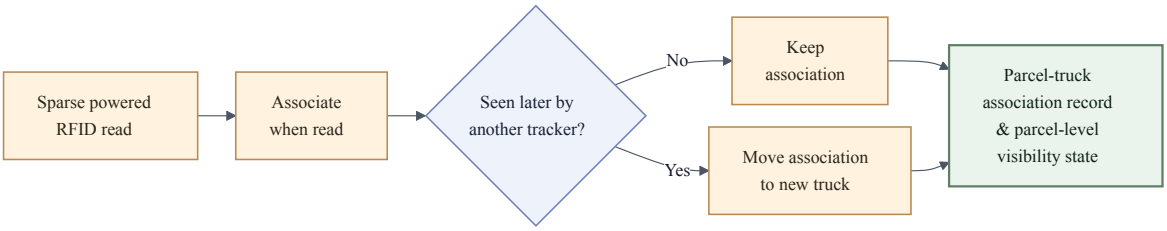

Figure 4.7: RFID-based association path from a sparse powered read to a parcel-level visibility state.

Source: Original representation.

The backend can also support system-side parcel-to-tracker association, where the logistics companies send their own association records to the backend. This association mechanism could either be taken at face value or, in the future, strengthen existing association evidence that is identifier-generated.

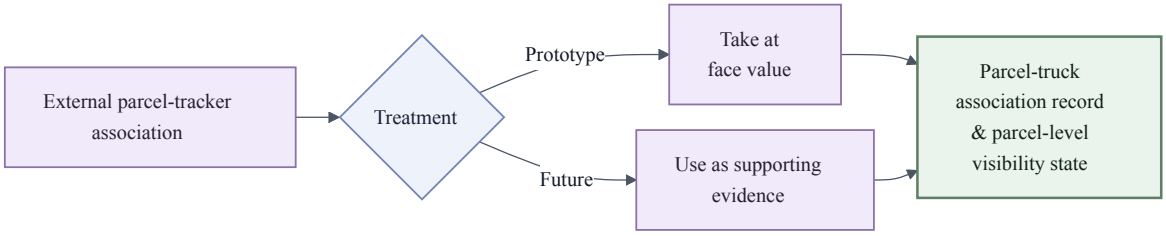

Figure 4.8: System-side association path from an external parcel-tracker record to a parcel-level visibility state.

Source: Original representation.

All of the association mechanisms discussed above should still be read as prototype-level proof only. They do not show a definitive, solved algorithm, as they would have to

determine optimal thresholds and prove universal RF robustness first. Then they would have to be further tested under dense yard conditions [42, 62, 95].

When association is finally determined, the second step of the association layer is to combine data captured from the tracker with information associated to the identifier. We now know where the truck is. We know what Identifiers are on that truck. We know what parcels are linked to those identifiers. Depending on supplier-provided information, we even know what items are in the individual parcels. This is where truck-level GNSS becomes parcel-relevant visibility. The truck remains the physical source of location evidence, but the backend relation makes that location meaningful for the parcel [25, 66, 82].

Table 4.3: Backend-derived outputs served to the dashboard.[2, 41, 73]

Derived output	Input trigger or basis	Operational purpose
Tracker route	GNSS fixes; cellular fallback; tracker time window	Show truck movement as an inspectable route before parcel-specific association is applied.
Tracker liveness	Latest accepted event; heartbeat; no-fix record	Distinguish an online tracker, a stale tracker, and an online tracker without a current location fix.
Beacon owner state	BLE/RFID scans; heartbeat; staleness; prior association state	Indicate which tracker most likely carries the identifier, including confidence, ambiguity, and conflict cues.
Beacon route path	Owner segments; tracker GNSS/cellular positions	Translate truck-level movement into the route history of the identifier or linked parcel.
Parcel live location	Active parcel–identifier assignment; beacon owner state; tracker location; liveness	Show where the parcel is likely located and whether this state is fresh, stale, or unavailable.
Parcel explorer	Parcel record; manifest items; active assignment; assignment history; route points; live state	Join transport evidence with the business object that receiving and operations users actually search for.
Sharing export	Selected path, scan, or assignment fields; chosen time window	Provide bounded partner-facing visibility without exposing complete internal telemetry.

The derived state outputs are summarized in Table 4.3. The table starts from the dashboard-facing outputs because several input events and association records can be combined into the same operational answer. In the flow of the artifact, this is the final backend step before information reaches the dashboard: preserved evidence is turned into read-side state that can answer operational questions.

4.3.4 Dashboard and Bounded Sharing Layer

The fourth and last implementation layer is the dashboard and bounded sharing layer. Its purpose is to display the state that was computed in the backend to the end users so that they can profit from this operational visibility. State information alone does not answer the questions that operators actually have. They need to be able to see the parcels or

trucks' locations on a map. Recipients need a way to access load information and to inspect transport progress. External partners need bounded status information without full access to the telemetry base. The layer therefore mainly addresses **R5** and **R7**, while also supporting **R2**, because arrival and load information becomes useful only if it can be interpreted and acted upon by the relevant parties [15, 82, 85].

The dashboard was additionally helpful to identify errors or to see direct information during the development and testing of the prototype. For stakeholder evaluation, the dashboard offered a user-friendly way to show the data's ability to generate familiar representations of that information. Chapter 5 returns to these surfaces when evaluating the operator-facing artifact and the stakeholder response to the concept.

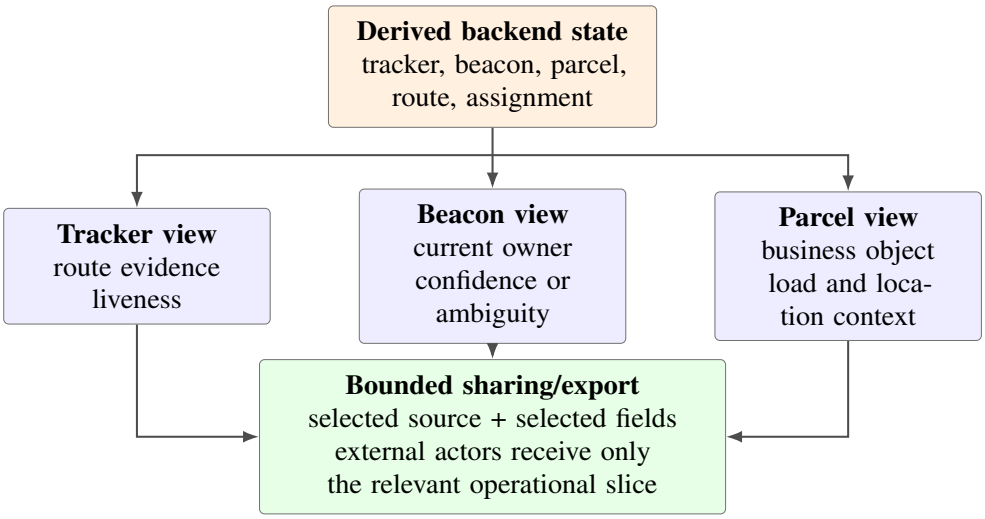

Figure 4.9: Tracker, beacon, and parcel dashboard surfaces and the distinct question each one answers.

Source: Original representation.

The parcel view is perhaps the most operationally relevant view, as it shows the location and state of the business object that operators and receiving actors ultimately care about. It displays the parcel on a map, showing its trajectory, and reveals the item manifest of items that are within that parcel.

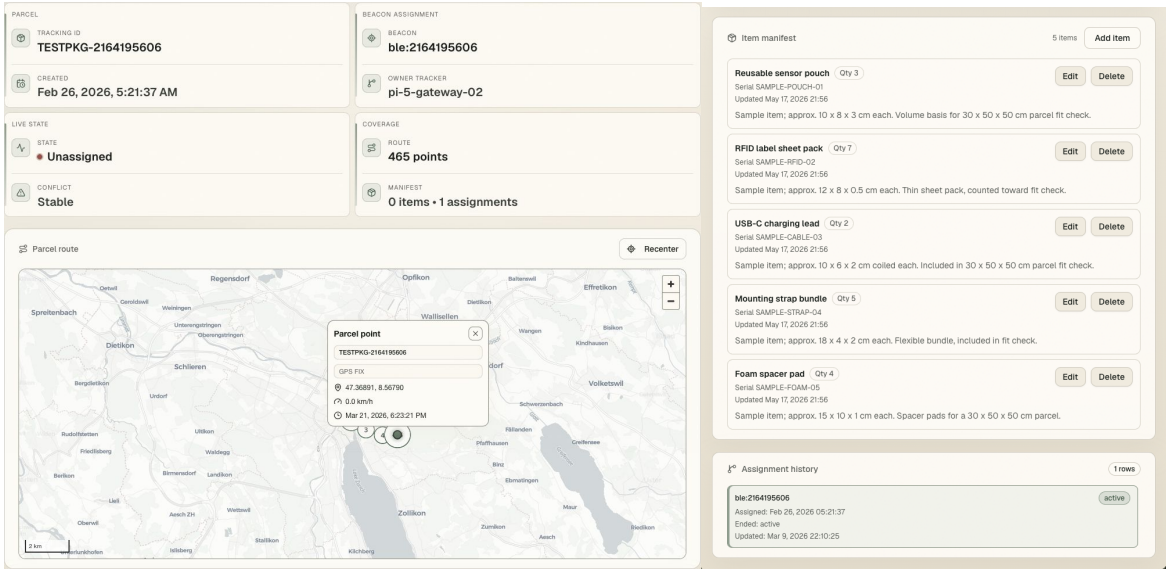

(a) Parcel overview and route map **(b)** Manifest and assignment history

Figure 4.10: *Parcel explorer* panels for the same test package, showing parcel metadata and route evidence alongside the stored item manifest and assignment history.
Source: Own Screenshot.

The beacon view is one abstraction layer above the parcel. It reveals the associated parcels that are linked to said beacon. It shows what truck it is on and displays information about the certainty and current status of the identifier-to-tracker association that is currently in effect. This view serves as an interface to diagnose parcel-to-truck associations by showing the user the abstract layer that sits between the tracker and the parcel.

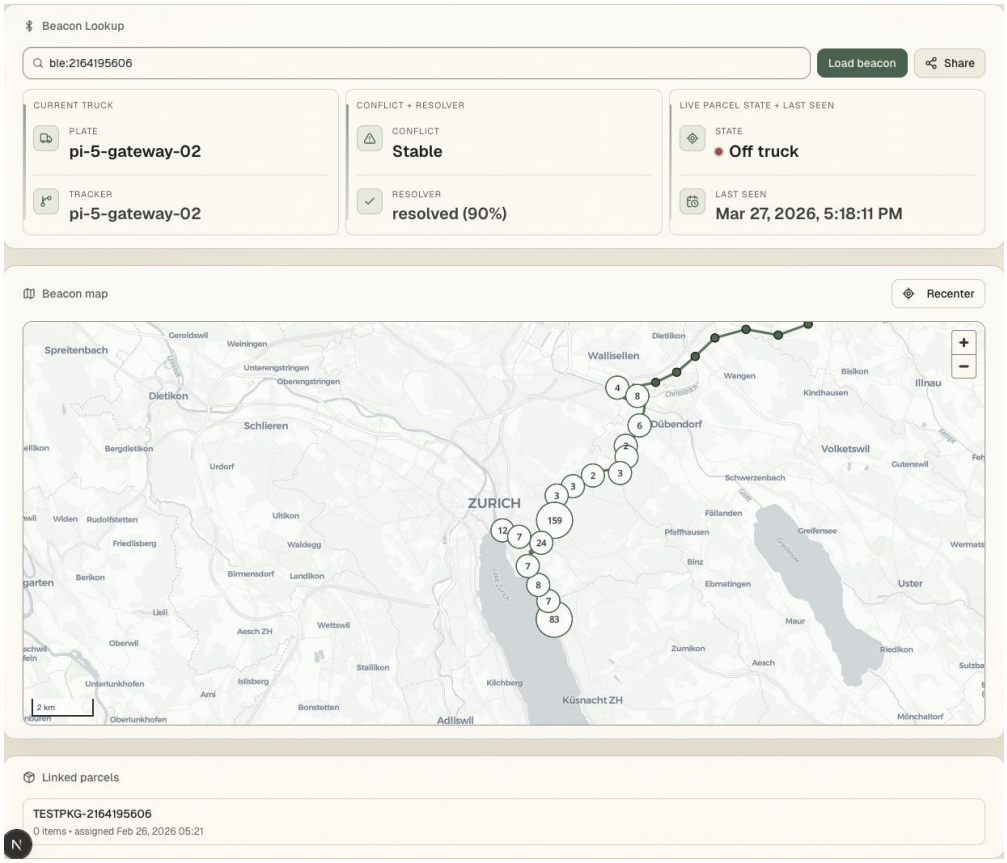

Figure 4.11: *Beacon explorer* with live-backend resolver state, map evidence, and linked parcel context.

Source: Own Screenshot.

The tracker view enables the observation of a truck's movements detached from any parcel association. It displays the map with the route of the truck, vehicle speed over time, and shows a running log of incoming sensing events for the inspected tracker. The individual points of the tracker's route on the map can be inspected further, whilst clicking on a sensing event routes the user to the beacon view of the appropriate beacon.

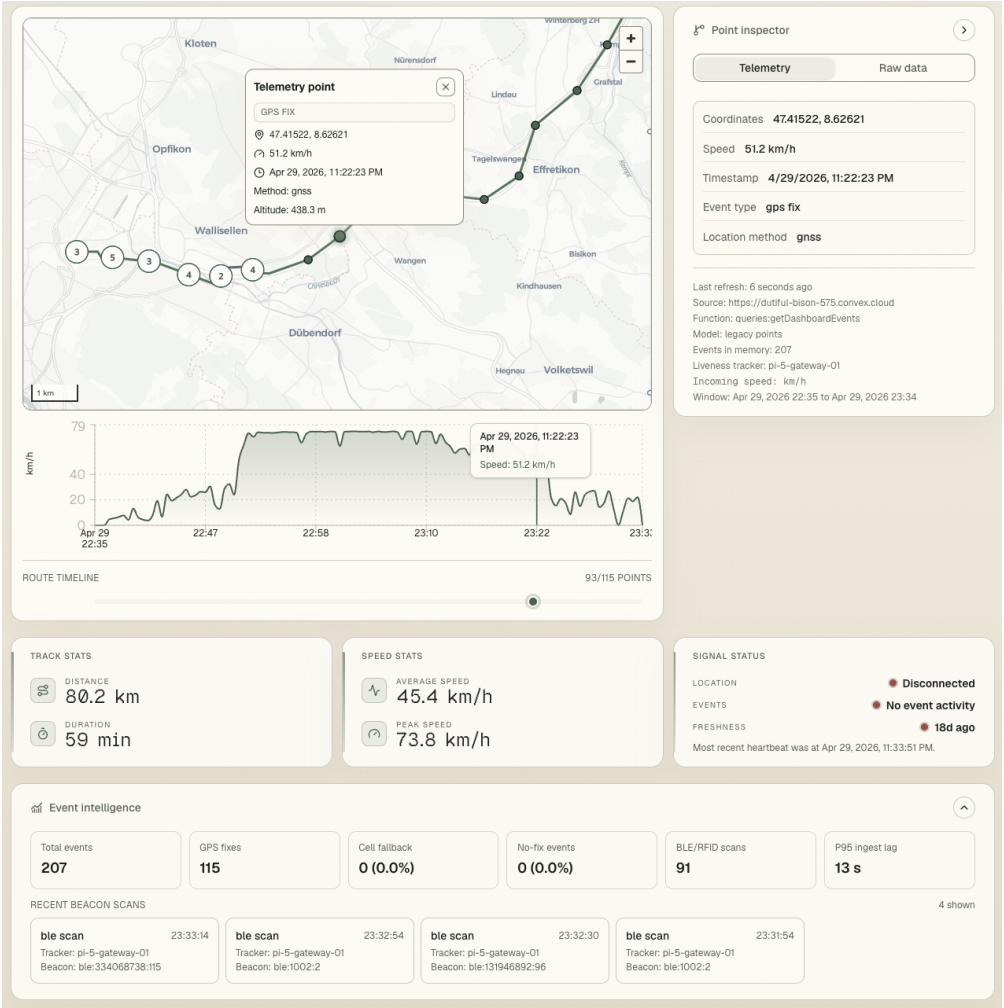

Figure 4.12: *Tracker explorer* showing route telemetry, point inspection, speed statistics, signal status, and event intelligence for an inspected tracker window.
Source: Own Screenshot.

Not all of those layers are adequate for every user type because the data density and access level of the beacon view, for example, can be too much for an internal phone operator. For data sharing across organizational boundaries, an export functionality has been introduced on a proof-of-concept level. Different presets of restricted views can be exported to a comma-separated values (CSV) document in the current implementation.

This bounded export functionality showcases the possibility for further data sharing opportunities through APIs and information sharing systems [15, 34, 85].

The layer prepares the **SQ3** discussion by showing that derived visibility can be packaged into controlled information products. Strong operational benefits, however, would likely require this visibility data to connect to downstream systems and processes that can trigger, or reinforce concrete decisions [98].

4.4 Implementation Scope, Trade-Offs, and Deployment Considerations

The implemented core is sufficient for the Chapter 5 evaluation because it covers truck capture, BLE/RFID identifier sensing, local persistence during uplink outages, backend ingest, derived tracker, beacon, and parcel state, dashboard views, and bounded sharing. It does not aim to present itself as a final product that is ready for full-scale adoption. The prototype was built as described to enable the following technical evaluation, requirement-oriented assessment, stakeholder reaction, and bounded business-feasibility reasoning.

Several trade-offs shaped the core design. The first being cost versus granularity. In response to **R4**, the concept avoids parcel-level GPS for every shipment unit and instead accepts association logic as the mechanism that recovers parcel relevance from cheaper identifiers. The second is power versus autonomy. In response to **R3**, the reusable intelligence sits on the truck and draws on vehicle power, which improves runtime stability but assumes viable truck-side installation for the artifact to function. The third is bounded sharing versus full transparency. In response to **R7**, the system aims to share derived and useful visibility slices, not unrestricted internal telemetry. The system's flexibility comes with a degree of architectural complication. Having support for BLE, RFID, buffered data ingest, along with multiple read models increases retrofitting ease and interpretability based on **R3**, **R5**, and **R6**. The complexity, however, is far broader than a single device solution design would have been. These compromises are sound, given the goal to enable evaluation as a bounded prototype.

The chapter, therefore, closes with the current implementation boundary rather than with a claim of completeness. It has established the requirement logic and architecture needed for **SQ1** and shown for **SQ2** how event and state data are captured, structured, interpreted,

and presented as operational visibility. Chapter 5 now assesses how far this bounded implementation supports the thesis-level feasibility claim.

Table 4.4: Implementation boundary of the current prototype

Capability or layer	Status	What exists now	Why the boundary remains
Truck GNSS capture and heartbeats (**R1/R6**)	Implemented	Continuous truck-side location, no-fix, and liveness events through the Pi runtime and backend ingest.	This is part of the core feasibility claim and therefore had to be implemented directly.
BLE parcel association path (**R1/R2**)	Implemented	BLE payload parsing, beacon identity derivation, event uplink, and backend ownership resolution.	Core parcel-association logic had to exist to make the truck-first concept meaningful.
RFID parcel association path (**R1/R2/R4**)	Implemented	EPC extraction and normalized `rfid_scan` events entering the same ownership engine.	Preserves a lower-cost passive association path without requiring a second backend logic stack.
Buffered delayed upload and replay-safe ingest (**R6**)	Implemented	SQLite outboxes, retries, dedupe by `eventId`, immutable raw history, and monotonic liveness handling.	The target field environment makes blackout tolerance a core requirement rather than a nice-to-have.
Derived beacon ownership and parcel live state (**R5/R6**)	Implemented / evolving	Live state, route sessions, conflicts, parcel assignments, and parcel explorer views exist today.	Resolver details can still mature without invalidating the architectural principle.
Grouped handling-unit logic (**R2/R4**)	Partly implemented	One identifier can already stand for a grouped unit in the current object model.	Credibility under split, repack, or partial-offload scenarios still needs stronger field validation.
Selective sharing and export (**R5/R7**)	Partly implemented	Bounded beacon-share preview and export surfaces exist.	Full stakeholder-specific enterprise integration was outside the scope of the prototype.
Partner-facing system integration (**R2/R7**)	Conceptual	No full ERP / TMS or customer platform rollout was built inside the thesis scope.	The thesis prioritizes demonstrating the bounded visibility concept over finishing every downstream workflow.

Chapter 5

Artifact Evaluation and Discussion

5.1 Evaluation Setup

The evaluation follows the design-science expectation that an artifact should be assessed against the problem environment and the objectives for which it was built [36, 68, 70]. In this thesis, the artifact is evaluated as a bounded prototype to address the research questions and the requirements derived in Chapter 4.

The evaluation has three strands. First, the **Technical Prototype Evaluation** section evaluates whether the implemented tracker and its surrounding architecture can support the visibility concept in practice. This strand mainly addresses **SQ2** and the technical aspects of the derived requirements. Second, the **stakeholder feedback section** evaluates whether practitioners recognize the problem-solution fit and how they qualify the artifact's operational value. Third, the **business-feasibility section** evaluates **SQ3** by asking whether the visibility state can plausibly create internal or external business value. The strands are then interpreted together because none of them alone can prove deployment feasibility.

5.2 Technical Prototype Evaluation

The technical evaluation asks whether the prototype can support the main visibility chain from identifier capture to dashboard display. The tests, therefore, ask whether a physical observation can reliably become displayable state information without information getting lost along the chain.

The prototype visibility system evaluated consists of a Raspberry Pi-based truck tracker, BLE and RFID identifiers, the Convex backend, the code powering the Dashboard, and the firmware of the different components. The tracker-side tests used two (identical in hardware) Raspberry Pi gateways during the evaluation period, named in the raw data as `pi-5-gateway-01` and `pi-5-gateway-02`. The technical evidence was collected from home-lab tests, road tests, backend upload tests, dashboard observations, and logging information from the tracker.

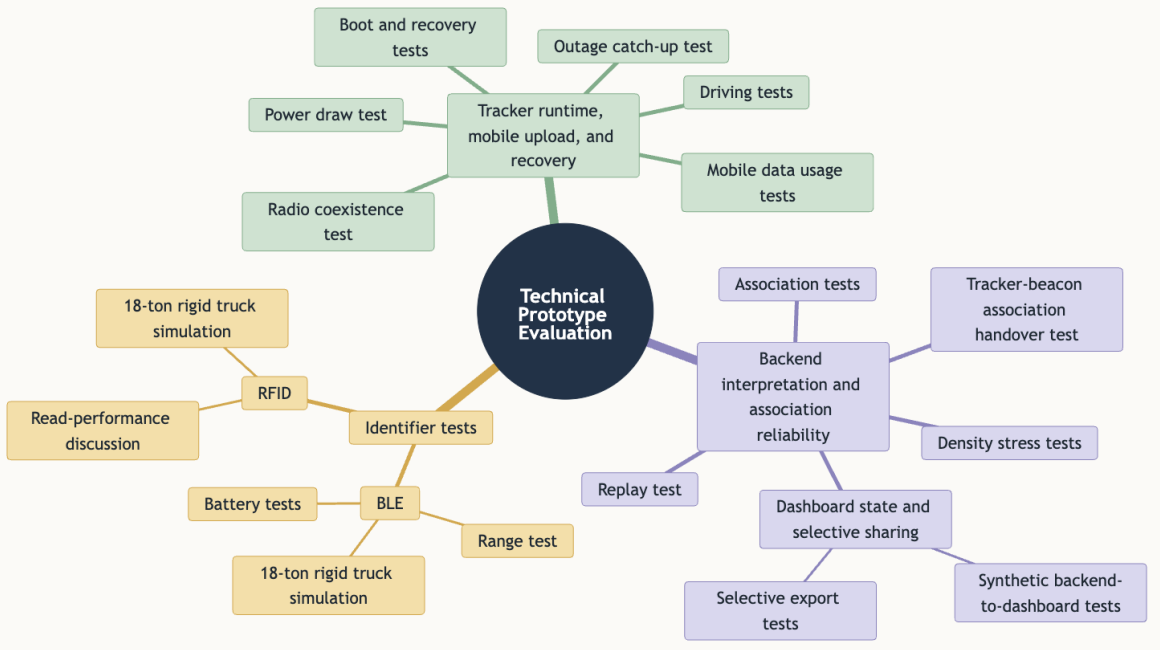

Figure 5.1: Tree diagram of the tests and evidence blocks used in the technical prototype evaluation.

Source: Original representation.

5.2.1 BLE and RFID Sensing Evidence

The sensing tests evaluate the ability for the identifiers to be sensed by the tracker. To produce desirable results, both components need to function with each other.

In the **outdoor BLE range test**, the tracker detected the BLE beacon in all tested distance windows up to 30 m. The result supports **R1** and **R6** because the beacon can create repeated presence evidence beyond formal scan points. At the same time, the received signal strength indicator (RSSI) values were noisy and environment-dependent. This matches prior BLE literature, where RSSI-based proximity is known to vary with orientation and human or material obstruction [13, 61, 71].

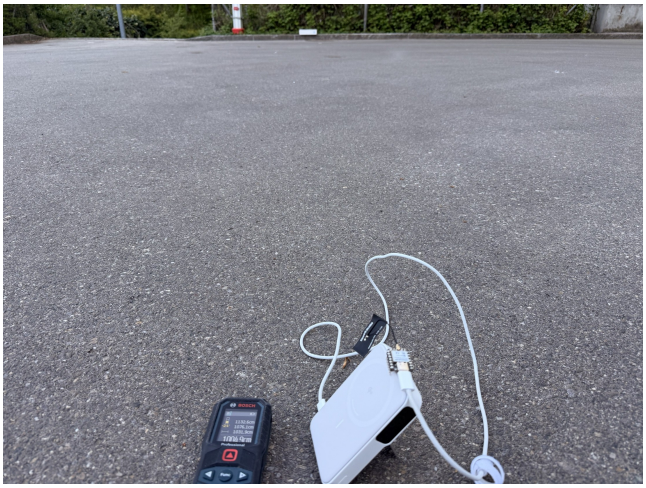

Figure 5.2: Outdoor BLE range-test setup with the BLE beacon connected to a battery pack and the distance meter in the foreground, and the `pi-5-gateway-01` on top of a cardboard box visible in the distance.

Source: Own Photograph.

This high range for the BLE identifier is not a uniquely positive outcome for the underlying scenario. Whilst it decreases the chances of a parcel being on a truck and not being detected by the tracker, it also increases the chances that a beacon that is nearby the truck falsely gets associated with the tracker. To address this, the backend association layer already handles ambiguous state and only associates a parcel to a truck if the linkage is unambiguous.

Table 5.1: BLE RSSI observation by distance. (Lower RSSI is better)

Distance	10 cm	1 m	3 m	6 m	10 m	15 m	20–25 m	30 m
Mean RSSI (dBm)	-31.0	-60.5	-56.3	-59.3	-67.6	-69.2	-73.9	-77.2

The RFID tests produced a narrower result. In the orientation test, RFID reads were strong in favorable geometry but weakened sharply outside the angles of the RFID antenna sensing cone.

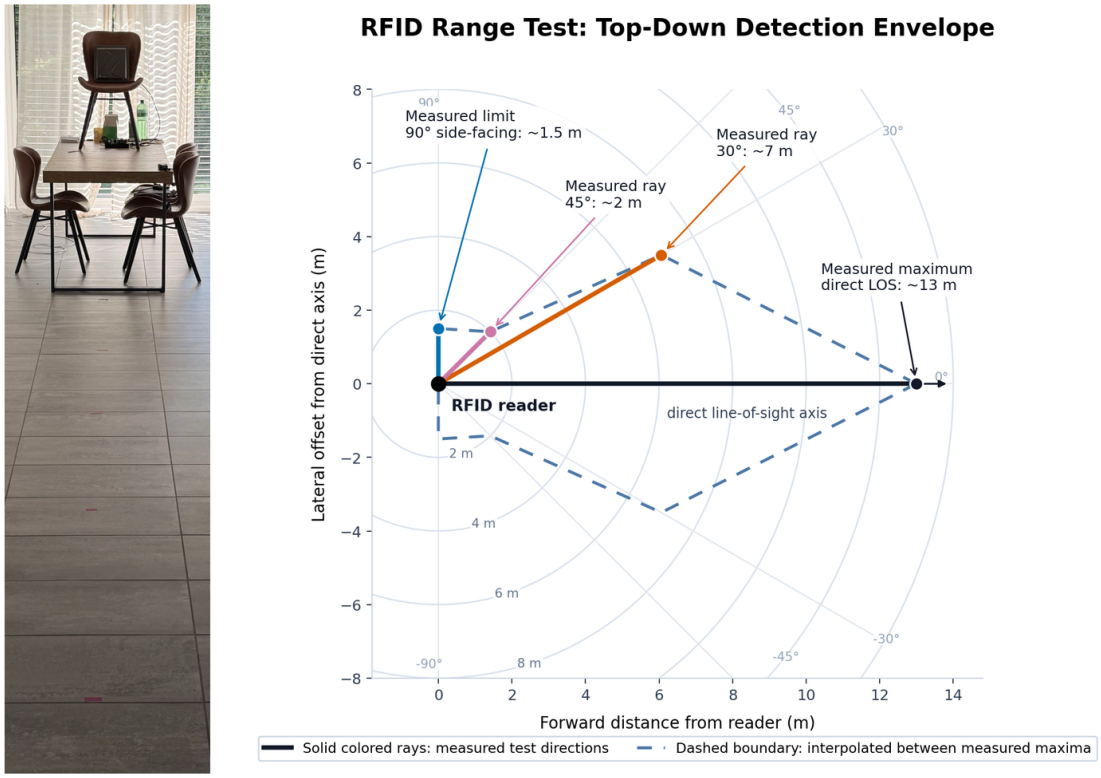

Figure 5.3: RFID angle-test setup and measured top-down detection envelope for the reader/tag configuration.

Source: Original representation.

To test the tracker's capability of sensing the RFID tags at a variety of angles relative to the antenna orientation, the tag got moved further and further away from the antenna at different angles. At a 90° angle to the antenna, the captures stopped coming in at a distance greater than 1.5 m. The range barely increased at an angle of 45° relative to the antenna, whilst at 30° the tag was able to get sensed by the tracker up to a distance of 7 m. The maximum range was 13 m, and it was achieved only when placing the tag at 0°. The range degradation away from the antenna's forward direction is consistent with UHF RFID behavior [42, 62]. With the following results, it can be said that our RFID antenna can recognize tags within a 60° field of view up to a distance of 7 m.

The simulated 18 t rigid-truck layout tested one parcel carrying both a BLE beacon and an RFID tag across 33 planned measurement positions. The backend recorded 578 events in the test window: 465 BLE scans and 113 RFID scans. BLE remained observable across the tested volume, while RFID became unreliable in the far end of the layout, especially around 7–8 m from the corner-mounted reader. Points P15, P28, P29, P30,

P31, and P33 were positions where BLE remained reliable, but RFID was not consistently sensed. This reaffirms the chapter 4 decision to support both modalities because different RF technologies suit different sensing roles. A more powerful or differently placed RFID antenna could expand the sensing capabilities of the prototype, but likely at higher cost. In the evaluated setup, RFID is therefore best interpreted as loading or portal confirmation evidence, not as continuous full-volume truck coverage.

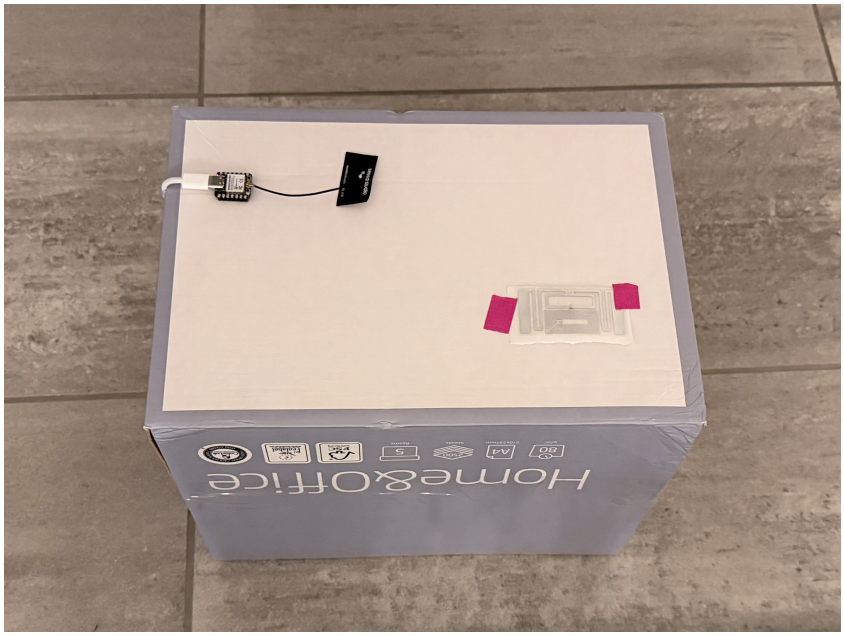

Figure 5.4: Box-mounted BLE beacon and RFID tag used in the simulated 18 t rigid-truck layout.

Source: Own photograph.

This result is consistent with RFID literature. Passive UHF RFID performance depends on antenna geometry, polarization, and nearby materials. Metal and water-rich objects can additionally reduce read performance [30, 42, 62]. Pallet-level RFID research also shows that readings depend on propagation inside the pallet, reader-tag distance, and the packaging materials and goods around the tagged item [3]. The practical implication for **R3** and **R6** is that a production installation should place the RFID antenna where every handling unit passes close to it, such as a load door. Conversely, the RFID tag should be placed on the topmost or outward-facing parcel of a pallet-style handling unit to reduce shielding by surrounding parcels.

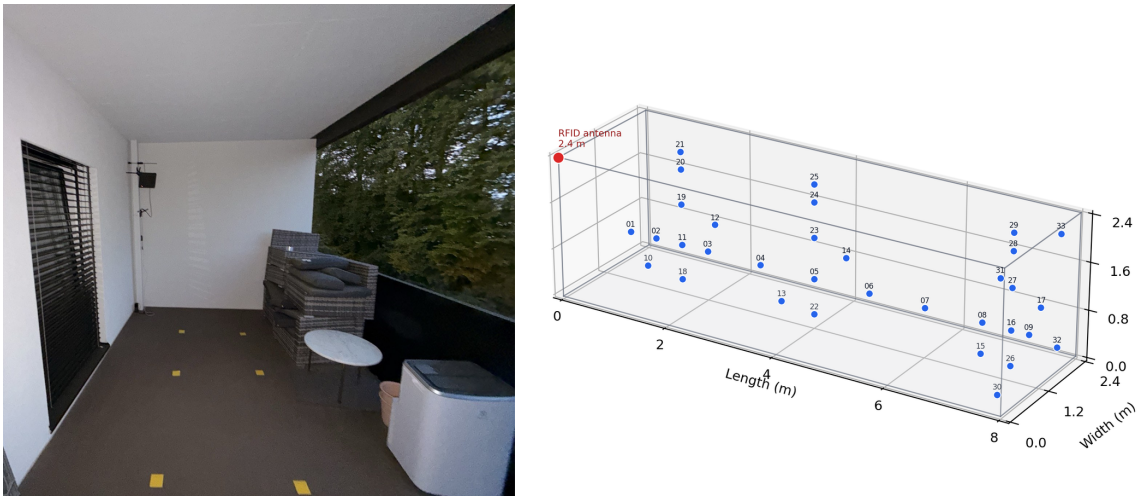

Figure 5.5: 18 t rigid-truck measurement grid on the right and the actual testing setup on the left, with the RFID antenna mounted at 2.4 m in the top-left corner of the 2.4 m-wide measurement area.

Source: Original representation.

The BLE parcel beacon was also evaluated as a cost-relevant identifier. Two **battery-runtime tests** used a short-wake BLE firmware profile with approximately one advertisement cycle per minute. In the first run, the battery only lasted for 23 h 27 min, and for the second 24 h 53 min. This demonstrates short-duration BLE feasibility, but it also limits the claim for **R4**. The tested two-cell CR2032 power path is not sufficient for robust multi-week parcel tracking in its current form. This finding is consistent with Coita's earlier low-cost parcel-tracking thesis, where battery life was also the main practical barrier and empirical power use diverged from theoretical expectations [17]. Even a simple linear scaling to one ping every 20 minutes suggests only an optimistic upper bound of about 20 days, and the observed failure pattern suggests that the repeated pulse-load and regulator behavior may be the limiting factor.

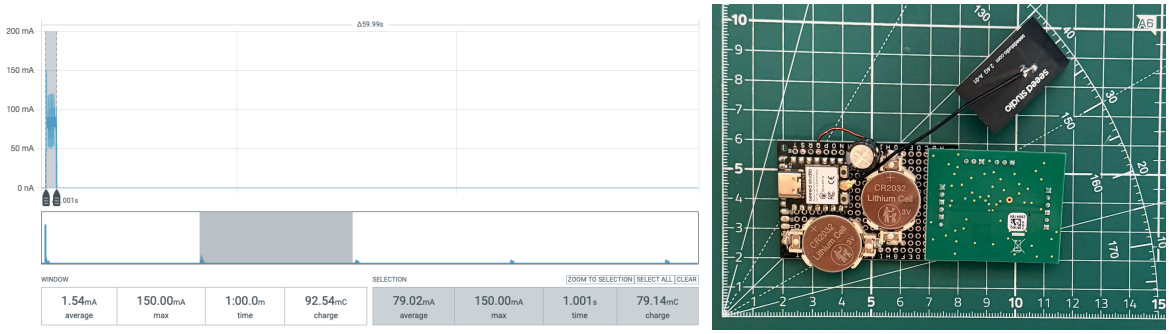

Figure 5.6: The NRF PPK2 screenshot on the left shows the one-minute beacon power profile and selected one-second wake interval, while the photo on the right shows the tested BLE beacon powered by two CR2032 cells.

Source: Original representation.

5.2.2 Tracker Runtime, Mobile Upload, and Recovery

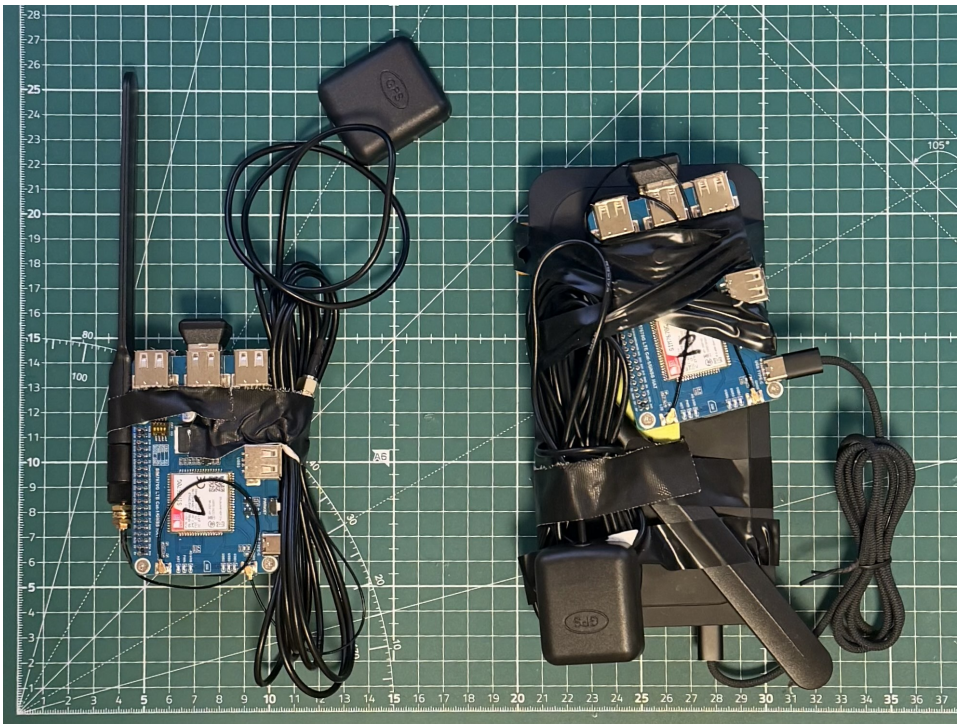

Figure 5.7: Two configurations used during Tracker testing. Tracker components only on the left and tracker with attached power bank on the right.

Source: Own photograph.

The tracker-side tests evaluate whether the sensing evidence can survive the operational conditions expected from a retrofit vehicle node. Across the multiple road tests performed,

the trackers covered 22 individual routes, including 18 sessions with BLE overlap. In aggregate, these sessions covered 76.91 h of logged road movement, 7812 GPS fixes, 1269 BLE scan events, and about 973.9 km of robust route movement. During these road tests, two things became apparent. The first was that the mobile data connection can often be interrupted during normal driving. Without a retry policy, this would have caused lost scan points. However, the retrospective view of the map (Figure 5.8) shows how even delayed uploads were able to be rebuilt as continuous route segments. The realization that longer tunnel segments cause a gap in location inputs was the second finding from these road tests.

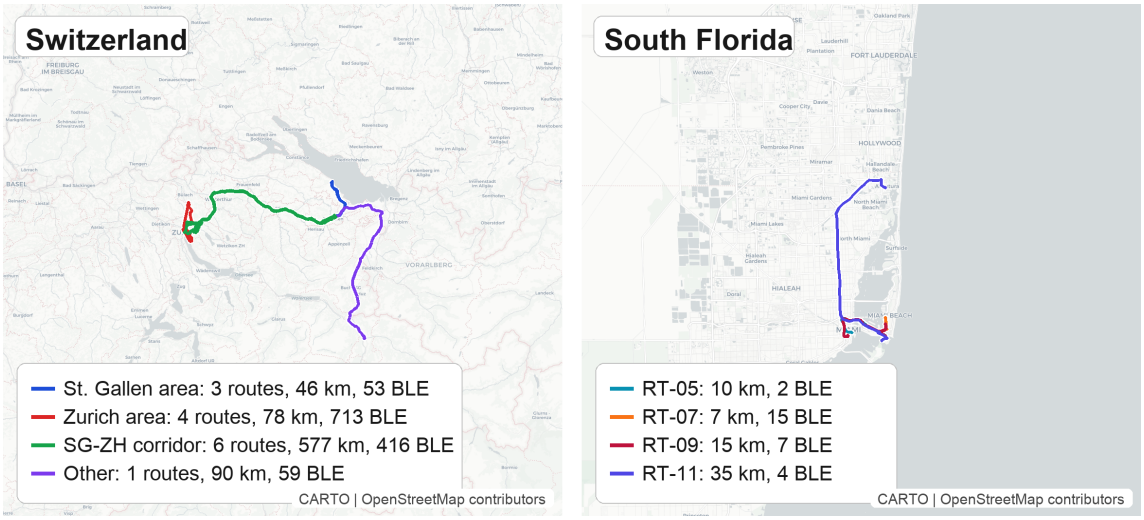

Figure 5.8: Recovered road-test route evidence with BLE overlap in Switzerland and South Florida.

Source: Original representation. Basemap from CARTO and OpenStreetMap contributors [14, 63]

The mobile data usage test used one explicitly selected session. This road run used 4,301,473 bytes, or 4.30 MB, over the corrected 2.28 h active-driving window summarized in Table 5.2, while the exact road window produced 715 backend-confirmed raw events including BLE scans, heartbeats, and GPS information. This supports **R4** because the prototype's visibility stream did not require large mobile-data volumes in this configuration. There remains, however, room for improvement or optimization. A configuration where captured events get stored in the tracker and only uploaded every 15 minutes would reduce communication overhead, with the drawback of inhibiting real-time information.

Table 5.2: Mobile-data session window for the selected road run.

Start time CEST	End time CEST	Duration	Data usage (bytes)	Network
29/04/26 22:32	30/04/26 00:49	02:17	4,301,473	Swisscom

Source: Data-usage export from `hologram.io` [37].

Boot and recovery tests measure the time between power loss and the next successful upload of an identifier scan to the backend. This was chosen as the metric for **R6** because boot time alone does not ensure that the capture-to-backend pipeline is already functioning. The two reboot runs produced local recovery times of 13.4–36.1 s, with an average of 24.7 s for the first new scan to appear in the backend after power to the tracker was restored. The post-restart queue-drain completion time was longer, 45.6–77.9 s, because queue processing added delay after local event creation. Interestingly, it seemed like the queue size had an effect on the time the tracker took to become operational again. The recovery time that is lower than 60 s is a result that supports the outbox design and robustness and shows that an intermittent power loss does not result in a substantial data gap.

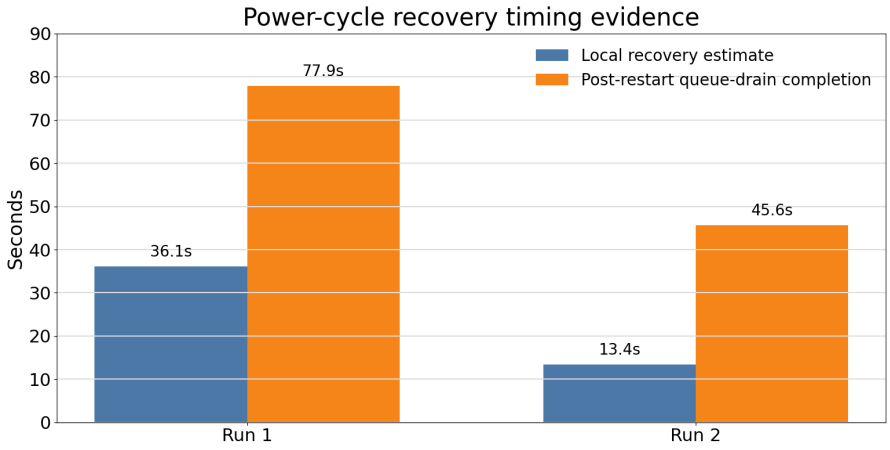

Figure 5.9: Power-cycle recovery timing comparison between local tracker recovery estimates and post-restart queue-drain completion.

Source: Original representation.

The radio-coexistence test is the strongest single stress test of simultaneous tracker operation. One BLE beacon's firmware was changed to emit 60 unique BLE identifiers, each running on their own one-minute cycle. At the same time, 16 RFID tags were held

against the RFID antenna. The tracker-side deduplication was addressing the constant inflow of repeated RFID tags, but still managed to forward the non-duplicate requests during this stress test.

During the coexistence test window, Convex received 540 events from `pi-5-gateway-02`: 334 BLE scans, 88 RFID scans, 58 GNSS fixes, one GNSS no-fix event, and 59 heartbeats. All 60 simulated BLE identities and 16 RFID EPCs appeared in the backend data. Median ingest lag was 8.95 s, while the 95th percentile was 14.98 s. The result supports **R1**, **R5**, and **R6** because all major sensing and upload functions stayed active concurrently. It should not, however, be interpreted as sub-second real-time performance under load.

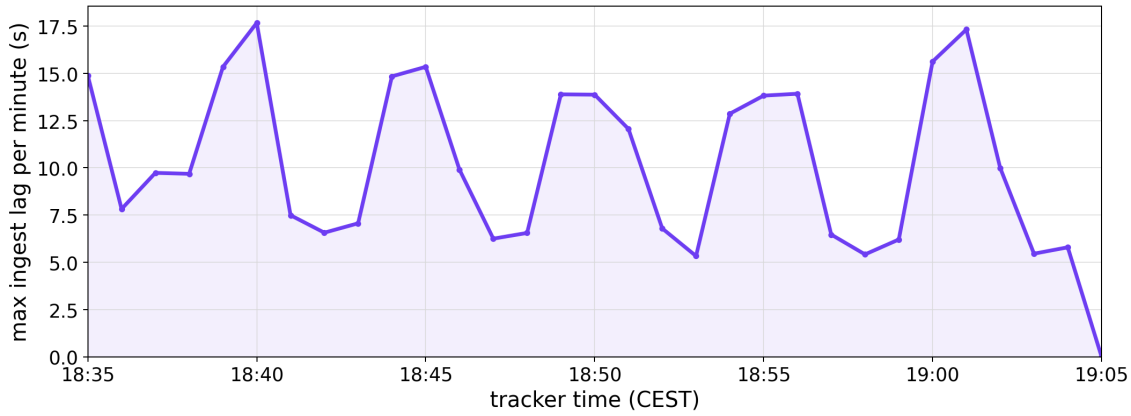

Figure 5.10: Maximum ingest-lag timeline during normal radio-coexistence operation.
Source: Original representation.

The data outage test, where the LTE connection was deliberately disabled, showed the outbox behavior more directly. While LTE was disabled, scans were intentionally triggered, creating a backlog of scans that had to be uploaded once the connection restarted. In the catch-up window, 69 delayed events reached Convex, with a median lag of 486.02 s and a maximum lag of 1004.42 s. Those lag numbers represent the difference between the event captured and the event uploaded to the backend. Because the upload was inhibited, earlier scans naturally have a larger delay, while scans closer to the restart of the SIM connection have a shorter delay. The core is therefore the successful restart of the upload sequence supporting the **R6** outbox concept.

The power and system load test was done during the same period. A USB PD power meter was inserted between the Raspberry Pi's power source and the Raspberry Pi itself during active performance tests. The USB meter captured approximately 5.10 V and

1.40 A, corresponding to 7.14 W. These measurements were the peak power draw rather than the average. The Raspberry Pi logs during this test showed low CPU load at around 1.2%, around 7% memory use, low CPU temperature (about 27.4 °C), and no active throttling flag (usually caused when insufficient power is supplied). This is enough to claim that a vehicle-powered prototype is plausible on a technical level, while the low system usage hints towards further optimization possibilities. It is explicitly not a production power budget as the Raspberry Pi 5 is a prototyping host and would likely be replaced or optimized in a production tracker.

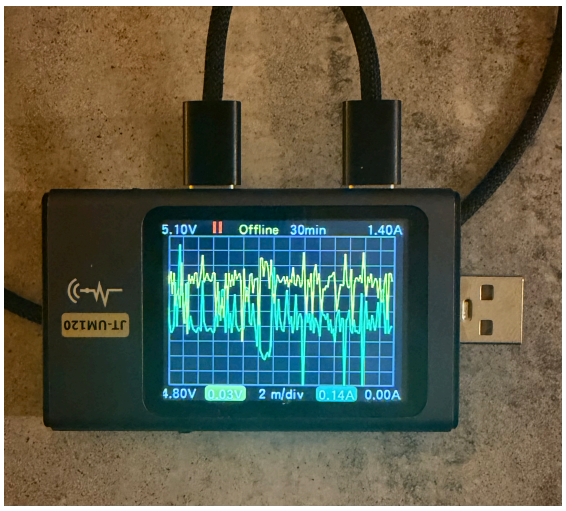

Figure 5.11: USB power-meter snapshot used for the Raspberry Pi tracker power observation.

Source: Own photograph.

5.2.3 Backend Interpretation and Association Reliability

The backend tests focus on **R5** and **R6**: raw events must become usable operational information without creating misleading assignments.

The density stress tests show that the data model preserved identifier identity under higher local event pressure. In a multi-identifier density test, one BLE beacon simulated 60 UUID-like identifiers, and the RFID reader observed 16 unique EPCs. The backend captured all 60 BLE identifiers and all 16 RFID identifiers. BLE and RFID cooldown values suppressed duplicate bursts while preserving distinct identifiers, which is the intended behavior for a tracker that may repeatedly see the same parcel in a truck.

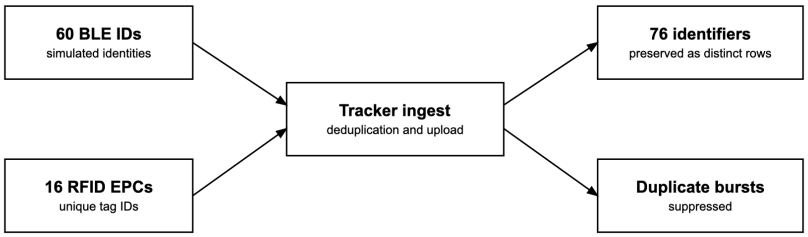

Figure 5.12: Identifier-density stress-test logic showing preservation of distinct BLE and RFID identifiers while duplicate bursts are suppressed.

Source: Original representation.

The simulated association tests then evaluated whether the backend made conservative ownership decisions from imperfect evidence. A synthetic upload test simulated edge cases to test the resolver behavior. Those edge cases covered four situation types: clear ownership, weak or equal competing evidence, handover pressure, and delayed or duplicate event arrival.

The BLE cases checked that weak neighboring detections do not create false ownership, that short signal spikes only create temporary uncertainty, and that a handover requires sustained stronger evidence before it is displayed as a handover. RFID cases, on the other hand, checked that a single read is not enough for assignment, repeated reads can assign, and stale ownership can hand over after sufficient new evidence. These tests do not prove field accuracy, but they directly test the functionality of the decision rules that convert event streams into parcel-truck associations.

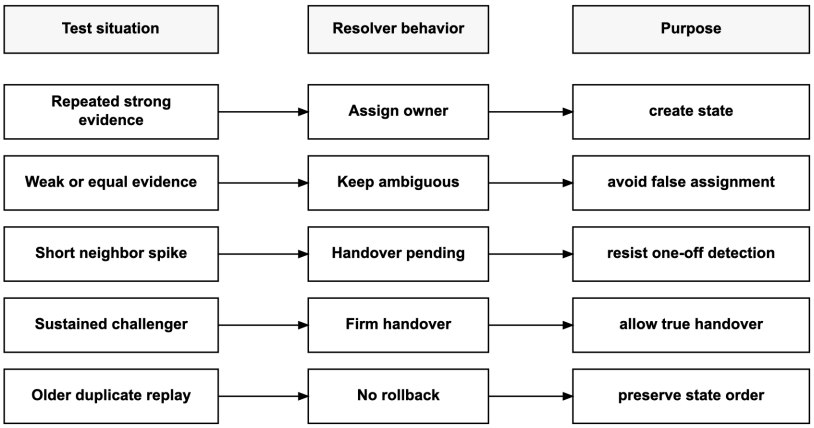

Figure 5.13: Association-resolver test logic showing how different evidence situations map to conservative ownership decisions.

Source: Original representation.

A second **physical BLE association test** complements this model evidence because the same BLE beacon was intentionally moved to be observed by two trackers. Although the beacon was temporarily seen by both trackers, the association did not flip after this short window of ambiguity. This physical repetition confirms that the kind of cross-tracker evidence that can create ambiguity is preserved for the resolver, while the harness tests show how the resolver treats weak, equal, and competing evidence. Together, these tests provide evidence for the internal consistency of **R5**.

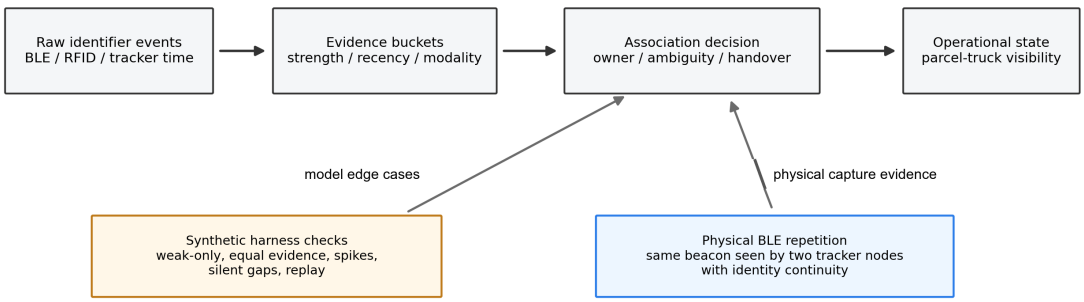

Figure 5.14: Evidence path used to evaluate how raw BLE and RFID events are transformed into conservative parcel-truck association state.

Source: Original representation.

Additional replay tests checked duplicate and out-of-order behavior. Delayed older events did not change the live location. It performed as expected and was inserted in the correct time slot, thereby completing the event history. Out-of-order events were also

handled correctly because the backend oriented itself from the tracker time instead of the ingest time to create the accurate event history. The backend, therefore, succeeded in creating a trustworthy and accurate representation of the events it stored.

5.2.4 Dashboard State and Selective Sharing

This chapter's dashboard tests are bounded to technical functionality, meaning whether backend state can be rendered and updated correctly, not whether the dashboard provides a polished user experience or has been validated with operators.

The synthetic backend-to-dashboard test showed that injected events appeared in the dashboard stream and that live tracker pages could represent no-fix, live, and stale states. This supports the technical precondition for **R2** and supports **R5** at prototype level because the dashboard exposes interpreted state instead of only raw events.

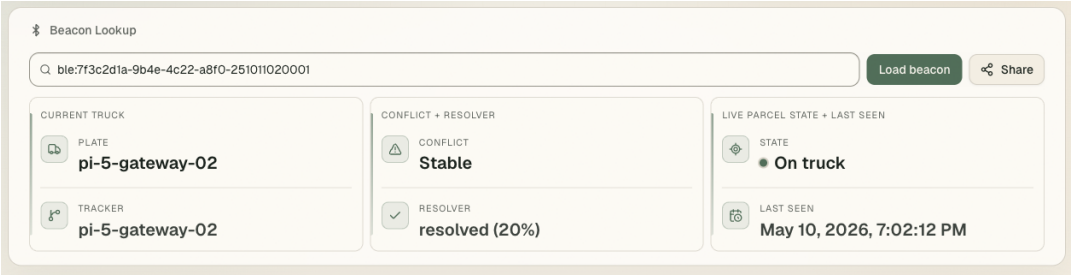

Figure 5.15: Beacon lookup view showing resolved tracker association, resolver state, parcel state, and last-seen evidence for a selected BLE beacon.

Source: Own screenshot.

A historical browsing test was performed to evaluate whether month-old route segments could still display this historical route information. This test successfully showed that the early 2026 route segments from February 20, February 21, February 28, and March 13 were successfully able to be displayed, proving that the dashboard does not only rely on recent or live data but can also be a tool where operators can retrieve information long after a shipment has taken place.

The selective export test supports **R7**. The beacon share mechanism could export selected visibility evidence. A separate leakage audit checked that the export surface was narrower than the full internal telemetry model and was able to confirm that only the selected fields in the interface could be exported.

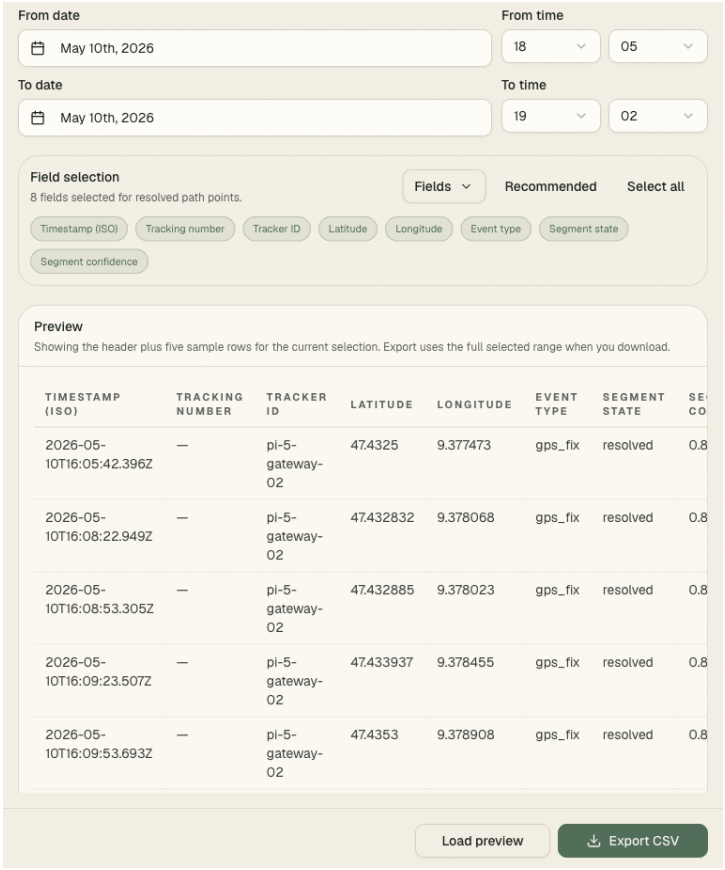

Figure 5.16: Beacon sharing dialog showing selectable export products, field selection, and a preview of the bounded export surface.

Source: Own screenshot.

The dashboard and export layer, therefore, support the principle of selective sharing and live or historical information display, while production deployment would still require authentication, role policies, audit logging, and contractual data rules.

5.2.5 Technical Prototype Evaluation Result

Table 5.3 summarizes the requirement-level result. The artifact is strongest where the thesis needs prototype evidence for the core visibility chain: sensing, local persistence, backend ingest, association-state logic, and selective display. The main unresolved technical boundary is how robustly the prototype behaves under production-scale hardware failure, long-duration radio interference, and real multi-party logistics operation.

Table 5.3: Requirement-linked summary of the technical evaluation evidence.

Evaluation block	Main requirements	Technical evidence strength	Technical result
BLE/RFID sensing	**R1**, **R3**, **R4**, **R6**	Directly evidenced	BLE provided repeated broad presence evidence; RFID provided strong but geometry-sensitive confirmation evidence and is better suited to portal-style loading points.
Tracker runtime and upload	**R1**, **R3**, **R4**, **R6**	Partly evidenced	Road tests, low data use, reboot recovery, coexistence testing, and a 7.14 W power observation support vehicle-side feasibility, but long-duration fleet reliability remains untested.
Outbox and recovery	**R1**, **R5**, **R6**	Directly evidenced	Delayed events after power and LTE disruption reached the backend and live uploads recovered.
Backend association logic	**R5**, **R6**	Directly evidenced in harness tests	Density, replay, and association tests preserved identifiers, suppressed duplicate bursts, avoided stale-state rollback, and handled ambiguous evidence conservatively.
Dashboard and sharing	**R2**, **R5**, **R7**	Partly evidenced	Dashboard state and selective export were demonstrated, but production sharing still needs access control, role policies, and audit governance.

Overall, **SQ2** can be answered positively at prototype level. The artifact demonstrates that a truck-side tracker can collect BLE, RFID, and location events, preserve them locally, upload them to the backend, and transform them into parcel-related state for dashboard use. The result is deliberately bounded. The prototype supports the technical feasibility of the architecture. Broader claims would require a longer pilot with integrated logistics workflows, more vehicles, controlled installation hardware, and operational key performance indicators (KPIs).

5.3 Stakeholder Feedback and Requirements Validation

Having assessed the artifact technically, the evaluation now turns to practitioner feedback on the artifact's operational meaning and relevance to their use cases. A two-minute demo video presenting the artifact's core functions and results was sent to the interview partners. Not all partners responded, and the level of detail in the responses varied. The material used here, therefore, consists of two written post-demo responses, from Quantifor and Valora, and one anonymized follow-up interview with two practitioners from the large international logistics company (Appendix A.7). For that reason, this section should be read as a qualitative plausibility check.

This is consistent with the design science logic of evaluating an artifact against the objectives derived from the problem environment [36, 68]. The feedback is therefore assessed against the seven requirement clusters developed from the earlier interviews and literature. Its value lies in being interpreted together with the technical feasibility evidence above and the literature-informed business-feasibility discussion below. The stakeholder feedback asks whether practitioners recognize the problem-solution fit and understand the artifact's value logic while also providing information on what is still missing for them.

R1: Visibility must reduce blind intervals between milestones. The feedback reiterates the importance of addressing this requirement. The Valora response described the concept as a prompt to reconsider how track-and-trace could be provided more continuously. It explicitly named RFID as one possible route to achieve such a goal (Valora email correspondence). In the follow-up interview, one practitioner described existing tracking as milestone- and scanpoint-based, while both practitioners clarified that, at best, handheld or stationary scanner device events still shape much of the operational visibility available to them today (Appendix A.7). The same interview also hinted at the cost of blind intervals. Static hub-to-hub running times become rough assumptions when congestion, weather events, or even breakdowns occur, and no live location feed to communicate this is available (Appendix A.7). Against this requirement, the artifact can be considered supported because the implemented truck-side tracker and backend event pipeline enable the reporting of location and parcel events between formal scan milestones. This directly addresses the blind-interval problem reiterated by the stakeholders. At the same time, the feedback qualifies the requirement. One practitioner cautioned that location tracking alone does not automatically produce a better ETA unless further contextual information is incorporated (Appendix A.7).

R2: Arrival and load information must support receiving preparation. Practitioners repeatedly framed visibility as more than an entry in a database. One practitioner identified ETA as the central customer-facing unit of information and described shipment location as a means toward better arrival estimation rather than as an end in itself (Appendix A.7). At the same time, the follow-up interview suggested that the strongest potential of the concept may lie in better internal execution and control, while both practitioners connected improved information to fewer phone and email clarifications about shipment status (Appendix A.7). This supports the thesis decision to treat visibility

as operational decision support. The example of time-critical deliveries also shows that load order information and shipment context matter for recipients that require the cargo to be unloaded as quickly as possible (Appendix A.7). The feedback, therefore, supports the current implementation path. The prototype provides arrival and load-context evidence supporting but not yet fully automating receiving preparation.

R3: Operational fit and retrofitability must constrain the design. The feedback is especially strong on the chances of adoption of such a system. The Valora response noted that practical usefulness depends entirely on whether an external logistics provider would participate (Valora email correspondence). The follow-up interview described the same issue from the logistics provider side. In subcontracted and sub-subcontracted networks, the difficulty is often organizational rather than purely technical because the actor arranging transport may not know until late which company will actually carry the goods (Appendix A.7). This uncertainty would require all partnering companies to be equipped with the same tracker for the logistics company to consider installing such a tracker-dependent identifier on their parcels. The practitioners appreciated the thought that went into retrofitability and ease of implementation of the self-contained system. However, the artifact cannot by itself solve partner willingness or subcontractor governance, which are broader issues in the logistics industry. Additional organizational measures were mentioned as prerequisites in practice (Appendix A.7). The issues that motivated the consideration of **R3** throughout the design process, therefore, remain some of the main adoption constraints and future-work areas.

R4: Cost must be justified by operational use. The feedback confirms that cost cannot be looked at as device price alone. The Valora response explicitly asked how effort and benefit would balance, pointing to likely IT changes and additional investment needs. They also mentioned the importance of company size when it comes to evaluating feasibility (Valora email correspondence). The follow-up interview made the same point from a network perspective. Scan automation can be attractive, but scaling such themes across heterogeneous operations is difficult, especially where hardware investment is involved (Appendix A.7). The feedback also helps identify where the cost-benefit balance may be more favorable. One practitioner described high-value and special transports as situations in which customers may explicitly require location tracking and accept an additional tracker. Conversely, the same interview cautioned that customers who invest in tracking may prefer to deploy their own devices and integrate them into their own

visibility platforms, which limits straightforward monetization by the logistics provider (Appendix A.7). These statements strengthen the importance of keeping the solution low-cost but warn that even a low-cost artifact needs to generate value quickly for adopters in order to be credibly implemented by the industry. The following cost analysis in Section 5.4 shows where these options become economically plausible.

R5: Raw events must become usable operational information. Stakeholders supported the thesis distinction between raw capture and usable state. The Quantifor written feedback valued backend storage because it enables active tracking and, in a later step, anomaly detection and optimization suggestions (Quantifor email correspondence). This directly fits the artifact's event-to-state architecture. In the follow-up interview, one practitioner described the goal as a more accurate digital representation of the physical shipment flow (Appendix A.7). The feedback also cautions against excessive data accumulation. One practitioner emphasized that future digital-twin-like systems should display critical states rather than overload users with information (Appendix A.7). The prototype addresses this principle by exposing the derived parcel state rather than showing raw capture data. However, with more parcels in the system, stakeholders along the supply chain would need more filtering options to surface only the relevant shipments.

R6: Visibility information must be reliable and trustworthy enough for use. The stakeholder feedback did not provide a field reliability test, so **R6** remains more strongly supported by the technical evaluation than by subjective comments. Nevertheless, the comments clarify what trustworthy visibility would have to support. The correspondent from Quantifor mentioned the importance of persisted backend data to trace the data that led to a more abstract representation in the dashboard (Quantifor email correspondence). The cautions about adoption and technology acceptance also point toward the need for trustworthy data, which is consistent with supply-chain information-sharing research showing that trust and distrust shape whether ICT-enabled sharing is actually used [50]. In some cases, it may be better not to show any data rather than showing uncertain states. The feedback also identified future extensions in which trustworthiness would matter especially strongly, including cold-chain monitoring and shock or tamper detection (Quantifor email correspondence; Appendix A.7). **R6** is therefore partially met at prototype level: the technical evaluation supports reliability of event validation, while stakeholder feedback confirms that these reliability properties are necessary for

operational trust. With deeper integration into customer systems, reliability and data cleanliness would have to be further expanded.

R7: Data must be shared selectively across ownership boundaries. The feedback most clearly validates the selective-sharing logic. One practitioner emphasized that customer information needs vary by use case, while the other argued that customers may only need to know that goods passed relevant scan points, warning against sharing information that reveals internal warehouse detail such as a specific loading position (Appendix A.7). The visibility information, therefore, must be further abstracted to a level that supports the recipient's decision-making process while not revealing too much detail. The same follow-up interview further noted that the additional information the artifact provides may have substantial internal value for the logistics provider, but that most of this data should remain internal (Appendix A.7). At the same time, different customers may demand different configurable status granularities (Appendix A.7). The prototype addresses this requirement at the principle level through reduced export surfaces and selectable sharing fields. Stronger policy logic and partner-specific role definitions remain future work beyond the evaluated artifact.

Stakeholder-feedback synthesis. Taken together, the stakeholder feedback supports the central feasibility claim of the thesis while narrowing its proper interpretation. The feedback confirms that the problem addressed by the main research question is recognizable in practice: milestone-based and partner-dependent visibility still leaves blind intervals, and truck-level sensing becomes useful when it can be related to shipment or handling-unit state.

For **SQ1**, the feedback supports the selected architecture and requirements, especially the need for retrofitability, configurable granularity, and cost-conscious deployment. It also shows that operational fit is the hardest part of the requirement set because partner willingness, subcontractor structures, and organizational standards cannot be solved by the artifact's design alone.

For **SQ2**, the feedback supports the event-to-interpretable-state logic of the prototype. The limited demo, however, did not cover the exact logic of this transformation. This point is therefore mainly supported by the technical evaluation in Section 5.2, the Technical Prototype Evaluation section.

For **SQ3**, the feedback generated additional evidence pointing toward concrete value paths such as fewer status inquiries, stronger internal control, high-value transport tracking, and data sharing. These value paths are built upon in Section 5.4, where business viability is evaluated further.

5.4 Business Feasibility Evaluation

The following section evaluates whether the implemented artifact creates credible value mechanisms and whether those mechanisms are strong enough to justify further deployment-oriented testing. In doing so, it adds the business-feasibility evaluation required to answer **SQ3** accordingly.

The analysis uses four actor terms consistently. The *focal logistics operator* operates the tracker and dashboard. The *sender* or *shipper* provides shipment and content data. The *receiver* or *consignee* consumes arrival and status information. *Partners* are other authorized firms that may receive selected data. The internal use concerns the focal logistics operator and is evaluated as internal process value. The external use concerns selected information shared with senders, receivers, consignees, or other partners. Its value comes from exposing bounded visibility state and, where available, relaying sender-provided shipment context under permission. Both uses draw on the same technical artifact, but they differ in maturity. Internal use is closer to the implemented prototype. External use is more conceptual and represents a literature-backed business model opportunity.

5.4.1 Cost Boundary and Identifier Scenarios

Cost is central to **SQ3** because the artifact is only feasible if it creates more value than it creates costs. The prototype deliberately concentrates a majority of the expensive components on the truck-side node, while the parcel side remains a lightweight identifier layer. Table 5.4 gives the truck- and parcel-side price anchors used for the business-feasibility calculation. These are public list-price or procurement anchors, not negotiated quotes for a whole-fleet deployment scale. Connectivity fees, cloud costs, installation labor, enclosure design, certification, support, maintenance, printed circuit board design, assembly, and enclosure production are excluded from the hardware table. These factors were considered in the design phase through the operational-fit and cost-value requirements in Sections 4.1.3 and 4.1.4, but they are omitted from the calculation here

because they depend strongly on the deployment model. The later break-even calculation therefore tests a bounded hardware-cost target rather than the full cost of implementing and operating a production system.

Table 5.4: Hardware price anchors used for the bounded cost scenarios.

Layer	Item and role	1-unit price	Scale price	Price source and interpretation
Truck host	Raspberry Pi 5, 4GB *Prototype SBC*	$110.00	quote required	Mouser list price [58]. *Tested host.*
Truck host alternative	Raspberry Pi Zero 2 W *Lower-cost SBC reference*	$15.00	not listed	Official Raspberry Pi price [72]. *Future optimization anchor.*
Truck connectivity	Waveshare SIM7670G *LTE Cat-1/GNSS HAT*	$44.99	$43.63 at 4+	Waveshare list price [89]. *No large-volume quote listed.*
Truck power	45 W vehicle USB power adapter *12 V to USB power anchor*	$24.99	not listed	B&H list price [7]. *Accessory-socket anchor.*
Truck cable	Anker 10 ft USB-C 100 W cable *3 m-class truck cable anchor*	$9.99	$7.42 at 10+	Anker list price and bulk discount [4].
RFID reader	Yanzeo SR791 *Integrated UHF reader/antenna*	$269.00	quote required	Procurement anchor. [21, 96]. *RFID scenarios only.*
BLE identifier	Seeed XIAO ESP32-C3 *Development-board beacon anchor*	$4.99	not listed	Mouser list price [53]. *Tested development board.*
BLE radio module	ESP32-C3-MINI-1-N4-A *Beacon radio module*	$3.58	$1.99 at 500	Mouser list and scale prices [56]. *Scale anchor used in the complete BLE beacon cost.*
BLE power management	TPSM83100SIUR *Buck-boost module proxy*	$7.27	$4.40 at 500	Mouser list and scale prices [59]. *Scale anchor for the low-power battery circuit.*
BLE identifier power	Two Panasonic CR2032 cells *Beacon battery anchor*	$0.72	$0.426 at 1,000	Mouser public US price per cell [55].
BLE battery contacts	Two CR2032 holders *Battery-holder anchor*	$1.22	$0.766 at 500	Mouser list and scale prices for BAT-HLD-003-SMT [54].
BLE buffer capacitor	Nichicon UKL0J471MPD *Low-leakage capacitor anchor*	$0.73	$0.315 at 100	Mouser list and scale prices [60]. *Scale anchor for the tested component.*
RFID identifier	Passive Murata UHF tag *RFID variable-cost proxy*	$0.63	$0.302 at 1,000	Mouser list and scale prices [57]. *Passive RFID variable-cost proxy.*

The Raspberry Pi 5 is a strong prototyping platform, but an expensive production assumption. The same design principle could be shifted toward a smaller Compute Module or

a microcontroller-plus-modem architecture. The Raspberry Pi Zero 2 W price anchor, therefore, illustrates the scale of this optimization opportunity, although it is not used in the scenario calculation because the evaluated tracker used the Raspberry Pi 5 and a corresponding LTE/GNSS HAT.

The parcel-side beacon cost is sensitive to bill-of-materials choices. The active BLE identifier used in the scenarios is therefore not priced as a bare radio chip. It is priced as a conservative complete battery beacon using the module-level ESP32-C3 scale price, the power-management component, two CR2032 cells, two holders, and the buffer capacitor:

$$C_{\mathrm{BLE}} = 1.99 + 4.40 + (2 \times 0.426) + (2 \times 0.766) + 0.315 \approx \$9.09$$

The passive RFID anchor is simpler because the priced passive tag itself is the parcel-side identifier.

$$C_{\mathrm{RFID}} = \$0.302$$

Table 5.5 translates the price anchors into four parcel-side scenarios with two different truck-side fixed-cost settings. The 5,000-parcel cycle is a conservative denominator for one truck-side tracker and means 5,000 unique shipment observations over the replacement period, or approximately three tracked shipments per day over five years. All scenarios assume that parcel-side identifiers are not recovered after the shipment cycle. This is conservative for BLE and realistic for passive RFID. The value of eight is a conservative planning value that is large enough to test grouped-use economics while still plausible for roll containers or small palletized groups.

Table 5.5: Illustrative one-cycle cost scenarios using the price anchors in Table 5.4.

Scenario	Truck fixed cost	Identifier assumption	Variable and one-cycle result
BLE identifier on every parcel	$190	$9.09 per complete BLE beacon	Variable cost: $9.09 per parcel. One-cycle result: $45,635 for 5,000 parcels, or **$9.13** per parcel including tracker cost.
BLE identifier on grouped handling unit	$190	$9.09 spread across 8 parcels	Variable cost: $1.14 per parcel. One-cycle result: $5,871 for 5,000 parcels, or **$1.17** per parcel including tracker cost.
RFID tag on every parcel	$459	$0.302 per RFID tag	Variable cost: $0.30 per parcel. One-cycle result: $1,969 for 5,000 parcels, or **$0.39** per parcel including tracker cost.
RFID tag on grouped handling unit	$459	$0.302 spread across 8 parcels	Variable cost: $0.04 per parcel. One-cycle result: $648 for 5,000 parcels, or **$0.13** per parcel including tracker cost.

The scenario result is clear. A BLE identifier on every parcel is not a plausible general-purpose cost path under the one-way loss assumption. BLE becomes interesting only when the identifier is reusable or assigned to a grouped handling unit. With the $9.09 complete-BLE-beacon cost anchor, grouped BLE falls below $1 variable cost per parcel only once the identifier represents at least ten parcels. It becomes cheaper than a passive RFID tag only when the BLE identifier represents at least 31 parcels. RFID, therefore, remains the stronger mass-cost path, while BLE remains useful where repeated observations, stronger active identity, or battery-powered extensions such as condition sensing justify the higher unit cost.

The business implication is that neither identifier technology dominates every use case. The deployment logic must choose the identifier path by object granularity. The truck-side tracker cost is less sensitive to parcel volume because it is reused across shipment cycles, while the parcel-side identifier cost scales with the number and granularity of represented objects.

The following internal benefit calculation uses the $0.39 RFID-every-parcel scenario as the bounded hardware-cost target because it is the more defensible mass-cost path. The grouped BLE scenario is treated as an additional higher-cost threshold at $1.17 per parcel-equivalent shipment, while BLE on every parcel is not used as the general baseline due to its prohibitive cost.

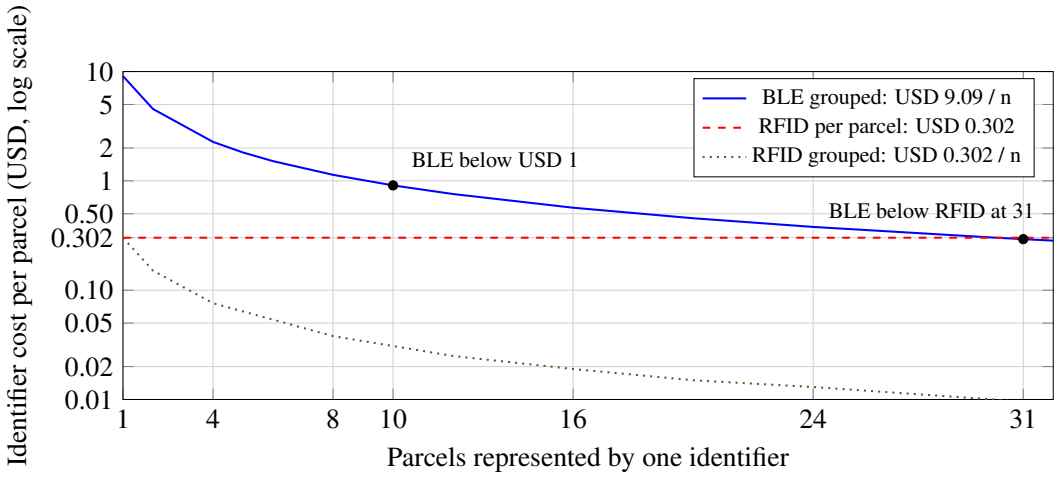

Figure 5.17: Variable identifier cost per parcel under grouped-identifier assumptions with logarithmic cost scaling.

Source: Original representation.

5.4.2 Internal Cost Justification Through Avoided Manual Status-Resolution Work

The internal use case concerns the logistics actor that operates the artifact directly. The relevant value mechanism is the operational answer the artifact maintains. What parcel or handling unit is associated with what truck, and what is the truck's last reliable position? Kalaiarasan et al. show that poor visibility limits deviation management and ETA-related decision-making, which converts directly into manual labor that the artifact can shorten [39]. This subsection, therefore, values the prototype by the labor it can remove from a tightly defined process: the manual resolution of "where is the truck with this shipment?" inquiries, a category commonly referred to as where-is-my-order (WISMO) inquiries.

For inseparable handling units, costs saved should be calculated on a per shipment basis and not a per parcel basis. If ten printers are shipped together in one inseparable handling unit, the customer does not generate ten separate WISMO inquiries, but one inquiry concerning the entire handling unit. Therefore, it would be wrong to use the grouped identifier cost for the following scenario. It is more sensible to use the conservative assumption of one identifier solving one potential inquiry case. Therefore, the $0.39 RFID-per-shipment path is used as the internal per-shipment hardware target.

5.4.2.1 Methodological Anchor: Avoided Status-Resolution Labor

Time-driven activity-based costing (TDABC) supports labor-minute substitution as the primary internal valuation logic. Afonso and Santana apply TDABC to logistics processes and use activity cost as the product of activity time and the cost per time unit. This results in minutes saved being the unit of value when evaluating the value that an information system creates by shortening activities [1].

The labor cost anchor used is the hourly wage of $36.57 used by the U.S. Federal Register for Cargo and Freight Agents [84]. This is a U.S. anchor and should be replaced by company-specific labor cost in future pilots. A per-shipment value target of $0.39 corresponds to about 38.4 seconds of labor that the artifact should save.

5.4.2.2 The Wycislak Case as an Inquiry-Frequency Baseline

The clearest baseline for this activity comes from Wycislak's case study of a real-time transportation visibility platform deployment in the European control tower of a multinational FMCG company [90]. The network covers 45 own factories and 60 warehouses, handles approximately 260,000 full truckloads per year through 110 transport service providers, and operates with 65% subcontracted shipments. The deployment goals include improving internal customer service, reducing demurrage, and replacing manual track-and-trace work.

For the internal break-even calculation, the relevant figures are the reported inquiry volume and annual shipment volume. Before automated visibility, transport planners in the control tower received an average of 200 emails and 50 calls per day about truck location [90]. Across 250 working days, this equals 62,500 status-resolution inquiries per year. Against 260,000 annual shipments, the documented inquiry frequency is therefore approximately 24%. This inquiry-frequency baseline is used below.

5.4.2.3 Per-Shipment Baseline and Break-Even Active Labor

The value calculation is framed as a break-even active-labor question. The annual hardware cost that has to be covered is:

$$C_{\text{annual}} = \text{hardware cost} \times 260{,}000$$

For the RFID-on-every-shipment scenario, where the per-shipment hardware target is \$0.39, this results in annual hardware cost of:

$$C_{\text{annual,RFID}} = 0.39 \times 260{,}000 = \$101{,}400$$

Using the \$36.57 hourly wage anchor, the active labor time that must be avoided per status-resolution inquiry is:

$$t_{\text{active,breakeven}} = \frac{\text{annual hardware cost}}{62{,}500 \times 36.57}$$

For the RFID-on-every-shipment scenario, the required active labor saving per status-resolution inquiry is therefore:

$$t_{\text{active,breakeven,RFID}} = \frac{101{,}400}{62{,}500 \times 36.57} \approx 0.0444 \text{ hours} \mathrel{\widehat{=}} 2.66 \text{ minutes}$$

The break-even point calculated has to be interpreted in a deliberately narrow manner. The artifact has to prevent about 2.66 minutes of active work per inquiry to justify its cost in this scenario.

Industry cases report 30% and 50% reductions in check-call volumes, while the academic real-time visibility platform case reports a 25% reduction in repetitive manual status-checking tasks and a 60% reduction in manual-task-related operational costs. These figures support the plausibility of material status-resolution savings, but they refer to different baselines and therefore should be treated as mechanism evidence rather than directly interchangeable reduction multipliers [81, 91].

Table 5.6 shows the per-shipment saving if only a limited share of that elapsed burden represents real working time that the artifact can remove. The calculation uses the loaded wage anchor of \$36.57 per hour.

Table 5.6: Per-shipment internal saving if a limited share of the 1.5-hour Wycislak elapsed resolution burden represents real working time that can be saved, evaluated against the $0.39 RFID target and the $1.17 grouped-BLE target.

Reduction in status-resolution burden	Real working time saved per shipment	Coverage of the RFID target	Coverage of the grouped-BLE target
3% RFID break-even	0.65 min per shipment; $0.40	100%	34%
5% low effect	1.08 min per shipment; $0.66	167%	56%
10% moderate effect	2.16 min per shipment; $1.32	335%	113%
15% high effect	3.25 min per shipment; $1.98	502%	169%

5.4.2.4 Bounded Internal Claim

The defensible thesis claim is intentionally illustrative. A European case baseline is combined with public USD hardware anchors and a U.S. cargo freight agent wage to test order-of-magnitude feasibility. Under that reference model, the RFID-every-parcel scenario clears the $0.39 per-shipment hardware target if the artifact prevents or shortens about 2.66 minutes of active status-resolution work per documented inquiry.

To apply such a scenario to a company-specific case, the data that is being used for it should first be gathered from that specific company. The company-specific inquiry frequency would have to be gathered, along with actual time savings that the artifact would provide in answering processes.

5.4.3 External Monetization Through a Visibility Service Layer

The external use case concerns selected information shared with shippers, receivers, or other authorized partners through a controlled service layer. The offerings such a service could offer include last observed parcel position, ETA or exception information, relevant ramp preparation context, and proof of custody for handoffs.

Such external data sharing services have two distinct value paths to the shipper. The first one is direct revenue. Partners who want access to this information need to have a high enough amount of willingness-to-pay to cover the costs that the information provider incurs in collecting the information. indirect operational savings because partners with self-service access produce fewer inquiries. The second value path to the shipper happens when the burden of providing shipment updates can be offloaded to the information user, and thereby generating value by causing fewer WISMO inquiries.

5.4.3.1 Substitution Effect

The internal saving above was valued on the assumption that the operator handles a baseline rate of manual status-resolution contacts and that the visibility layer reduces the labor per contact. When the same data is exposed to the interested parties through a partner-facing portal, those partners may answer their own questions without contacting the shipper in the first place. Summing the internal saving with an external subscription fee can therefore double-count the same underlying inquiry pool. The internal and external value mechanisms should therefore be interpreted separately.

5.4.3.2 Three Monetizable Products on the Same Artifact

The artifact provides the technical preconditions for three monetizable products. They differ in the data they expose, in the value they create for the partner, and in the maturity of the academic and market evidence that supports them. A customer solution could then mix and match the three products to tailor its offering.

Product A: Historical-Data ETA Prediction. The first product is an ETA prediction service that could be trained on the visibility data the artifact accumulates over time. The artifact already collects a ledger of a variety of data, including vehicle and parcel location, truck speed, the routes the truck has taken, the stopping times at different stops, and more.

This input data is in the same class as the data used by Balster et al. in their machine-learning ETA model for intermodal freight transport networks. Their study trains one machine-learning model for each leg using historical transport data together with additional external data. The algorithms predict processing times at logistics nodes and transport times on road and rail routes. Their model then combines the per-leg predictions into an overall ETA model for full ETA coverage of the entire network [8]. Historical visibility data of the kind the artifact generates is therefore a credible input basis for a per-leg ETA model, although the accuracy reachable on a specific lane mix would have to be measured empirically once a deployment dataset exists and once the large amount of data collected to craft such an algorithm is gathered.

Product B: Live ETA Through Mapping-API Integration. The second product is a live ETA status. It would have to combine the most recent parcel location with the destination of the next stop. The artifact yields the most recent credible coordinate of

the parcel. Combined with the destination on the shipment record and a routing service that incorporates live traffic, the system could return a continuously refreshed ETA that responds to congestion, road closures, and time-of-day effects thanks to the external routing service.

This product is most pragmatic in short-term paths. It does not require a trained ETA model and can be offered from the first day of a deployment. Its limit is that mapping-API ETAs are usually calculated for road travel and may ignore freight realities such as mandatory rest periods, loading and unloading dwell times, and regulated stops. The mapping-API ETA is therefore most accurate when the truck is close to the destination, where the remaining leg is short and dominated by road-network conditions. It is best operated alongside Product A. The learned baseline of expected node-to-node times could be overlaid with live road-network adjustments for the active leg.

Product C: Selective Sender-Provided Data Relay. The third product builds on the relay channel established in the architecture. The artifact accepts authorized sender-provided shipment information, such as parcel content context, handling instructions, and downstream-consignee preparation hints. It could relay this data to authorized downstream recipients along with the live status. Troeger et al. identify reusable, standardized event data as the basis for value-added services and describe fTRACE as an example of a commercially operated B2B or B2G visibility service with priced usage rights [82].

This type of data relay service is most valuable where no preceding joint data exchange platform is already in use. The interviews indicated that this can still be the case in fragmented B2B settings today. The relay product is a software offering on top of the same artifact. Its marginal cost per shipment is low beyond connectivity and storage, but it requires sender willingness to share content data and recipient willingness to consume it through the platform.

5.4.3.3 Service-Layer Framing in the Literature

Academic literature supports the general visibility-service framing. Chu et al. introduce a service-oriented supply-chain visibility artifact and explicitly formulate it as Visibility-as-a-Service, with cloud-based metering through mobile services to make visibility more affordable and expandable [16]. Moller et al. taxonomize data-driven logistics business

models and identify visibility services as a central class, where tracking devices and mobile phones generate data that is processed and resold as actionable insight [52]. Heinbach et al. argue that telematics-enabled data can be offered to shippers, consignees, transport operators, and insurance companies, each valuing different attributes [33].

5.4.3.4 Willingness-to-Pay and Market Price Anchors

External pricing has to be anchored in evidence about what comparable customers pay or state they would pay. Industry survey evidence points toward positive shipper demand for visibility and control. Logixboard and FreightWaves report that 47% of shippers consider visibility software a critical part of evaluating a freight forwarder, that more than 80% would pay higher rates for better technology, visibility, and control, and that 20% say they would pay more than 20% higher rates [48].

The clearest external numbers come from current visibility-platform vendors.

Table 5.7: Current visibility-platform price anchors and included services from Descartes MacroPoint, Loadbase, and Shiptail [18, 19, 47, 78].

Vendor	Price per shipment	Services included at this price
Descartes MacroPoint	From $0.10 per load at high volume	Aggregation of carrier, ELD/GPS, TMS, and mobile-app visibility data into real-time tracking, status updates, predictive ETAs, automated notifications, CO_2 reporting, temperature monitoring, customer portals, and backend integrations.
Loadbase	$3.00 per load, plus $50 per month	Aggregation of carrier and driver-provided shipment-status data into real-time tracking, email notifications, multi-stop handling, scheduled send-outs, customer tracking links, simple carrier setup, and no required contract.
Shiptail	$1.00 per load after the first five loads	Aggregation of driver-app and shared tracking-link data into load tracking, load history, customer support, load-status updates, and location visibility.

The general direction is clear: shipment-level visibility prices in current commercial deployments can sit at or above the $0.39 per-shipment cost target, and the near-live positioning data the artifact produces would additionally elevate this specific product compared with other visibility services.

5.4.3.5 The Zero-Willingness-to-Pay Argument

Even if downstream partners place no positive willingness to pay on the visibility service, the logistics operator can still benefit from operating it because partners with self-service

access generate fewer inquiries to the logistics operator. Each deflected inquiry saves the operator-side labor cost of the contact.

This means the external service can have positive expected value to the logistics operator even at zero external revenue, as long as it deflects a non-zero share of inbound contacts. The logistics operator's decision is therefore not only whether customers will pay directly, but how much to invest in service delivery, such as portals, APIs, and partner onboarding, given that some of the value returns indirectly through deflected inquiries.

5.4.3.6 Bounded External Claim

The defensible thesis claim is therefore that selected visibility and sender-provided shipment information can plausibly be monetized through a partner-facing service layer at price points compatible with the $0.39 per-shipment cost target and that the same service layer can create value to the shipper even at zero direct external revenue through inquiry deflection.

The claim rests on peer-reviewed support for visibility-as-a-service architecture. ETA modeling theoretically functions on the data the artifact generates [8, 16, 52, 82]. Industry willingness-to-pay is evidenced by real-world price anchors pointing towards a comparable price [18, 19, 47, 48, 78].

Before an external business offering could be launched, however, a further market validation of different price targets for such a solution should be conducted to gain better market readiness data on the value of the visibility that the artifact creates.

5.4.4 Closing the SQ3 Argument

The cost-boundary analysis shows that passive RFID remains the defensible mass adoption cost path, while active BLE is conditional because the conservative hardware model does not support a one-BLE-identifier-per-parcel deployment as a general baseline. The internal cost-justification analysis then demonstrates that the $0.39 RFID-per-shipment target is covered if the artifact prevents or shortens about 2.66 minutes of active status-resolution work per documented inquiry. The external service-layer analysis demonstrates, through different visibility-as-a-service business models, how value could be gained from third parties. The artifact generates the kind of data that a partner-facing visibility service has to operate on top of. Similar services are already commercially viable at comparable

or even higher price points. There even exist scenarios where the service could have positive value to the shipper when offering the service for free.

The bounded answer to **SQ3** is therefore that the artifact creates credible internal and external value mechanisms whose per-shipment value plausibly covers the per-shipment hardware cost under the RFID cost path. The active BLE path should be read more cautiously. It only becomes plausible when the identifier represents a grouped shipment or when the shipment's high value justifies the marginal gain in data accuracy. A future scenario where BLE could become viable is when the technology gets inexpensive enough to support such a deployment. The validation of the metrics used in the above scenarios, including specific reduction percentages, willingness-to-pay levels, BLE hardware costs, and the degree of partner adoption, remains a deployment-stage question out of scope for this thesis.

5.5 Integrated Discussion

The evaluation sections examined the artifact from technical, stakeholder, and business-feasibility perspectives. This integrated discussion brings these findings back to the thesis goal: Assessing whether a low-cost, retrofittable visibility artifact can reduce blind intervals in B2B 3PL handoffs.

From data collection to usable visibility. The central takeaway from this thesis is that demand for better information in logistics does exist, and it is a demand primarily for usable visibility, not raw tracking data. Additional location points only increase utility when they mitigate uncertainty early enough for an actor to respond. What is coming? Where is it? Is it on track? When will it arrive? Are the critical questions. This is where the conceptual distinction introduced earlier in the thesis becomes important. Tracking can generate position or signal evidence, tracing can reconstruct a history, and visibility turns this evidence into operational knowledge. The artifact is strongest when it operates in the middle of this space by creating visibility from the tracking and tracing data it collects.

Logistics operations can fail informationally before they fail physically. The artifact addresses this informational gap. The prototype's possible value paths are therefore fewer status inquiries, earlier receiving preparation, and stronger exception handling. Without this, the shipment may still be moving through the network, but if the relevant actor

does not know its current state, processes must be planned without the necessary level of certainty.

Visibility data as future option value. At the same time, the visibility data stream has the potential to deliver value beyond current operational use cases. The structured history of transport events collected can become the input base for later analytics like ETA learning, anomaly detection, flow optimization, digital-twin-style representations, and AI-assisted decision support. Other industries have already shown that collected data can become more valuable once later analytical capabilities make new uses possible. The big data trend motivated many companies to collect large amounts of data before this data was able to be used to target users with advertisements [20], for example.

Truck-first division of labor. The strongest architectural lesson is that the cost problem, that heavily impacts adoption odds, has to be addressed through technical division of labor. The truck-first architecture reduces burden by placing the heavier technical functions on reusable truck-side hardware and keeping the task of the parcel-side element limited to emitting presence information.

Besides saving on cost, this architecture also adds strategic value because it keeps the sensing layer adaptable. The retrofittable truck-side tracker could, in the future, allow additional identifiers to be added with minimal extension needed. This expandability is especially important to keep in mind because there is currently no universal, low-cost identifier that can be attached to every single parcel, and the industry standard could still change.

BLE and RFID as complementary boundaries. The evaluation also makes the technology boundary clearer. BLE is technically feasible under the tested conditions, but the economics of every-parcel BLE are weak under the current cost model. Mass production could reduce the beacon cost, and a less powerful chip could reduce power consumption. These improvements could make BLE more attractive, but as things stand today, the evaluated cost level is too high for ordinary, one-way, mass parcel tracking.

Passive RFID, on the other hand, is economically attractive and much closer to a mass-deployable parcel identity mechanism. In the tested truck-like setup, however, RFID performance did not perform well at desired ranges. Dense parcel environments create additional shielding problems. Passive RFID evidence could therefore be useful at

controlled read points, but it has proven itself to be less convincing as a single technology for reliable full truck volume visibility.

Economic interpretation. The thesis supports the claim that visibility data is valuable, but each hardware configuration still has to earn its place economically. The strongest economic case appears when visibility removes uncertainty that would lead to human resolution work. Even modest task reductions can justify low-cost RFID-style deployments. The same logic is hard to justify for disposable BLE on every parcel because the identifier cost is simply too high. The value of the information stream is therefore dependent on the cost of the identification mechanism.

Infrastructure and bounded sharing. The near-term role of the artifact is likely to serve as an information infrastructure. The installing party can equip trucks with the reusable sensing system and make the resulting visibility base available to the focal logistics operator, and he can then decide what parcels deserve to have an identifier mounted to them and what identifier it would be. This avoids forcing a blanket identifier strategy onto every parcel. The $459 truck-side tracker is, therefore, easier to justify than a non-retrofittable system that requires extensive replacement of vehicle hardware or changes in third-party processes.

Bounded sharing is part of that same business logic. The access layer exposes operational state while hiding internal details, which makes the system more acceptable and more useful. The bounded sharing groundwork also creates the basis for visibility services, but the commercial value of such services enabled by the prototype will depend on the specific company setting and integration work.

Explaining the adoption gap. Today's absence of universal parcel-level real-time visibility reflects real adoption hurdles rather than conservatism for its own sake. Hardware cost is critical in low-margin parcel flows. Battery life becomes a concern when identifiers require constant servicing. Device recovery is an issue when hardware travels through open networks. RF reliability is challenging in dense and variable loading conditions. Partner adoption is central when visibility depends on subcontractors. Integration effort increases when existing customer systems must consume the data. Benefit distribution is complex because the paying actor and the main beneficiary can differ.

This explains the cautious adoption observed in practice. The prototype shows that a truck-first visibility architecture is technically plausible and economically promising under specific conditions. The prototype, at the same time, shows why broad rollout remains difficult. The more realistic conclusion is that adoption will begin where uncertainty is costly enough to justify the effort and cost associated with overcoming the hurdles described above.

Bounded contribution. The contribution of the artifact is therefore practically relevant while also being bounded. It does not demonstrate a one-size-fits-all visibility solution for low-cost parcel-level tracking. Instead, it demonstrates a realistic prototype architecture for improving visibility under current technological and economic constraints. This prototype further exposes the exact barriers a future productized solution would have to overcome to be realistically adopted in the logistics industry.

Chapter 6

Conclusion and Future Work

6.1 Summary of Contributions

This thesis addressed visibility gaps in B2B road-freight shipments handled through 3PL contexts. When goods leave the focal company's direct operational control, visibility is often reduced to scan events and milestone traces. These mechanisms can document that a shipment passed a known point, but they leave blind intervals during transport between those predetermined points. Guided by a DSR approach, the thesis identified requirements for this problem, translated these requirements into an architectural concept, implemented a functional prototype, and evaluated it as a bounded feasibility contribution. In doing this, the thesis makes four primary contributions.

First, the thesis provides a requirements-centered framing of operational visibility needs. It shows that the relevant user demand is not for raw movement data alone, but for decision-relevant state information. This requirement spectrum for a visibility artifact is captured in seven clusters: blind-interval reduction, receiving preparation, retrofitability, cost discipline, event interpretation, trustworthiness, and selective information sharing.

Second, the thesis develops the truck-first visibility architecture as a prototype that addresses the central trade-off between visibility scope and operational reality. Instead of turning every parcel into a battery-powered tracking device, the architecture places connectivity and reusable sensing capabilities on the truck side. The shipment side remains focused on lightweight presence evidence. This division of labor reduces parcel-side complexity while preserving the possibility of object-level visibility by asserting state information in the backend of the prototype architecture.

Third, the thesis presents a functional prototype that supports the technical feasibility of this architecture at prototype level. The artifact collects GNSS, BLE, RFID, liveness, and recovery event data using a reusable truck tracker. It preserves events during connectivity disruptions and forwards them to the backend. The backend then combines that information and makes it visible to different parties through the dashboard.

Fourth, the thesis evaluates the artifact from technical, stakeholder, and business feasibility perspectives. The technical evaluation shows that the core process chain from sensing to the dashboard worked under the tested prototype conditions. Stakeholder feedback confirms that the identified problem and proposed solution direction are recognizable in individual logistics settings. The business-feasibility evaluation identifies credible value mechanisms through reduced status-resolution work and potential partner-facing visibility services. At the same time, the evaluation clarifies the boundary of the contribution. The thesis supports a plausible prototype-level architecture and is not a production-ready system or a proven full-scale business case.

6.2 Final Considerations

The answer to visibility gaps offered by the thesis is architectural rather than dependent on one specific technology. A low-cost visibility prototype can reduce handoff-related blind intervals by combining lightweight shipment-side identification with reusable truck-side sensing and conservative backend interpretation. In this architecture, raw positional and identifier signals become operational visibility only when they are transformed into state and shared with relevant actors.

In response to **SQ1**, the thesis shows that effective visibility prototypes must address blind intervals while respecting retrofitability, cost pressure, trustworthiness, and data-sharing boundaries. These constraints determine whether a visibility concept can plausibly fit the individual fragmented logistics environments.

In response to **SQ2**, the implementation and evaluation show that low-level tracking signals can be captured, stored during connectivity losses, uploaded to the backend, interpreted into decision-relevant state information, and presented in an accessible form. The most important technical lesson is that visibility emerges from the chain that connects sensing, association, user-facing presentation, and interpretation, not from collecting the data alone.

In response to **SQ3**, the thesis identifies credible internal and external value mechanisms. Under a passive RFID or grouped-identifier cost path, these could generate positive ROI. It does not, however, prove realized value in ongoing operations. On the other hand, the use of battery-powered disposable identifiers for every parcel was not able to be justified under the evaluated cost assumptions.

The takeaway is therefore that logistics visibility is not a consequence of data acquisition alone. It arises when collected evidence can be interpreted into a state that stakeholders can rely on and share at the level of detail needed for effective decision-making. The prototype is most impactful as an information infrastructure layer. It does not replace existing logistics systems, but augments them with additional evidence during the blind intervals for the parcels that need this additional data.

This thesis does not present a solution for ubiquitous real-time parcel-level tracking. Instead, it outlines a pragmatic path toward better handoff visibility within the current technical, economic, and organizational landscape. It also helps explain why broad parcel-level real-time visibility has not yet become standard despite a long-standing interest in adopting the technology. These adoption barriers include technical, economic, organizational, and incentive-related constraints.

6.3 Future Work

Future research could first move the developed artifact from the laboratory and prototype testing environment to a supervised field pilot. Essential pending measurements include whether the system reduces status inquiries and decreases the time spent seeking and resolving shipment status information. It could further analyze whether the collected data improves ETA prediction, supports earlier exception detection, or improves pre-arrival preparation for recipients. This research would transform the current plausibility claim into concrete evidence about the artifact's impact on real operations.

Once a reliable repository of historical visibility data becomes available, further work could investigate future use cases that are built on top of the same visibility data that the prototype collects today.

The hardware and identifier strategy should also be hardened for deployment scenarios. A production-ready version would require a ruggedized casing, stable vehicle power integration, improved antenna design or better placement, and radio-frequency interfer-

ence testing. The identifier strategy could be refined through systematic comparison of a variety of association technologies across different shipment types. These alternatives could then be further compared by cost and operational fit.

The information layer should also be integrated more deeply with existing logistics management systems. Future research could evaluate whether such an integration could allow visibility data to trigger operational workflows and support more efficient exception handling. This step would be a necessary addition to this work in order to test whether visibility data changes decisions rather than only improving observability.

Bibliography

[1] Paulo Afonso and Alex Fabiano Bertollo Santana. “Application of the TDABC Model in the Logistics Process Using Different Capacity Cost Rates”. In: *Journal of Industrial Engineering and Management* 9.5 (2016), pp. 1003–1019. DOI: 10.3926/jiem.2086.

[2] Jaehyun Ahn, Haifa Gaza, Juhyeok Lee, Hyeongchan Kim, and Jaewook Byun. “Oliot EPCIS: An Open-Source EPCIS 2.0 System for Supply Chain Transparency”. In: *SoftwareX* 23 (2023), p. 101477. DOI: 10.1016/j.softx.2023.101477.

[3] Ignacio Angulo, Asier Perallos, Leire Azpilicueta, Francisco Falcone, Unai Hernandez-Jayo, Asier Moreno, and Ignacio Julio Garcia Zuazola. “Towards a Traceability System Based on RFID Technology to Check the Content of Pallets within Electronic Devices Supply Chain”. In: *International Journal of Antennas and Propagation* 2013 (2013), pp. 1–9. DOI: 10.1155/2013/263218.

[4] Anker. *New Nylon USB-C to USB-C 100W Cable (10 ft)*. 2026.

[5] Natapat Areerakulkan and Detcharat Sumrit. “Elucidating Critical Success Factors for Digital Transformation in Logistics Service Providers Using a Grey-DANP Approach: Insights from the Thai SMEs Road Freight Transportation Sector”. In: *Journal of Open Innovation: Technology, Market, and Complexity* 11.2 (2025), p. 100549. DOI: 10.1016/j.joitmc.2025.100549.

[6] Aloÿs Augustin, Jiazi Yi, Thomas Clausen, and William M. Townsley. “A Study of LoRa: Long Range & Low Power Networks for the Internet of Things”. In: *Sensors* 16.9 (2016), p. 1466. DOI: 10.3390/s16091466.

[7] B&H Photo Video. *Xcellon 45W Dual USB PD Car Charger*. 2026.

[8] Andreas Balster, Ole Hansen, Hanno Friedrich, and André Ludwig. “An ETA Prediction Model for Intermodal Transport Networks Based on Machine Learning”. In: *Business & Information Systems Engineering* 62.5 (2020), pp. 403–416. DOI: 10.1007/s12599-020-00653-0.

[9] Mark Barratt and Adegoke Oke. “Antecedents of Supply Chain Visibility in Retail Supply Chains: A Resource-Based Theory Perspective”. In: *Journal of Operations Management* 25.6 (2007), pp. 1217–1233. DOI: 10.1016/j.jom.2007.01.003.

[10] Rahul C. Basole and Maciek Nowak. “Assimilation of Tracking Technology in the Supply Chain”. In: *Transportation Research Part E: Logistics and Transportation Review* 114 (2018), pp. 350–370. DOI: 10.1016/j.tre.2016.08.003.

[11] Ygal Bendavid and Luc Cassivi. “Bridging the Gap Between RFID/EPC Concepts, Technological Requirements and Supply Chain E-Business Processes”. In: *Journal of Theoretical and Applied Electronic Commerce Research* 5.3 (2010). DOI: 10.4067/S0718-18762010000300002.

[12] Mario Henrique Callefi, Gilberto Miller Devós Ganga, Moacir Godinho Filho, Elias Ribeiro da Silva, Lauro Osiro, and Vasco Reis. “Illuminating the Road Ahead: Unlocking the Potential of ICTs for Enhanced Data Visibility in Road Transportation”. In: *Industrial Management & Data Systems* 124.2 (2023), pp. 786–819. DOI: 10.1108/IMDS-02-2023-0115.

[13] Vicente Canton Paterna, Anna Calveras Auge, Josep Paradells Aspas, and Maria Perez Bullones. "A Bluetooth Low Energy Indoor Positioning System with Channel Diversity, Weighted Trilateration and Kalman Filtering". In: *Sensors* 17.12 (2017), p. 2927. DOI: 10.3390/s17122927.

[14] CARTO. *CARTO Basemaps*. Accessed 2026-05-11. 2026. URL: https://carto.com/basemaps/.

[15] Bangho Cho, Sung Yul Ryoo, and Kyung Kyu Kim. "Interorganizational Dependence, Information Transparency in Interorganizational Information Systems, and Supply Chain Performance". In: *European Journal of Information Systems* 26.2 (2017), pp. 185–205. DOI: 10.1057/s41303-017-0038-1.

[16] Sung-Chi Chu, Jerrel Leung, Waiman Cheung, and Gang Chen. "Visibility Cloud: A Supply Chain Perspective". In: *ICEB 2015 Proceedings*. 2015, pp. 507–512.

[17] David Coita. "Are We There Yet? Low-Cost Parcel Tracking System". Bachelor's Thesis. University of Zurich, Department of Informatics, 2025.

[18] Descartes MacroPoint. *Understanding Descartes MacroPoint Recurring Costs*. 2025.

[19] Descartes MacroPoint. *What is the Cost of Load Tracking? Visibility Platform Pricing*. 2022.

[20] José van Dijck. "Datafication, Dataism and Dataveillance: Big Data Between Scientific Paradigm and Ideology". In: *Surveillance & Society* 12.2 (2014), pp. 197–208. DOI: 10.24908/ss.v12i2.4776.

[21] eBay. *Yanzeo SR791 UHF 10dbi Long Range RFID Reader Listing*. 2026.

[22] EM Microelectronic. *em|echo-V (EM4425)*. 2026.

[23] Nadine Farquharson, Joash Mageto, and Hemisha Makan. "Effect of Internet of Things on Road Freight Industry". In: *Journal of Transport and Supply Chain Management* 15 (2021). DOI: 10.4102/JTSCM.V15I0.581.

[24] FedEx. *How to Embrace Incremental Innovation in Your Business*. 2026.

[25] Vernon Francis. "Supply Chain Visibility: Lost in Translation?" In: *Supply Chain Management: An International Journal* 13.3 (2008), pp. 180–184. DOI: 10.1108/13598540810871226.

[26] Stephan L. K. Freichel, Johannes Wollenburg, and Daniel A. Wulf. "Challenges of Supply Chain Visibility in Distribution Logistics: A Literature Review". In: *Logistics Research* 15 (2022).

[27] Olivier Gallay, Kari Korpela, Tapio Niemi, and Jukka K. Nurminen. "A Peer-to-Peer Platform for Decentralized Logistics". In: *Proceedings of the Hamburg International Conference of Logistics*. 2017.

[28] George A. Giannopoulos. "Towards a European ITS for Freight Transport and Logistics". In: *European Transport Research Review* 1.4 (2009), pp. 147–161. DOI: 10.1007/s12544-009-0022-5.

[29] Daniel Christopher Goll and Nils-Ole Bolte. "Potential Analysis of Track-and-Trace Systems in the Outbound Logistics of a Swedish Retailer". Master's Thesis. Jönköping University, 2020.

[30] Joshua D. Griffin, Gregory D. Durgin, Andreas Haldi, and Bernard Kippelen. "RF Tag Antenna Performance on Various Materials Using Radio Link Budgets". In: *IEEE Antennas and Wireless Propagation Letters* 5 (2006), pp. 247–250. DOI: 10.1109/LAWP.2006.874072.

[31] GS1. *EPCIS Standard, Release 2.0*. 2022.

[32] Abubaker Haddud, Arthur DeSouza, Anshuman Khare, and Huei Lee. "Examining Potential Benefits and Challenges Associated with the Internet of Things Integration in Supply Chains". In: *Journal of Manufacturing Technology Management* 28.8 (2017), pp. 1055–1085. DOI: 10.1108/JMTM-05-2017-0094.

[33] Christoph Heinbach, Jan Beinke, Friedemann Kammler, and Oliver Thomas. "Data-Driven Forwarding: A Typology of Digital Platforms for Road Freight Transport Management". In: *Electronic Markets* 32.2 (2022), pp. 807–828. DOI: 10.1007/s12525-022-00540-4.

[34] Christoph Heinbach, Pascal Meier, and Oliver Thomas. "Designing a Shared Freight Service Intelligence Platform for Transport Stakeholders Using Mobile Telematics". In: *Information Systems and e-Business Management* 20.4 (2022), pp. 847–888. DOI: 10.1007/s10257-022-00572-5.

[35] Petri Helo and Vinh V. Thai. "Logistics 4.0: Digital Transformation with Smart Connected Tracking and Tracing Devices". In: *International Journal of Production Economics* 275 (2024), p. 109336. DOI: 10.1016/j.ijpe.2024.109336.

[36] Alan R. Hevner, Salvatore T. March, Jinsoo Park, and Sudha Ram. "Design Science in Information Systems Research". In: *MIS Quarterly* 28.1 (2004), pp. 75–106. DOI: 10.2307/25148625.

[37] Hologram. *Data Usage by Session Export*. 2026.

[38] Maria-Eugenia Iacob, Gilang Charismadiptya, Marten van Sinderen, and Jean Paul Sebastian Piest. "An Architecture for Situation-Aware Smart Logistics". In: *2019 IEEE 23rd International Enterprise Distributed Object Computing Workshop*. 2019, pp. 108–117. DOI: 10.1109/EDOCW.2019.00030.

[39] Ravi Kalaiarasan, Tarun Kumar Agrawal, Jan Olhager, Magnus Wiktorsson, and Jannicke Baalsrud Hauge. "Supply Chain Visibility for Improving Inbound Logistics: A Design Science Approach". In: *International Journal of Production Research* 61.15 (2022), pp. 5228–5243. DOI: 10.1080/00207543.2022.2099321.

[40] Ravi Kalaiarasan, Tarun Kumar Agrawal, Magnus Wiktorsson, Jannicke Baalsrud Hauge, and Jan Olhager. "Requirements on Supply Chain Visibility: A Case on Inbound Logistics". In: *Advances in Production Management Systems*. 2021, pp. 115–122. DOI: 10.1007/978-3-030-85902-2_13.

[41] Iurii Konovalenko and André Ludwig. "Event Processing in Supply Chain Management: The Status Quo and Research Outlook". In: *Computers in Industry* 105 (2019), pp. 229–249. DOI: 10.1016/j.compind.2018.12.009.

[42] Hugo Landaluce, Laura Arjona, Asier Perallos, Francisco Falcone, Ignacio Angulo, and Fernando Muralter. "A Review of IoT Sensing Applications and Challenges Using RFID and Wireless Sensor Networks". In: *Sensors* 20.9 (2020), p. 2495. DOI: 10.3390/s20092495.

[43] Chang-Woo Lee, Dong-Gyun Sohn, Min-Gyu Sang, and Chulung Lee. "Empirical Analysis of Barriers to Collaborative Information Sharing in Maritime Logistics Using Fuzzy AHP Approach". In: *Sustainability* 17.4 (2025), p. 1721. DOI: 10.3390/su17041721.

[44] Hau L. Lee and Özalp Özer. "Unlocking the Value of RFID". In: *Production and Operations Management* 16.1 (2007), pp. 40–64. DOI: 10.1111/j.1937-5956.2007.tb00165.x.

[45] Ming K. Lim, Widya Bahr, and S. C. H. Leung. "RFID in the Warehouse: A Literature Analysis (1995–2010) of Its Applications, Benefits, Challenges and Future Trends". In: *International Journal of Production Economics* 145.1 (2013), pp. 409–430. DOI: 10.1016/j.ijpe.2013.05.006.

[46] Rong Liu, Akhil Kumar, and Wil M. P. van der Aalst. "A Formal Modeling Approach for Supply Chain Event Management". In: *Decision Support Systems* 43.3 (2007), pp. 761–778. DOI: 10.1016/j.dss.2006.12.009.

[47] Loadbase. *Loadbase Pricing*. 2026.

[48] Logixboard and FreightWaves. *The Power of Customer Experience in Freight Forwarding*. 2021. URL: https://pages.logixboard.com/hubfs/Whitepapers/Logixboard-Freightwaves-Survey-The-Power-of-Customer-Experience.pdf.

[49] Luca Mainetti, Francesca Mele, Luigi Patrono, Francesco Simone, Maria Laura Stefanizzi, and Roberto Vergallo. "An RFID-Based Tracing and Tracking System for the Fresh Vegetables Supply Chain". In: *International Journal of Antennas and Propagation* 2013 (2013), pp. 1–15. DOI: 10.1155/2013/531364.

[50] Kristijan Mirkovski, Robert M. Davison, and Maris G. Martinsons. "The Effects of Trust and Distrust on ICT-Enabled Information Sharing in Supply Chains". In: *The International Journal of Logistics Management* 30.3 (2019), pp. 892–926. DOI: 10.1108/IJLM-06-2017-0155.

[51] Gonzalo D. Molero et al. "Key Factors for Implementation and Integration of Innovative ICT Solutions in Multimodal Transport of Dangerous Goods". In: *European Transport Research Review* 11 (2019), p. 28. DOI: 10.1186/s12544-019-0363-8.

[52] Frederik Möller, Maleen Stachon, Christina Hoffmann, Henrik Bauhaus, and Boris Otto. "Data-Driven Business Models in Logistics: A Taxonomy of Optimization and Visibility Services". In: *Proceedings of the 53rd Hawaii International Conference on System Sciences*. 2020, pp. 5377–5386. DOI: 10.24251/HICSS.2020.661.

[53] Mouser Electronics. *113991054 Seeed Studio XIAO ESP32C3*. 2026.

[54] Mouser Electronics. *BAT-HLD-003-SMT CR2032 Battery Holder*. 2026.

[55] Mouser Electronics. *CR2032 Panasonic Battery Coin Cell Battery*. 2026.

[56] Mouser Electronics. *ESP32-C3-MINI-1-N4-A Espressif Systems*. 2026.

[57] Mouser Electronics. *LXMSJZNCMH-225 Murata Electronics UHF RFID Tag*. 2026.

[58] Mouser Electronics. *SC1111 Raspberry Pi 5/4GB*. 2026.

[59] Mouser Electronics. *TPSM83100SIUR Texas Instruments Power Management Module*. 2026.

[60] Mouser Electronics. *UKL0J471MPD Nichicon Aluminum Electrolytic Capacitor*. 2026.

[61] Sharareh Naghdi and Kyle O'Keefe. "Detecting and Correcting for Human Obstacles in BLE Trilateration Using Artificial Intelligence". In: *Sensors* 20.5 (2020), p. 1350. DOI: 10.3390/s20051350.

[62] Pavel V. Nikitin and K. V. S. Rao. "Antennas and Propagation in UHF RFID Systems". In: *2008 IEEE International Conference on RFID*. 2008, pp. 277–288. DOI: 10.1109/RFID.2008.4519368.

[63] OpenStreetMap contributors. *OpenStreetMap Copyright and License*. Accessed 2026-05-11. 2026. URL: https://www.openstreetmap.org/copyright.

[64] Sebastian Opriel, Frederik Möller, Gero Strobel, and Boris Otto. "Data Sovereignty in Inter-organizational Information Systems". In: *Business & Information Systems Engineering* 67.6 (2025), pp. 833–853. DOI: 10.1007/s12599-024-00893-4.

[65] Ertunga C. Özelkan and Agnes Galambosi. "When Does RFID Make Business Sense for Managing Supply Chains?" In: *Ubiquitous and Pervasive Computing: Concepts, Methodologies, Tools, and Applications*. 2010, pp. 1250–1283. DOI: 10.4018/jisscm.2008010102.

[66] Shenle Pan, Damien Trentesaux, Duncan McFarlane, Benoit Montreuil, Eric Ballot, and George Q. Huang. "Digital Interoperability in Logistics and Supply Chain Management: State-of-the-Art and Research Avenues Towards Physical Internet". In: *Computers in Industry* 128 (2021), p. 103435. DOI: 10.1016/j.compind.2021.103435.

[67] John Patsavellas, Rashmeet Kaur, and Konstantinos Salonitis. "Supply Chain Control Towers: Technology Push or Market Pull – An Assessment Tool". In: *IET Collaborative Intelligent Manufacturing* 3.3 (2021), pp. 290–302. DOI: 10.1049/cim2.12040.

[68] Ken Peffers, Tuure Tuunanen, Marcus A. Rothenberger, and Samir Chatterjee. "A Design Science Research Methodology for Information Systems Research". In: *Journal of Management Information Systems* 24.3 (2007), pp. 45–77. DOI: 10.2753/MIS0742-1222240302.

[69] Anna Pernestål, Albin Engholm, Marie Bemler, and Gyözö Gidófalvi. "How Will Digitalization Change Road Freight Transport? Scenarios Tested in Sweden". In: *Sustainability* 13.1 (2020), p. 304. DOI: 10.3390/su13010304.

[70] Nicolas Prat, Isabelle Comyn-Wattiau, and Jacky Akoka. "A Taxonomy of Evaluation Methods for Information Systems Artifacts". In: *Journal of Management Information Systems* 32.3 (2015), pp. 229–267. DOI: 10.1080/07421222.2015.1099390.

[71] Ramiro Ramirez, Chien-Yi Huang, Che-An Liao, Po Ting Lin, Hung-Wei Lin, and Ssu-Han Liang. "A Practice of BLE RSSI Measurement for Indoor Positioning". In: *Sensors* 21.15 (2021), p. 5181. DOI: 10.3390/s21155181.

[72] Raspberry Pi Ltd. *Raspberry Pi Zero 2 W*. 2026.

[73] Henrik Anders Ringsberg and Vahid Mirzabeiki. "Effects on Logistic Operations from RFID- and EPCIS-Enabled Traceability". In: *British Food Journal* 116.1 (2013), pp. 104–124. DOI: 10.1108/BFJ-03-2012-0055.

[74] Vivek Roy. "Contrasting Supply Chain Traceability and Supply Chain Visibility: Are They Interchangeable?" In: *The International Journal of Logistics Management* 32.3 (2021), pp. 942–972. DOI: 10.1108/IJLM-05-2020-0214.

[75] Per Runeson and Martin Höst. "Guidelines for Conducting and Reporting Case Study Research in Software Engineering". In: *Empirical Software Engineering* 14.2 (2008), pp. 131–164. DOI: 10.1007/s10664-008-9102-8.

[76] Seeed Studio. *Getting Started with Seeed Studio XIAO ESP32C3*. 2026.

[77] Seeed Studio. *Seeed Studio XIAO ESP32-C3*. 2026.

[78] Shiptail. *Shiptail Pricing*. 2026.

[79] Sirirat Somapa, Martine Cools, and Wout Dullaert. "Characterizing Supply Chain Visibility – A Literature Review". In: *The International Journal of Logistics Management* 29.1 (2018), pp. 308–339. DOI: 10.1108/IJLM-06-2016-0150.

[80] Morgan Swink, Igor Sant'Ana Gallo, Cliff Defee, and Andrea Lago da Silva. "Supply Chain Visibility Types and Contextual Characteristics: A Literature-Based Synthesis". In: *Journal of Business Logistics* 45.1 (2024), e12366. DOI: 10.1111/jbl.12366.

[81] Transporeon. *Five Myths About Real-Time Visibility*. 2026.

[82] Robert Tröger, Rainer Alt, and Markus Zillmann. "Design Options for Supply Chain Visibility Services: Learnings from Three EPCIS Implementations". In: *Electronic Markets* 27 (2017), pp. 265–282. DOI: 10.1007/s12525-016-0231-4.

[83] u-blox AG. *GNSS antennas: RF design considerations for u-blox GNSS receivers*. 2019.

[84] U.S. Customs and Border Protection. *Electronic Refunds*. 2026. URL: https://www.govinfo.gov/content/pkg/FR-2026-01-02/pdf/2025-24171.pdf.

[85] Mbali Cynthia Valashiya and Rose Luke. "Enhancing Supply Chain Information Sharing with Third Party Logistics Service Providers". In: *The International Journal of Logistics Management* 34.6 (2022), pp. 1523–1542. DOI: 10.1108/IJLM-11-2021-0522.

[86] Ilias Vlachos. "Implementation of an Intelligent Supply Chain Control Tower: A Socio-Technical Systems Case Study". In: *Production Planning & Control* 34.15 (2021), pp. 1415–1431. DOI: 10.1080/09537287.2021.2015805.

[87] Samuel Fosso Wamba and Akemi Takeoka Chatfield. "A Contingency Model for Creating Value from RFID Supply Chain Network Projects in Logistics and Manufacturing Environments". In: *European Journal of Information Systems* 18.6 (2009), pp. 615–636. DOI: 10.1057/ejis.2009.44.

[88] Qiang Wang, Baofeng Huo, Xiande Zhao, and Zhaohui Hua. "What Makes Logistics Integration More Effective? Governance from Contractual and Relational Perspectives". In: *International Journal of Production Economics* 228 (2020), p. 107625. DOI: 10.1016/j.ijpe.2020.107625.

[89] Waveshare. *SIM7670G LTE Cat-1/GNSS HAT*. 2026.

[90] Slawomir Wycislak. "From Real-Time Visibility to Operational Benefits – Tensions on Unfinished Paths". In: *The International Journal of Logistics Management* 34.5 (2022), pp. 1446–1474. DOI: 10.1108/IJLM-03-2022-0126.

[91] Slawomir Wycislak and Atif Akhtar. “Real-Time Visibility as a Catalyst for Operational Enhancements”. In: *LogForum* 20.2 (2024), pp. 161–174. DOI: 10.17270/J.LOG.001010.

[92] Sławomir Wyciślak. “Exploring Real-Time Visibility Transportation Platform Deployment”. In: *LogForum* 18.1 (2022), pp. 109–121. DOI: 10.17270/J.LOG.2022.660.

[93] Sławomir Wyciślak. “Real Time Visibility in a Transportation Network of a Complex Supply Chain”. In: *International Journal of Supply Chain Management* 10.3 (2021), pp. 44–52.

[94] Sławomir Wyciślak and Pourya Pourhejazy. “Supply Chain Control Tower and the Adoption of Intelligent Dock Booking for Improving Efficiency”. In: *Frontiers in Energy Research* 11 (2023). DOI: 10.3389/fenrg.2023.1275070.

[95] Jian Yang, Christian Poellabauer, Pramita Mitra, and Cynthia Neubecker. “Beyond Beaconing: Emerging Applications and Challenges of BLE”. In: *CoRR* abs/1909.11737 (2019). DOI: 10.48550/arXiv.1909.11737.

[96] Yanzeo Smart Co., Ltd. *Yanzeo SR791 SR792 SR793 Long Distance Outdoor UHF RFID Reader*. 2026.

[97] Yanzeo Smart Co., Ltd. *Yanzeo SR791 UHF RFID Reader*. 2026.

[98] Xin Zhang, Xiaolong Guo, Wei Thoo Yue, and Yugang Yu. “Servitization for the Environment? The Impact of Data-Centric Product-Service Models”. In: *Journal of Management Information Systems* 39.4 (2022), pp. 1146–1183. DOI: 10.1080/07421222.2022.2127454.

Appendix A

Interview Attachments

Declaration of Authorship

‘I hereby declare,

- that I have written this thesis independently;
- that I have written the articles and contributions using only the aids listed in the index;
- that all parts of the thesis produced with the help of aids (incl. AI-Tools) have been precisely declared;
- that I have mentioned all sources used and cited them correctly according to established academic citation rules; this also includes the comprehensible disclosure of all personal publications;
- that I have acquired all immaterial rights to any materials I may have used, such as images or graphics, or that these materials were originally created by myself;
- that I am aware of the legal provisions regarding the publication and dissemination of parts or the entire thesis and that I comply with them accordingly;
- that I am aware that the University will prosecute a violation of this Declaration of Authorship and that disciplinary as well as criminal consequences may result, which may lead to expulsion from the University or to the withdrawal of my title.’

By submitting this thesis, I confirm through my conclusive action that I am submitting the statutory declaration, that I have read and understood it, and that it is true.

St. Gallen, May 19, 2026 Signature:

Declaration of Auxiliary Aids

The creation of this dissertation involved a lot of third-party software, including AI-based tools. I declare all uses of third-party software in the dissertation text that contributed non-trivially to the dissertation's content and are non-standard in such research. Further, by the University of St. Gallen's decree II.B.1.20 section 5.3, I explicitly declare the following uses:

Aid/tool	**Intended use**	**Affected parts**
OpenAI ChatGPT, Anthropic Claude	Spell and grammar checking; rewording; translations; support with research work.	Complete paper
Grammarly, LanguageTool, DeepL Write	Spell and grammar checking; rewording.	Complete paper
Research Rabbit	Source exploration.	Research work
Elevenlabs Scribe v2	Interview transcription.	Appendix
OpenAI Codex	Coding assistance during prototype/software development.	Prototype/software code
OpenAI ImageGen 2	Transformation of hand drawn explanatory figures into final figures.	Two Figures (cited)