Strategic positioning of MOOS in the retail technology ecosystem

An ecosystem-based diagnosis of actionable shelf insight and its role between shelf reality and store action

Master's Thesis · Strategic Product Design · Technical University Delft

Delft University of Technology

May 2026

Master's Thesis

MSc. Strategic Product Design

Faculty of Industrial Design Engineering

Author

Gijs de Meester - 4876237

Graduation committee

Chair (TU Delft) - Dr. S. (Shahrokh) Nikou PhD

Mentor (TU Delft) - N. Kyriakopoulos PhD

Company mentor (MOOS) - Dr. M. Reunis

Delft, the Netherlands

Preface

This thesis marks the final step of my MSc Strategic Product Design at Delft University of Technology. The project was carried out in collaboration with the company MOOS and focused on the strategic positioning of shelf-sensing technology in the retail technology ecosystem.

The project started from a practical question: how can a small shelf-level sensing tech start-up become relevant in a highly competitive and crowded landscape? During the research, this question developed into a broader ecosystem problem. Shelf insight only becomes valuable when it can be trusted, routed, owned, and acted on by the right actors. This made the project both analytical and strategic. It required moving between store-level and expert perspectives, academics, and design reasoning.

I would like to thank my graduation committee for their guidance throughout the project. Dr. Shahrokh Nikou and Nikolaos Kyriakopoulos helped me keep the academic contribution clear and grounded, while also supporting me throughout the process. I am also grateful to Dr. Marc Reunis and the rest of the MOOS team for their openness, feedback, and involvement during the project.

I also want to thank the store respondents, expert interview participants, and co-creation/validation participants who contributed their time and perspectives. Their input made it possible to move beyond a technology-centred view and understand how shelf sensing is evaluated in practice.

Finally, I want to thank my close relatives and friends. Their encouragement helped me complete a process that was at times demanding, but also highly valuable as a learning experience.

Gijs de Meester

Delft, May 2026

Note on the public repository version

This public repository version preserves the academic structure and main argument of the thesis, but some company-sensitive material has been removed, shortened, or generalised. This concerns customer names, proof cases, detailed roadmap content, specific partner and integration routes, supplier-specific competitor analysis, and confidential validation feedback. The full internal version was available for assessment by TU Delft and MOOS.

Abstract

This thesis examines how MOOS, a shelf-sensing component provider, can position itself in the retail technology ecosystem around actionable shelf insight. The study starts from the observation that shelf-level problems such as out-of-stocks and theft are not only visibility problems. Shelf information becomes valuable only when it is reliable, timely, connected to a workflow, owned by a decision-maker, and able to trigger store action.

The research uses ecosystem-as-structure as the main theoretical lens. This lens helps analyse the activities, actors, positions, links, and alignment conditions required to turn shelf-condition information into action. Co-innovation risk, adoption-chain risk, and internal readiness are used as supporting concepts.

The empirical research combines 41 store visits with a ranking exercise and eight semi-structured expert interviews. The findings show that store teams mainly need reliable signals, timely actionability, and low additional workload. At ecosystem level, shelf sensing depends on fragile links between store teams, IT, loss prevention, point-of-sale systems, workflow systems, partners, and ownership structures. Theft-related use cases are mainly shaped by urgency and response risk, while on-shelf availability use cases depend more strongly on integration and backend closure.

The design phase translates this diagnosis into a positioning direction and roadmap. The thesis concludes that MOOS can most credibly be understood as a bounded, partner-ready shelf intelligence and trigger layer between shelf reality and store action. The revised roadmap should be used as a strategic decision-support tool, not as a fixed rollout plan or proof of market success.

Keywords: actionable shelf insight, retail technology, ecosystem-as-structure, shelf sensing, strategic positioning, roadmap, MOOS

List of abbreviations

AI
Artificial intelligence

CCTV
Closed-circuit television

DEPEST
Demographic, Economic, Political, Ecological, Social, and Technological analysis

ERP
Enterprise resource planning

ESL
Electronic shelf label

IT
Information technology

LP
Loss prevention

OOS
Out-of-stock

OSA
On-shelf availability

POS
Point of sale

RFID
Radio-frequency identification

RQ
Research question

SQ
Subquestion

Executive summary

Retailers increasingly need better information about what happens at the shelf. Out-of-stocks, theft, shrinkage, and store execution problems show that the physical shelf remains a weak point in retail operations. Many systems know what should be in store, but not always what is actually happening at the shelf. MOOS addresses this gap with shelf-sensing technology that can detect shelf events and turn them into usable signals.

The central problem in this thesis is that shelf visibility does not automatically create value. A shelf signal only becomes useful when it can be trusted, routed to the right actor or system, connected to a workflow, owned by a decision-maker, and followed by store action. This makes MOOS's positioning problem broader than a product or competitor question. MOOS must position itself in an ecosystem of retailers, store teams, IT departments, loss-prevention teams, point-of-sale systems, workflow tools, enterprise systems, and technology partners.

The main research question is:

How can MOOS, as a shelf-sensing component provider, position itself in the retail-technology ecosystem around actionable shelf insight?

The thesis uses ecosystem-as-structure as the main theoretical lens. This lens defines an ecosystem around a focal value proposition and examines the activities, actors, positions, links, and alignment conditions needed to realise that proposition. In this thesis, the focal value proposition is actionable shelf insight. Co-innovation risk, adoption-chain risk, and internal readiness are used as supporting concepts to assess where value creation may break down and what MOOS must be able to execute internally.

The empirical research consists of two parts. Part 1 includes 41 store visits with a structured ranking exercise and short contextual conversations. This part identifies store-level adoption thresholds. Part 2 consists of eight semi-structured expert interviews with actors across the retail and retail technology ecosystem. This part examines ecosystem-level dependencies, ownership issues, integration conditions, and adoption risks. The two parts were analysed separately first and then synthesised through the conceptual framework.

The store-level findings show that store teams mainly care about three conditions: reliable signals, fast enough actionability, and low additional workload. Respondents did not evaluate shelf sensing as an abstract technology. They evaluated it through daily store work. A system that produces many false alerts, arrives too late, or adds another app or process is unlikely to be accepted, even if the sensing technology itself is promising.

The ecosystem-level findings show that the path from shelf signal to store action is fragile. Several actors may control different parts of the action chain. Store teams may feel the operational pain, but IT, loss prevention, operations, data teams, or external vendors may control the systems needed to act on the signal. This means that responsibility is spread across several actors. The main challenge is not only whether MOOS can detect shelf events, but whether those events can be connected to the right routines, systems, and decision owners.

The synthesis shows an important difference between theft-related use cases and on-shelf availability use cases. Theft-related use cases are mainly shaped by urgency, routing, response risk, and safe escalation. They may create value when a trusted shelf event can trigger a timely and targeted response. On-shelf availability use cases depend more strongly on integration with replenishment, stock data, task ownership, and backend correction. This makes availability more dependent on wider system alignment.

The design phase translates these findings into positioning options and roadmap logic. The preferred positioning direction is that MOOS should become a bounded, partner-ready shelf intelligence and trigger layer between shelf reality and store action. Bounded means that MOOS should not try to become the full retail operating system. Partner-ready means that MOOS should be able to connect selectively to adjacent systems and partners. Shelf intelligence and trigger layer means that MOOS should focus on reliable shelf events, interpretation, and action routing.

The initial roadmap was developed as a pre-validation design output. It was then validated for clarity, relevance, credibility, and usefulness. The validation did not prove market success, technical performance, commercial viability, partner willingness, or long-term operational effectiveness. It tested whether the roadmap logic was understandable and useful as a strategic decision-support tool.

The validation led to a clearer revised roadmap. The main shift was from a time- and segment-led expansion path toward a maturity-based decision logic. MOOS should start from its existing evidence base, understand which conditions make this evidence transferable, define repeatable traction patterns, and then move toward adjacent propositions only when comparable conditions are present. Deeper ecosystem embedding should come later and selectively, when evidence, internal readiness, and dependency load justify it.

The practical implication is that broad expansion across retail contexts should be assessed carefully rather than assumed as the next step. It should use the roadmap as a decision filter. New opportunities should be assessed through lookalike conditions, use-case fit, integration burden, partner dependency, internal readiness, and evidence strength. Theft and on-shelf availability should remain on different scaling paths because they depend on different action chains. Partners should be assessed not only by commercial access, but also by the dependency load they create.

The thesis concludes that MOOS's strongest position is not as a stand-alone sensor provider and not as a full store operating system. Its most credible role is between those extremes: a bounded shelf intelligence and trigger layer that makes shelf reality usable for store action. This position is promising, but it requires disciplined choices about proof, repeatability, partnerships, integration, and internal readiness.

Table of contents

1. Introduction 14
1.1 Background and motivation 16
1.2 Problem definition and research gap 16
1.3 Research objective and research questions 17
1.4 Scope and boundaries 18
1.5 Theoretical lens and research approach 18
1.6 Expected outputs 19
1.7 Thesis structure 19
2. Ecosystem Context Analysis 22
2.1 Internal context: MOOS as the focal case 24
2.2 Retail and retail technology environment 25
2.3 Store technology trajectories and scaling constraints 28
2.4 Retail technology actor categories and control points 30
2.5 Alternative solution routes around shelf-level visibility 34
3. Literature Review 40
3.1 Theory background 42
3.2 Core building blocks from the “ecosystem-as-structure” and related literature 43
3.3 Operationalising the lens: risks, bottlenecks, and internal readiness 45
3.4 Integrated conceptual framework 47
4. Methodology 52
4.1 Case context and use-case focus 54
4.2 Research process overview 54
4.3 Data collection Part 1: structured in-store visits with ranking exercise 55
4.4 Data collection Part 2: semi-structured expert interviews 56
4.5 Ethics, consent, and confidentiality 57
4.6 Data preparation 57
4.7 Data analysis 58
4.8 Synthesis of Part 1 and Part 2 59
4.9 AI-assisted analysis support 59
4.10 Co-creation as part of the design process 59
4.11 Validation logic 60
4.12 Methodological limitations 60
5. Results 62
5.1 Store level findings 64
5.2 Ecosystem-level findings 68
5.3 Cross-dataset synthesis 72
5.4 Framework refinement 75
5.4.1 What became more specific through the findings 76
..... 77
..... 77
6. Design 80
6.1 From diagnosis to design 82
6.2 Design process and role of co-creation 84
6.3 Co-creation outputs as design input 84
6.4 Design criteria 84
6.5 Future vision 85
6.6 Ideation and development of positioning directions 86
6.7 Preferred direction 86
6.8 Public-level roadmap 86
6.9 Tactical level workstream 87
6.10 Conclusion 88
7. Validation & roadmap revision 90
7.1 Object of validation 92
7.2 Participants and procedure 92
7.3 Validation criteria and analysis method 93
7.4 Validation findings 93
7.5 Roadmap revisions after validation 94
7.6 Final revised roadmap logic 95
7.7 Limitations of validation 98

8. Discussion 100
8.1 Theoretical contribution 102
8.2 Methodological and design contribution 102
8.3 Practical implications for MOOS 103
8.4 How the roadmap should be used 103
8.5 Broader practical relevance 103
8.7 Future research 104
8.6 Limitations 104
8.8 Closing reflection 105

9. Conclusion 106
9.1 Answer to the main research question 108
9.2 Answer to SQ1: structure and actionability 108
9.3 Answer to SQ2: dependencies and alignment 108
9.4 Answer to SQ3: risks and feasibility 109
9.5 Answer to SQ4: points of influence and readiness 109
9.6 Final remarks 109

Appendix 110
Appendix A. Official project documentation 112
A.1 Approved project brief 112
A.2 Human Research Ethics Committee approval 118
Appendix B. Supplementary Context Analysis 120
B.1 Full Depest 121
..... 121
C.Store Visit Research Material 122
C1. Store Visit Protocol 122
C2. Store Visit Question list 124
D. Expert Interview Research Material 128
D.1 Interview Protocol 128

List of Figures

Figure	Title
Figure 1.1	Research logic
Figure 1.2	Report structure
Figure 2.1	MOOS in store operations
Figure 2.2	DEPEST overview
Figure 2.3	Store operations pressure map
Figure 2.4	Actor categories and control points in the retail technology environment
Figure 2.5	Data flow and integration bottlenecks around shelf events
Figure 2.6	Solution routes around shelf-level visibility
Figure 3.1	Integrated diagnostic framework for actionable shelf insight
Figure 3.2	Focal value proposition
Figure 3.3	Basic activity chain
Figure 3.4	Alignment conditions
Figure 3.5	Ecosystem risks
Figure 3.6	Bottlenecks and leverage points
Figure 3.7	Internal readiness
Figure 3.8	Diagnostic outputs
Figure 4.1	Research process overview from framework development to conducted validation
Figure 5.1	Empirical analysis flow
Figure 5.2	Store adoption threshold model
Figure 5.3	Main store-level reasons for sensor interest based on open responses
Figure 5.4	Exploratory branch versus franchise comparison
Figure 5.5	Exploratory higher-tech versus lower-tech store comparison
Figure 5.6	Retail blindness and the split between OOS and theft pathways
Figure 5.7	Signal-to-action failure map
Figure 5.8	Fragmented buying centre around shelf sensing adoption
Figure 5.9	Two-layer integration view for shelf sensing
Figure 5.10	Cross-dataset synthesis chain
Figure 5.11	Refined diagnostic framework after the empirical findings
Figure 6.1	Design process from refined framework to roadmap logic
Figure 7.1	Public-level strategic maturity logic after validation

List of Tables

Table	Title
Table 2.1	Store technology trajectories and scaling constraints
Table 2.2	Representative actor groups and control point examples in the retail technology environment
Table 2.3	Category-level comparison of shelf-related solution routes
Table 4.1	Actor perspectives included and less visible in the expert interview sample
Table 5.1	Overview of the store visit sample
Table 5.2	Average ranking of store adoption attributes
Table 5.3	Store users and organisational gatekeepers compared on value, risk, evidence, and rejection
Table 6.1	Design brief overview
Table 6.2	Public-level design criteria
Table 6.3	Public-level future role boundary
Table 6.4	Structured ideation methods
Table 6.5	Public-level positioning elements
Table 6.6	Composition of the preferred positioning direction
Table 6.7	Public-level roadmap maturity logic
Table 6.8	Public-level tactical workstreams
Table 7.1	Validation criteria used to assess the pre-validation design output
Table 7.2	Main public-level roadmap revisions after validation
Table 7.3	Revised tactical workstream logic after validation
Table 7.4	Decision options for using the revised roadmap logic

Chapter 1

Introduction

1.1 Background and motivation
1.2 Problem definition and research gap
1.3 Research objective and research questions
1.4 Scope and boundaries
1.5 Theoretical lens and research approach
1.6 Expected outputs
1.7 Thesis structure

This chapter introduces the positioning problem addressed in this thesis. It explains why shelf-level information matters for retail execution, but also why visibility alone is not sufficient. Shelf information only becomes valuable when it is reliable, timely, connected to a workflow, owned by a decision-maker, and able to trigger accountable store action. The chapter then positions MOOS as the focal case for studying this problem, defines the research gap and research questions, and outlines the scope, theoretical orientation, expected outputs, and structure of the thesis.

1.1 Background and motivation

Retail and wholesale are undergoing structural change driven by digitalisation, sustainability requirements, and increasing pressure on labour and margins (McKinsey & EuroCommerce, 2022). As retailers integrate online and offline channels, in-store execution becomes more important because pricing, replenishment, fulfilment, and availability depend on timely and usable information from the shop floor (Deloitte, 2025).

Two persistent operational problems illustrate why store-level information remains economically material. First, out-of-stock (OOS) situations continue to cause lost sales and customer dissatisfaction despite decades of process optimisation and information system investment (ECR Europe, 2012). Inventory inaccuracy is a structural driver of OOS situations because system records frequently diverge from physical shelf reality, which makes execution decisions less reliable at store-level (Rekik et al., 2019). Second, shrinkage imposes high financial and organisational costs on retailers and also affects employee safety and daily store operations (National Retail Federation, 2023).

Both issues are linked to gaps between what systems assume and what actually occurs at the shelf in stores. Greater visibility into shelf conditions can therefore support better execution and decision-making, but only when that visibility can be translated into timely and accountable store action (ECR Europe, 2012).

Within this context, the case company, MOOS, develops shelf-sensing technology that produces real-time insights into shelf conditions using AI-based sensors (MOOS, 2025). The company frames its value proposition around making shelf events visible and actionable for store operations, with the aim of supporting availability execution and reducing theft-related losses (MOOS, 2025).

However, the value of shelf-sensing does not emerge from the component alone. A shelf event must be reliable enough to trust, timely enough to matter, routed into a workflow, connected to a decision owner, and followed by accountable store action. Impact at scale therefore depends on integration with enterprise systems, alignment with store processes, and coordination with partners across the retail technology landscape (EY, 2023). This creates a strategic positioning challenge that goes beyond product performance.

MOOS therefore functions as the focal case through which this thesis examines how a shelf-sensing component provider can seek a viable position when value creation depends on actors, interfaces, and routines beyond its direct control.

1.2 Problem definition and research gap

Although retailers deploy a growing number of in-store technologies, these initiatives are often implemented as isolated solutions rather than as integrated system changes, which can create silos across data, processes, and organisational ownership (Verhoef et al., 2021). This can lead to fragmented technology landscapes in which interoperability and data integration remain difficult to achieve across channels and store systems (Mirzabeiki & Saghiri, 2020). This fragmentation matters because shelf-level problems such as out-of-stocks and shrinkage are not solved by visibility alone. Shelf information becomes valuable only when it can move across systems, roles, and routines in a way that supports store action.

In such settings, control over interfaces, data access, and decision rights can become contested because value creation and value capture depend on how governance is organised across interdependent actors (Jacobides et al., 2018). For a shelf-sensing component provider such as MOOS, this creates a specific positioning problem. The component may detect shelf events, but the value of those events depends on whether other actors can integrate, interpret, route, and act on the signal. Strategic positioning therefore cannot be assessed only by asking whether the technology works. It must also be assessed by asking which actors, links, and routines are needed to turn shelf information into accountable execution.

Research on innovation increasingly emphasises that outcomes depend not only on technological capabilities, but also on how firms align their roles and activities within ecosystems (Adner, 2017). An

ecosystem-as-structure perspective is useful here because it defines an ecosystem around a focal value proposition and examines the activities, actor positions, and dependency links required to realise that proposition (Adner, 2017). However, existing retail technology and operations literature does not fully explain how cross-actor dependencies, ecosystem coordination, and link-level alignment shape the feasibility of actionable shelf insight in store execution settings.

Research on inventory accuracy and on-shelf availability (OSA) often examines specific mechanisms or single technologies, such as inventory record inaccuracy in analytical supply chain models (Fleisch & Tellkamp, 2005) or the effects of RFID-enabled visibility on OOS events (Hardgrave et al., 2013). Retail technology research also often frames the issue as an adoption or acceptance decision, for example through organisational acceptance constructs (Venkatesh et al., 2003) or retailer adoption conditions shaped by shopper attitudes and privacy concerns (Inman & Nikolova, 2017). Studies of omnichannel execution explain process realignment within the retailer and its supply chain (Ishfaq et al., 2022), while digital transformation scholarship explains constraints related to legacy systems and organisational change (Vial, 2019). Together, these streams explain important parts of the problem, but their dominant units of analysis are different from the problem addressed in this thesis: the inventory record, the technology, the adopter, the retailer process, or the transformation programme. They do not sufficiently explain how shelf insight becomes feasible when value depends on links between multiple actors that control different systems, interfaces, routines, and decision rights. This leaves a gap in explaining why a technically promising shelf-sensing solution may still fail to create value when it cannot be aligned across the ecosystem required for store execution.

This thesis contributes by operationalising ecosystem-as-structure for a component provider in a retail execution setting. The contribution is not a new ecosystem theory, and it is not limited to MOOS as a single case. It lies in using ecosystem-as-structure as a diagnostic approach that makes activities, actor positions, dependency links, alignment conditions, and risks observable around a focal value proposition. This allows the feasibility of actionable shelf insight to be assessed before it is translated into positioning options and roadmap logic.

In this thesis, actionable shelf insight means shelf-condition information that can trigger a relevant store decision or action. It is actionable when the signal can be trusted, arrives in time, fits a workflow, reaches a clear decision owner, and does not create unacceptable extra workload. This definition keeps the analysis focused on operational value rather than on sensing outputs in isolation. It is motivated by evidence that gaps between system records and physical shelf reality create real execution losses and require better detection and response mechanisms (Rekik et al., 2019).

1.3 Research objective and research questions

The objective of this project is to diagnose the ecosystem conditions that shape how a shelf-sensing component provider can position itself around actionable shelf insight, and to translate that diagnosis into positioning and roadmap logic, using MOOS as the case company.

In this thesis, actionable shelf insight is examined across two related use-case variants that MOOS aims to support. The first variant concerns availability execution and OOS reduction. The second variant concerns the execution of loss prevention (LP) and the detection of theft-related events. The ecosystem configuration and feasible positioning options may differ across these variants because data access, governance constraints, and operational ownership can vary by use case.

Therefore, the central research question is*:*

'How can MOOS, as a shelf-sensing component provider, position itself in the retail-technology ecosystem around actionable shelf insight?'

This question is addressed through four thematic subquestions that move from ecosystem structure, to dependencies and alignment, to risk and feasibility, and finally to leverage and readiness.

SQ1. **Structure and actionability** - *What activity structure, actor clusters, and actionability conditions are required to turn shelf-condition information into store action?*

SQ2. **Dependencies and alignment** - *Where do key dependencies, bottlenecks, and misalignments occur across ecosystem positions and links?*

SQ3. **Risks and feasibility** - *How do co-innovation risk and adoption-chain risk affect the feasibility of*

different positioning options for a shelf-sensing component provider?

SQ4. **Leverage and readiness** - *Which leverage points and internal readiness conditions should guide MOOS in selecting and sustaining a viable ecosystem position?*

1.4 Scope and boundaries

The focal value proposition of this study is actionable shelf insight, defined as shelf-condition information that can trigger relevant store decisions or actions through reliable, timely, integrated, and accountable execution. This scope keeps the analysis centred on the ecosystem required to realise this value proposition, rather than on the full retail-technology market or on MOOS as a firm in isolation.

The unit of analysis is the ecosystem-alignment structure around this value proposition, rather than a single firm or a single technology in isolation. This means that the analysis focuses on the conditions under which actionable shelf insight can be realised across actors, systems, and routines, rather than on evaluating MOOS' product performance in isolation. Adoption is included only insofar as it manifests as staged, multi-gate adoption-chain risk within the ecosystem-alignment structure, and not as a stand-alone consumer or user-adoption problem.

The empirical setting focuses on the retail-technology context described in the context chapter, with emphasis on enterprise retailers and their partner landscapes in Europe. The thesis does not aim to provide a complete go-to-market strategy, pricing model, or customer-adoption plan, except where such elements are necessary to analyse incentives, governance, and alignment. Observations from retail-technology trade fairs are used solely as contextual input to support understanding of actor categories and recurring problem framings and are not treated as formal empirical data because of the absence of systematic data collection.

1.5 Theoretical lens and research approach

The core theoretical lens of this study is ecosystem-as-structure. It is used to examine how activities, actors, positions, and links must fit together around a focal value proposition (Adner, 2017). In this thesis, that value proposition is actionable shelf insight. The lens is used to study why shelf sensing may fail to create value when the links between actors, systems, and routines are not aligned.

Two supporting lenses are used in a limited way. Co-innovation risk and adoption-chain risk help identify where the path from pilot to wider use may break when several actors must act together (Adner & Kapoor, 2010). Dynamic capabilities are used as an internal readiness layer to assess whether MOOS can execute and adapt a chosen position as conditions change (Teece, 2007). In short, ecosystem-as-structure maps the external alignment problem, risk concepts show where delivery may break down, and internal readiness assesses whether a position is feasible for MOOS.

The research approach starts with a focused literature review, which is used to develop the conceptual framework in Chapter 3. That framework informed the empirical design, including the store visit protocol, interview protocol, and analytical coding scheme.

Empirical data were collected in two parts. The first part consisted of structured store visits with a ranking exercise and short contextual conversations. These store visits were used to identify store-level adoption thresholds and practical actionability conditions. The second part consisted of semi-structured expert interviews with selected actors across relevant ecosystem clusters. These interviews were used to examine links, alignment failures, governance mechanisms, bottlenecks, and adoption gates. The combined diagnosis then became the basis for positioning options and roadmap logic.

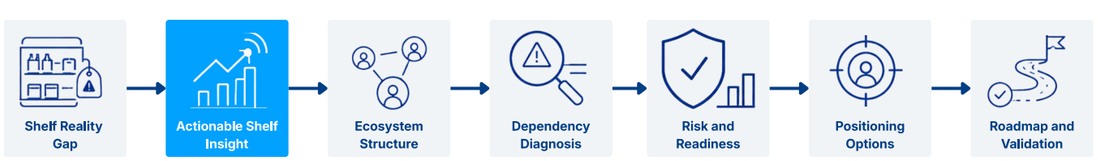

Figure 1.1. Research logic.

Figure 1.1 provides a simplified preview of the research logic. It shows that actionable shelf insight defines the ecosystem boundary, while ecosystem structure and dependency diagnosis make the positioning problem visible. Positioning is treated as an output of ecosystem diagnosis, not as a starting assumption.

The figure shows how the thesis moves from actionable shelf insight to ecosystem structure, dependency diagnosis, risk and readiness assessment, positioning options, roadmap logic, and validation. It does not show the full conceptual framework, which is developed in Chapter 3.

1.6 Expected outputs

The expected output of this thesis is a strategic design synthesis rather than a product design, technical specification, or go-to-market plan.

The academic output is an operationalisation of ecosystem-as-structure for diagnosing how a component provider can position itself around a focal value proposition in a retail execution setting. This operationalisation connects ecosystem structure, dependency links, alignment conditions, risk diagnosis, and internal readiness in one diagnostic logic.

The empirical output is a structured diagnosis of the actionable shelf insight ecosystem. It identifies bottlenecks, alignment failures, risk patterns, governance issues, and differences between availability and loss-prevention use cases.

The design output translates this diagnosis into positioning options and roadmap logic, using MOOS as the case company.

The validation output assesses the clarity, relevance, credibility, and usefulness of the roadmap and supports a final revision of the strategic direction. The roadmap is not presented as a rollout plan or as evidence of market success. It is a decision-support tool for strategic positioning.

1.7 Thesis structure

The remainder of the thesis follows the logic from problem framing to context, framework, method, results, design, validation and roadmap revision, discussion, and final conclusion.

Figure 1.2 summarises the chapter flow. Chapter 2 grounds the positioning problem in MOOS and the retail-technology environment. Chapter 3 develops the analytical lens, Chapter 4 explains the method, Chapter 5 presents the empirical diagnosis, Chapter 6 translates that diagnosis into a pre-validation roadmap, and Chapter 7 validates and revises the roadmap. Chapter 8 discusses the meaning and limits of the work, and Chapter 9 gives the final answer to the research question.

Figure 1.2 shows how the report moves from problem framing and context to framework, method, empirical diagnosis, design translation, conducted validation and roadmap revision, discussion, and final conclusion. The thesis proceeds as follows:

Chapter 1 introduces the research context, research problem, research questions, scope, theoretical orientation, expected outputs, and thesis structure.

Chapter 2 presents the ecosystem context analysis. It introduces MOOS as the focal case and explains the retail technology setting, store technology trajectories, actor categories, control points, and alternative solution categories around shelf-level visibility.

Chapter 3 reviews the academic literature and develops the conceptual framework. It explains ecosystem as structure as the core lens and shows how risk diagnosis and internal readiness support the empirical analysis.

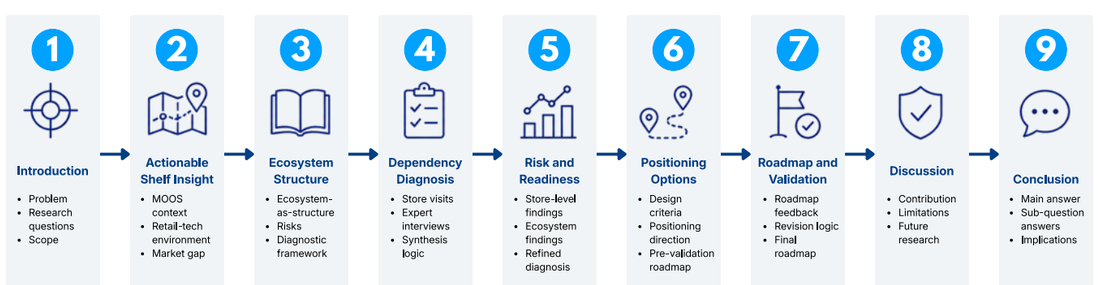

Figure 1.2. Report structure.

Chapter 4 outlines the research methodology. It explains the case context, use case focus, store visits, expert interviews, ethics, data preparation, and analytical procedure.

Chapter 5 presents the empirical results in the same order as the research design. It reports the store level findings, then the ecosystem level findings, and finally synthesises both datasets into a diagnosis of dependencies, alignment failures, risk patterns, governance gaps, and framework refinements.

Chapter 6 develops the design phase. It translates the empirical diagnosis and co creation input into design criteria, a preferred positioning direction, and a pre validation strategic and tactical roadmap for MOOS.

Chapter 7 presents the conducted validation and roadmap revision. It assesses whether the pre validation roadmap is understandable, relevant, credible, and useful for MOOS and relevant stakeholders, and explains how validation findings informed the final revised roadmap.

Chapter 8 discusses the meaning of the findings. It reflects on the conceptual contribution, methodological and design contribution, practical implications for MOOS, limitations, and future research.

Chapter 9 concludes the thesis. It answers the main research question and subquestions, then closes with the final positioning answer and the main implication for MOOS and for ecosystem based positioning in retail execution.

Chapter 2

Ecosystem Context Analysis

2.1 Internal context: MOOS as the focal case
2.2 Retail and retail technology environment
2.3 Store technology trajectories and scaling constraints
2.4 Retail technology actor categories and control points
2.5 Alternatives and competitors around shelf-level visibility

This chapter establishes the market and ecosystem context for examining how MOOS can position itself within retail technology systems. The chapter does not yet present the formal theoretical framework. Instead, it explains why MOOS' positioning problem cannot be understood as a single product, buyer, or direct-substitute problem.

Retail and wholesale are undergoing structural change driven by digitalisation, sustainability pressure, and skills constraints, which reshape how stores are designed, operated, and managed within omnichannel systems (McKinsey & EuroCommerce, 2022). Consumer expectations also push retailers to treat technology as part of the core store experience rather than as a set of supporting tools (Dhingra et al., 2023). These changes matter for MOOS because shelf sensing only becomes valuable when it connects shelf reality to system data, workflows, decision ownership, and accountable follow-up action.

The analysis is based on desk research using published sector reports, industry material, and organisational documents. It moves from MOOS as the focal case to the wider retail technology environment, store technology trajectories, actor control points, and alternative routes to shelf-level visibility. Across these sections, one central point is developed: MOOS operates in a retail technology ecosystem in which shelf reality, store systems, actor control points, and workflow ownership do not automatically align. This creates the need for the ecosystem lens developed in Chapter 3.

2.1 Internal context: MOOS as the focal case

MOOS is used in this thesis as the focal case of a company developing shelf-level sensing for physical retail environments. Its relevance does not lie only in the sensing technology itself, but in the fact that it addresses operational problems that remain structurally difficult to solve in stores, namely OSA, shrink, and execution workload (MOOS, 2025). In that sense, MOOS is relevant not simply as a technology supplier, but as a case through which the relation between shelf events, store action, and ecosystem dependencies can be examined.

This distinction is important because the same sensing capability may enter different operational chains depending on the use case. Prevention-oriented applications require fast signals, clear response routines, and ownership over immediate action. Diagnosis-oriented applications require evidence quality, comparison with other data sources, and integration with review or investigation routines. The use-case split therefore already points to a positioning problem rather than a single product logic.

MOOS is positioned around persistent shelf-level problems, such as OOS events, theft-related losses, and inefficient execution on the shop floor (MOOS, 2025). Rather than presenting value through one single performance indicator, the offering is framed as a set of operational improvements that depend on making shelf events visible, interpretable, and actionable within existing store routines and decision moments.

The internal framing of the offering suggests that shelf sensing can support both immediate operational follow-up and later diagnostic interpretation. In both cases, the value does not come from sensing alone. It depends on whether shelf-event information can be trusted, interpreted, routed to the right actor or system, and connected to existing store routines. This means that MOOS is best understood as a contributing layer within broader retail operations, rather than as a stand-alone system that owns the full workflow.

Two characteristics follow from this internal problem framing. First, the product sits at the boundary between hardware deployment in stores and software integration with operational and analytic systems, which makes scaling dependent on partners, retailer IT maturity, and process change. This dependence is consistent with broader sector observations that retailers need to overhaul architecture and operating models to exploit technology investments fully (Bick et al., 2022). Second, the primary user group is not only the shopper, but also store operations teams, loss-prevention teams, and retail IT. This means that the value of shelf sensing is shaped by cross-functional incentives and by the ability to integrate new data streams into existing decision routines (Dhingra et al., 2023). For MOOS, the strategic question is therefore not only whether it can generate accurate shelf signals. The question is where those signals should connect, who should act on them, and which partners or systems are needed to make the signal operationally useful at scale.

Figure 2.1 shows the path from a physical shelf event to interpretation, routing, and store action. It highlights that MOOS' value does not depend only on sensing accuracy, but also on the systems and routines that receive, interpret, and act on the signal.

2.2 Retail and retail technology environment

Retail and wholesale in Europe face transformation pressures that combine digitalisation, sustainability, and workforce constraints, which require both investments in new channels and reinvention of the physical store role (McKinsey & EuroCommerce, 2022). In this setting, technology adoption tends to create value only when it is embedded in operating models, governance arrangements, and routines that make new data usable in execution (Vial, 2019). Omnichannel operations further increase interdependence because store-level processes depend on shared data infrastructures and cross-functional coordination across fulfilment, pricing, and replenishment activities (Ishfaq et al., 2022). These pressures matter for MOOS because its product sits inside store operations, where integration quality, execution readiness, and ownership of follow-up actions influence whether shelf-level sensing can become a scalable position rather than a pilot-level technology.

This shift increases the operational demands placed on the store. Digitalisation expands the number of processes that require reliable data, such as pricing integrity, fulfilment, replenishment, and availability management, which are hard to execute when store-level information is incomplete or delayed (McKinsey & EuroCommerce, 2022). At the same time, retailers are expected to reduce friction and keep store processes consistent across channels (Dhingra et al., 2023). This makes shelf-level execution a more visible constraint. Availability, shrink, and workload are not only store problems. They are also data and coordination problems. For MOOS, this means that shelf sensing becomes strategically relevant only when the signal fits into an operational rhythm: who receives it, when it is received, which action follows, and whether that action can be repeated without adding too much workload.

Two store-level problems are particularly relevant for MOOS: OOS events and shrinkage. A sector perspective on availability notes that it is not unusual for a supermarket to miss out on around 5 percent of sales through OOS events, which frames availability as a recurring economic loss rather than a marginal operational issue (Oliver Wyman, 2012). The same source links availability performance to replenishment routines, forecasting, and in-store processes (Oliver Wyman, 2012). Better sensing is therefore not enough on its own. The shelf signal must connect to a decision point, an accountable actor, and a practical follow-up action. This fits the broader argument that retail technology creates value when it is integrated into operating model change, rather than treated as a standalone upgrade (Bick et al., 2022).

The implication for MOOS is that availability is not only a detection problem. A shelf signal only becomes valuable when it leads to replenishment ownership, task prioritisation, or another clear action. This makes OOS a difficult positioning case. MOOS may win where retailers need faster shelf-level evidence, but it may struggle where the main problem sits in forecasting, replenishment logic, backend stock correction, or unclear store ownership.

Shrink is a second structural driver because it combines economic loss with safety and

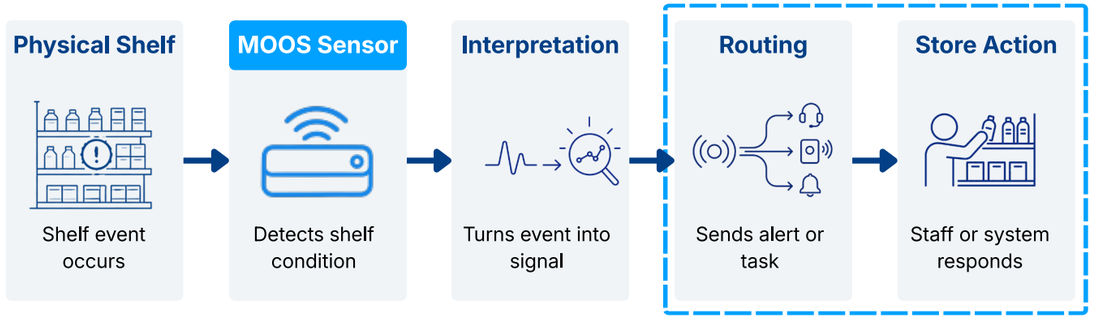

Figure 2.1. MOOS in store operations.

operational burden. The National Retail Security Survey reports that shrinkage as a percentage of sales increased from 1.4 percent to 1.6 percent, and it estimates total retail shrinkage at $112.1 billion, signalling a level of loss that is material for both margin and managerial attention (National Retail Federation, 2023). The same report highlights that organised retail crime and violence have become central concerns for retail leaders, suggesting that loss-prevention investments are shaped not only by financial calculations but also by safety and reputational considerations (National Retail Federation, 2023).

For a shelf-sensing offer, this matters because theft-reduction use cases are likely to be assessed within broader security strategies and evidence routines. MOOS company material describes links between shelf events, POS data, and video review, which suggests that the product is not positioned as a replacement for camera-based security, but as a signal that can make review and response more targeted (MOOS, 2025). This creates a different positioning logic from OOS. In LP, MOOS may be strongest where the retailer needs a timely shelf-level trigger and supporting evidence. It may be weaker where the main need is broad visual surveillance, deterrence, staff safety infrastructure, or investigation after the event.

Cost and capability constraints also shape the environment. The EU sector analysis emphasises that upgrading e-commerce, expanding supply chain networks, and digitally enabling stores require significant investment and organisational capacity, which implies that retailers will prioritise technologies that fit into coherent roadmaps rather than adding isolated tools (McKinsey & EuroCommerce, 2022). A related point in retail technology transformation work is that retailers often need to modernise their technology architecture and reorganise IT to deliver faster time-to-market and manage technological complexity (Bick et al., 2022). These constraints are consistent with a pattern in which operational technologies must show clear links to performance outcomes and be implementable within limited transformation bandwidth. For MOOS, this creates a practical constraint on positioning. A technically strong shelf-sensing proposition may still struggle if it requires too much integration effort, unclear ownership, or a business case that competes with higher-priority transformation initiatives. The positioning problem is therefore partly about reducing perceived implementation burden while making the operational value specific enough to justify attention and investment.

A DEPEST view helps to condense these pressures into a smaller set of contextual forces that shape store technology decisions. The purpose here is not to present a full trend study, but to identify which external forces increase the relevance of shelf-level visibility and which forces constrain technology adoption.

Demographic and social factors include ageing populations, smaller households, time pressure, and expectations of convenient, low-friction shopping. These forces increase pressure on stores to remain easy to operate and reliable for shoppers, while labour availability becomes more constrained (McKinsey & EuroCommerce, 2022).

Economic factors include margin pressure, cost inflation, value-seeking consumer behaviour, and the need to improve productivity. In this context, technologies are likely to be evaluated against clear operational benefits, short payback expectations, and their ability to reduce manual workload or prevent avoidable losses (Deloitte, 2025).

Political and legal factors are visible in the increased attention to retail crime, employee safety, data protection, and traceability requirements. These forces can support investment in loss-prevention and tracking technologies, but they also make privacy, evidence quality, and data governance more sensitive evaluation criteria (National Retail Federation, 2023).

Ecological factors include waste reduction, supply chain transparency, and the sustainability transition of retail operations. These forces make better inventory visibility relevant because inaccurate or delayed store information can contribute to waste, inefficient replenishment, and avoidable operational losses (McKinsey & EuroCommerce, 2022).

Technological factors include the growing maturity of AI, cloud infrastructure, edge computing, connectivity, and in-store automation. These technologies make real-time sensing and decision support more feasible, but they also increase the need for coherent data architectures, cybersecurity, and integration standards (McKinsey & Company, 2025).

Taken together, the DEPEST forces suggest that the opportunity for MOOS is created by pressure on availability, labour, shrink, waste, and data-driven execution. However, these same forces also create constraints. Retailers are unlikely to adopt another point solution unless it fits existing systems, supports clear operational ownership, and can be justified within broader transformation priorities. The implication is that MOOS should not be positioned only as a more accurate sensor. Its relevance depends on whether shelf-level data can reduce execution uncertainty in situations where existing store data, staff routines, or security systems are insufficient. The full DEPEST analysis is included in Appendix B.1.

Figure 2.2 summarises the DEPEST forces that make shelf-level visibility more relevant for retailers. The figure should be read as a context filter: it translates broad external pressures into store-level conditions that matter for shelf sensing.

Figure 2.3 on the next page then connects these external pressures to the operational pain points around availability, labour, cost, shrink, security, and data fragmentation. The figure shows why shelf technology becomes relevant when it reduces store burden and turns shelf events into action.

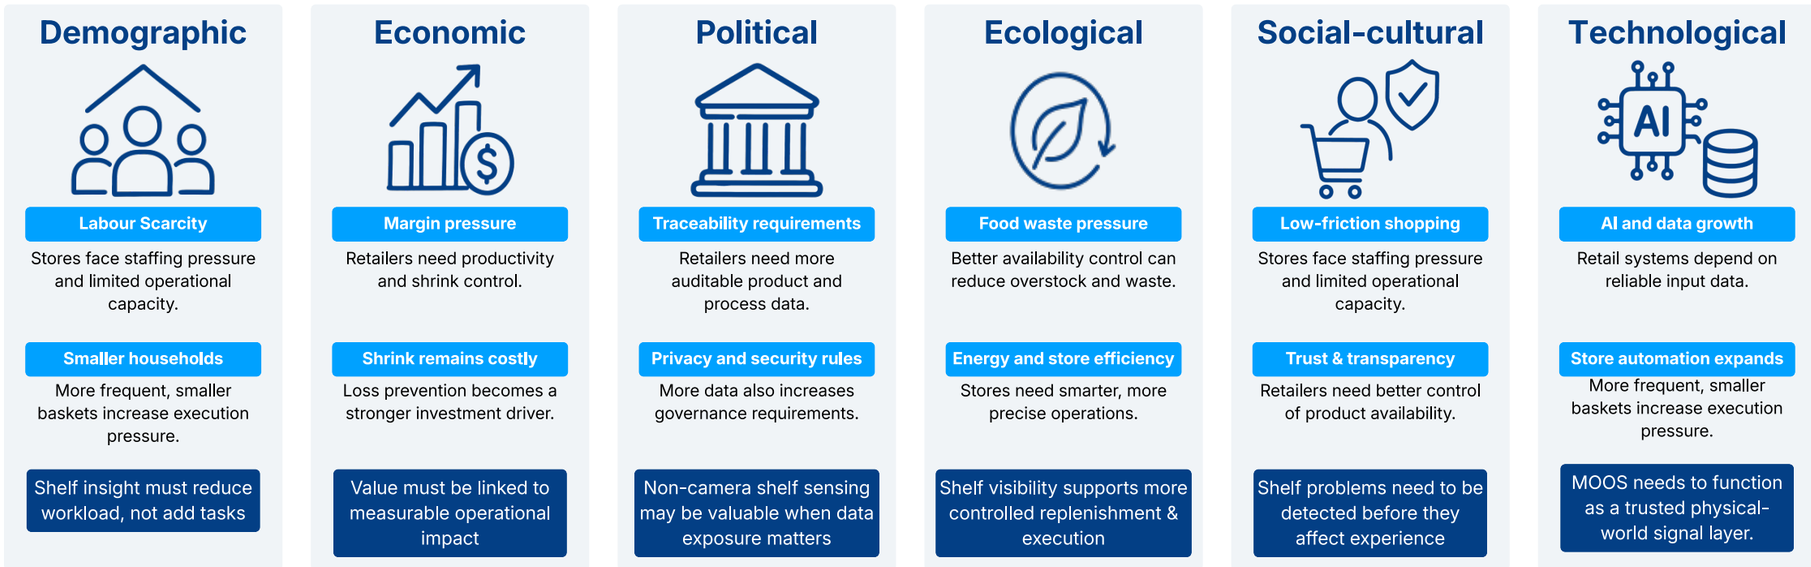

Figure 2.2. DEPEST overview.

2.3 Store technology trajectories and scaling constraints

Store technologies are discussed here as contextual evidence for how in-store technologies move from pilots to scalable use, not as a separate adoption study. Retail technology in stores often develops through cycles of experimentation, infrastructure build-up, and gradual integration into broader operating models. This pattern is consistent with research on digital transformation, which suggests that technologies create value only when they are embedded in routines, governance, and architecture rather than treated as isolated upgrades (Vial, 2019). For MOOS, these trajectories are relevant because they show that technical performance is only one part of scaling. The more strategic question is whether a technology can become part of the store's existing systems, roles, and decision routines.

As a visible example, self-checkout has scaled widely over several decades, yet it continues to generate operational frictions and follow-on innovation waves such as checkout-less concepts (EY, 2023). Self-checkout illustrates that scaling a store technology can shift rather than remove operational problems. A checkout intervention may reduce labour at the transaction point, but it can also increase the need for exception handling, loss-prevention routines, and monitoring. For MOOS, this suggests that a shelf-sensing proposition should not only claim efficiency gains, but also clarify how new signals will be governed, acted upon, and connected to existing control processes.

The store technology base is also broadening beyond checkout. In the European retail context, sector evidence points to technologies such as RFID, smart carts, and traffic sensors as building blocks for more instrumented store workflows (McKinsey & EuroCommerce, 2022). These technologies do not all solve the same problem. RFID strengthens item identification and inventory visibility where tagging infrastructure is available. Smart carts and traffic sensors generate data about shopper movement, basket behaviour, and store flow. Shelf-level sensing, including the type of capability developed by MOOS, is different because it focuses on what happens at the product location itself. The relevance of this comparison is that each technology creates value from a different control point in the store.

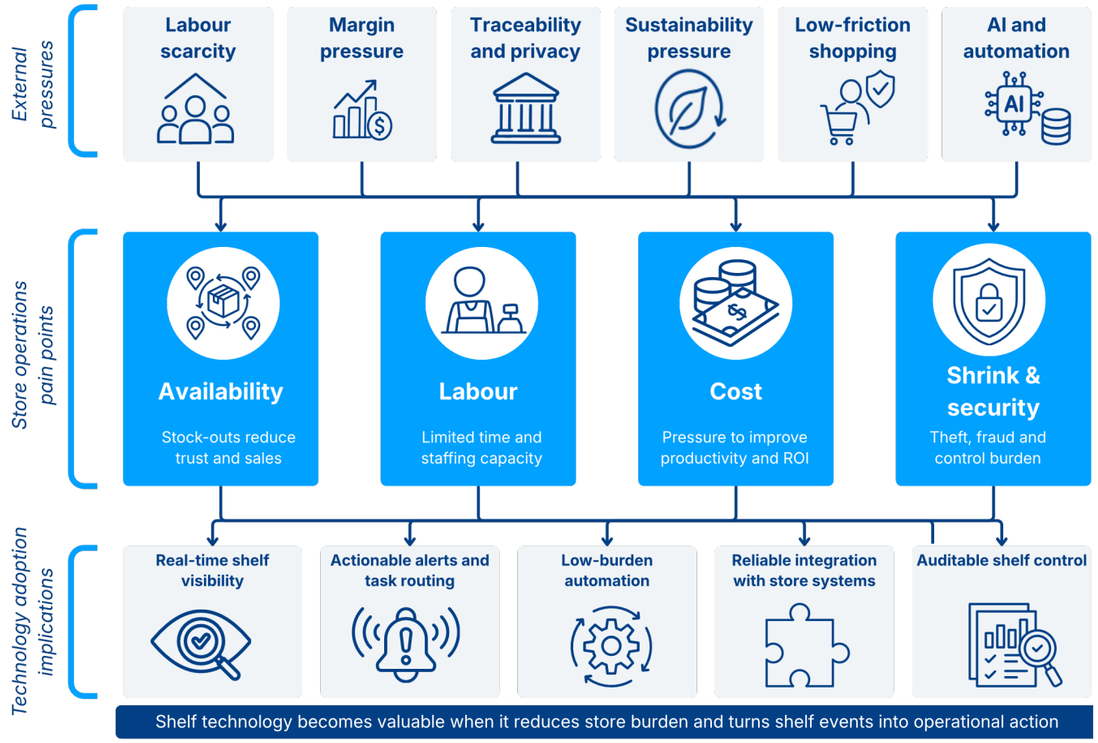

Figure 2.3. Store operations pressure map.

At the same time, general technology developments in AI, connectivity, and cloud and edge computing increase what is technically feasible in near-real-time store sensing, while also raising expectations for cybersecurity, reliability, and integration quality (McKinsey & Company, 2025). This matters because the technical feasibility of sensing is no longer the only differentiator. As AI, connectivity, and edge infrastructure become more mature, more actors can claim to generate store data. The more difficult question becomes which signals are accurate enough, timely enough, and operationally usable enough to be trusted in execution. For MOOS, this shifts the positioning problem from *"can the shelf be sensed?"* to *"which shelf signals are reliable and useful enough to become part of operational decision-making?"*

Electronic Shelf Labels (ESLs) are a relevant infrastructure class because they connect pricing and promotion execution to central systems and change the shelf from a static fixture into a managed digital surface (Deloitte, 2025). Their relevance for MOOS is not that they solve shelf sensing directly, but that they show how the shelf can become part of a broader store technology stack. ESLs infrastructures may provide physical attachment points, power, communication layers, pricing interfaces, and, in some cases, media or analytics functions. This makes ESL providers important because they can become partners, hosts, or gatekeepers for other shelf-level technologies. For MOOS, this creates a positioning tension: the company may benefit from being embedded in shelf infrastructure, but it may also become dependent on actors that control access to that infrastructure.

Across these technologies, integration and operating model readiness repeatedly appear as scaling constraints. Academic work on digital transformation highlights that legacy constraints and organisational routines can become binding, even when individual technologies perform adequately in isolation (Vial, 2019). Sector publications echo this by emphasising that retailers must modernise architecture and ways of working to scale store technologies beyond pilots (Bick et al., 2022). For MOOS, this means that technology trajectories should be read as warnings about positioning. Shelf sensing may create a valuable signal, but the signal will compete for integration capacity, workflow attention, and ownership within already crowded store technology roadmaps.

Table 2.1 summarises the selected technology trajectories as a scaling-constraint matrix. The purpose of the table is not to compare all technologies in detail, but to show how different technologies create value from different store control points and how each trajectory exposes integration, governance, or operating-model constraints that are relevant for MOOS.

Three implications for MOOS follow from these technology trajectories. First, store technologies tend to create value after they become part of wider store systems, not while they remain isolated

Technology route	*Store fit*	*Data quality*	*Integration load*	*Ownership*	*Routine change*	*Implication*
ESL and shelf edge	High store fit	Medium data dependency	High integration load	Shared ownership	Medium routine change	MOOS should not try to own the full infrastructure here, but connect where shelf events add actionable information.
Computer vision	Variable store fit	High data dependency	High integration load	Complex ownership	High routine change	MOOS can stand out through a more focused and less privacy-sensitive sensing route than camera-based solutions.
RFID and item tagging	Physical tagging needed	High data quality need	High integration load	Supply-chain ownership	Medium routine change	RFID shows that full item-level traceability can be heavy; MOOS may be more attractive as a lighter, shelf-focused route.
POS and inventory analytics	Indirect shelf fit	Dependent on transaction data	Backend integration	Commercial or IT ownership	Low shelf visibility	MOOS can add value where backend data remains indirect and gives too little insight into physical shelf reality.
Shelf sensing	Direct shelf fit	Event-level data	Selective integration	Shared ownership	Action routine needed	This is MOOS' most logical core route, provided signals are reliable and can be connected to store action.
Workflow and task tools	Strong store fit	Depends on input signal	Medium integration load	Operations ownership	High routine dependency	MOOS should not replace these tools, but feed them with better shelf-event signals and trigger logic.

Table 2.1. Store technology trajectories and scaling constraints.

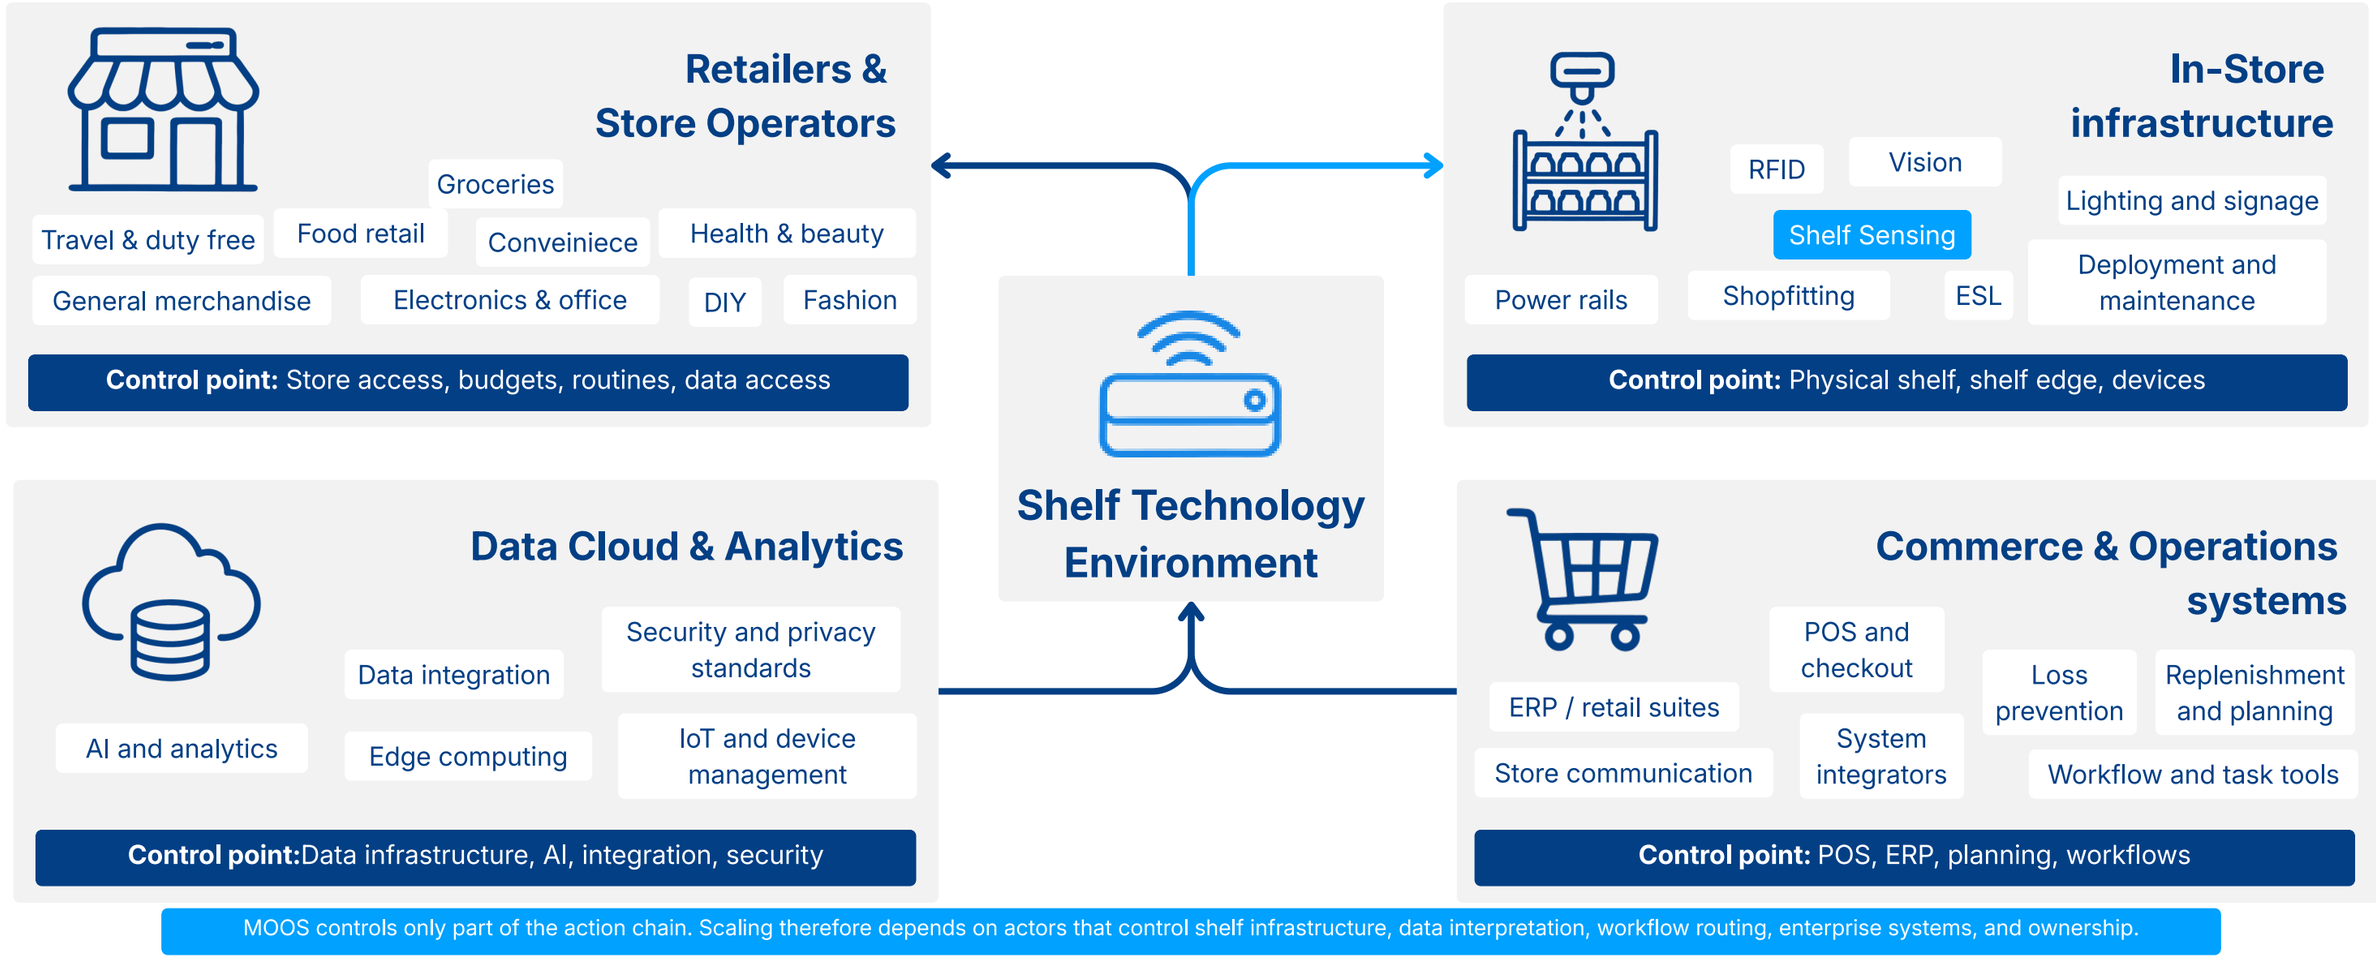

Figure 2.4. Actor categories and control points in the retail technology environment

tools. Shelf sensing therefore needs to be understood as a shelf-event signal that may connect to broader operational system categories, such as transaction, replenishment, loss-prevention, or store execution systems. Second, competition is likely to occur around control points rather than only around direct substitutes. ESL providers, checkout systems, RFID infrastructures, camera platforms, and data platforms each control a different route to store-level insight. Third, repeatability matters. Shelf sensing is more likely to scale when installation, interpretation, and follow-up can be repeated across stores. This narrows the positioning problem: The positioning problem is therefore shaped by where shelf-event signals can create operational value, which actor categories may need to act on them, and which system links may be required at a category level.

2.4 Retail technology actor categories and control points

After the technology trajectories discussed in Section 2.3, the relevant question is not only which technologies can detect, display, or process store information. For MOOS, the more strategic question is which actors control the systems that are needed to turn shelf-level information into operational action.

Retail technologies are embedded in multi-actor systems that connect retailers, vendors, integrators, infrastructure providers, and store teams through shared data flows and operational handoffs. In omnichannel retail, the store becomes part of a wider channel system in which physical and digital operations depend on shared information (Verhoef et al., 2015). This makes coordination more demanding, because replenishment, pricing, fulfilment, and customer-facing processes need consistent data across several functions and systems (Ishfaq et al., 2022). Sector evidence in the European context also frames store enablement as part of a broader transformation agenda that spans channels, logistics, and store systems (McKinsey & EuroCommerce, 2022).

To make this environment useful for MOOS, this section first separates actor categories from individual companies. The aim is not yet to compare direct competitors, but to identify which types of actors control the main interfaces around store execution. These control points include the physical shelf, the shelf edge, transaction data, enterprise systems, replenishment logic, loss-prevention workflows, data infrastructure, integration decisions, and store routines. Figure 2.4 therefore provides a category-level map before representative actor categories are discussed at a category level later in the section.

This figure maps the main actor categories around store technology and shows which control points they occupy. It visualises how control is distributed across the physical shelf, shelf edge, POS and checkout, enterprise systems, planning, LP, data platforms, integration partners, and store execution.

The figure should be read as a structural overview rather than a vendor comparison. The physical shelf is shaped by shopfitting and display suppliers. The shelf edge is increasingly occupied

No.	Actor group	Typical control point	Relevance for actionable shelf insight	Public-safe examples of system category
1	Retail organisation	Store operations, decision rights, budgets	Determines whether shelf information is acted on	Store operations, category management, IT, LP
2	Store teams	Daily execution and response routines	Receive or ignore shelf-related signals	Replenishment routines, shelf checks, incident response
3	Shelf-edge infrastructure providers	Shelf-edge hardware and display layer	May control shelf-edge data or installation space	Electronic shelf labels, digital shelf infrastructure
4	Sensing and vision providers	Detection of shelf, item, or behaviour signals	Provide alternative or complementary visibility routes	Shelf sensing, camera analytics, RFID
5	Transaction and checkout systems	Sales and checkout event data	Provide indirect evidence of shelf and product movement	POS, self-checkout, transaction analytics
6	Enterprise and replenishment systems	Stock records and planning logic	Needed for OSA-related closure and backend correction	ERP, inventory, replenishment systems
7	Workflow and task tools	Task routing and operational follow-up	Translate signals into store action	Task management, workflow tools
8	Loss-prevention systems	Incident handling and investigation	Support theft-related escalation and evidence routines	Security systems, case management, CCTV
9	System integrators	Integration and implementation decisions	Shape whether data flows can be connected	Integration, implementation, consulting
10	Data and cloud infrastructure	Data storage, access, and analytics	Enables interpretation and governance of shelf signals	Cloud, APIs, analytics platforms

Table 2.2 Representative actor groups and control point examples in the retail technology environment

by ESL vendors. Product identification and item visibility are supported by RFID and barcode infrastructures. Checkout and transaction data are controlled by POS, self-checkout, and payment-related systems. Replenishment and planning decisions are shaped by retail planning platforms and enterprise systems. Loss-prevention systems influence how shrink and theft signals are interpreted. Above these layers, cloud providers, data platforms, and systems integrators influence how information is connected, secured, and scaled across store networks. These control points matter because a new sensing technology does not enter an empty environment. It enters a store architecture in which other actors already control interfaces, data access, routines, and commercial relationships.

Representative companies can be placed within these actor categories because they control or influence specific parts of the store technology system. ESL vendors, for example, occupy the shelf edge. POS and checkout providers control transaction data. ERP and retail suite providers shape system architecture and master data. Shopfitting firms influence the physical shelf environment. Loss-prevention vendors shape security routines and evidence logic. Integrators influence whether new technologies can be connected to existing retailer systems. Table 2.2 translates the category-level map into representative actor groups and control point examples. The table should not be read as a competitor ranking, but as a control-point overview that shows where MOOS may encounter partners, gatekeepers, infrastructure owners, or adjacent solution providers.

The table shows that the retail technology environment is organised around multiple control points rather than around one dominant technology category. For MOOS, this means that shelf sensing is not positioned in a neutral space. It must connect to actors that already control store access, shelf infrastructure, transaction data, planning systems, security routines, data platforms, and implementation decisions. The relevance of these actor categories lies not only in their individual company profiles, but in the interfaces they occupy around shelf-level information and store execution.

The implication for MOOS is that the competitive field should be read as a set of partial substitutes and possible complements. ESL systems are strong where retailers need shelf-edge infrastructure, pricing execution, and digital shelf communication. They are weaker when the problem is direct product-event detection. Camera platforms are strong where broad visual context and security monitoring matter. They may be weaker when the retailer needs a non-camera signal that points to a specific shelf event. RFID is strong where tagging infrastructure and item-level control already exist, but it is less suitable where product economics or category conditions do not support tagging. POS data is strong for transaction evidence, but weak for real-time shelf reality. Manual routines remain flexible, but they are labour-intensive and inconsistent. This creates a positioning space in which MOOS can be framed less as a stand-alone replacement and more as a specialised shelf-event layer that strengthens selected store technology routes.

The actor categories can also be grouped by their practical role in store technology decisions. Some actors organise broader systems, such as POS platforms, ERP suites, planning platforms, large ESL ecosystems, cloud infrastructures, and retailer-owned technology stacks. Other actors provide specialised components, such as sensors, RFID tags, labels, cameras, antennas, fixtures, or analytics modules. A third group focuses on integration and implementation, especially when retailers need to connect new tools to legacy systems, data platforms, security rules, and store workflows (Bick et al., 2022; Vial, 2019). These roles are not fixed. The same company may act as a component supplier in one context and as a broader system organiser in another. For MOOS, this means that actor categories cannot be assessed only as partners or competitors. Their role depends on which control point is relevant in a specific retailer setting.

For MOOS, the most important interfaces are not only physical shelf interfaces, but also informational and organisational interfaces. A shelf event creates operational value only when it can be routed to the right system or function, such as an alerting interface, a staff workflow, a POS record, a replenishment system, a loss-prevention process, or an analytics environment. In practice, these systems are often owned by different internal functions or external vendors. This makes the pathway from shelf event to store action a potential bottleneck.

Figure 2.5a and b compare two data flow situations around the shelf. The first flow shows the existing store technology environment without MOOS. In this situation, shelf reality is inferred through systems such as POS, inventory management, shelf media, ESL, camera vision, staff tools, team communication, and shopfitting infrastructure. These systems can each provide part of the picture, but they are not automatically connected to the same action route or decision owner. The full actor maps and stakeholder overview are provided in Appendix B.2.

The second flow shows where MOOS's shelf sensor enters this environment. MOOS adds a physical shelf-event signal between shelf reality and adjacent store systems. This creates several possible connection routes, but each route also brings dependencies around data access, interface ownership, workflow fit, and follow-up action. The figure therefore shows why MOOS's positioning problem is broader than sensor performance alone.

A signal may originate at the shelf, but its value is realised elsewhere: in a replenishment action, a security response, a POS comparison, a category-management decision, or a store task. The signal must be trusted, routed to the right system or actor, and embedded in a routine that can be repeated across stores. This is why the next chapter develops an ecosystem-as-structure lens rather than treating MOOS's position as a simple product comparison.

For MOOS, this means that technical performance is only one condition for relevance. The company also needs to determine which actors it must connect to, which interfaces it should support, and where it can create value without taking ownership of the full store technology stack. Section 2.5 therefore narrows the analysis from actor categories and control points to the alternative solutions and competitors that address shelf-level visibility, availability, and shrink.

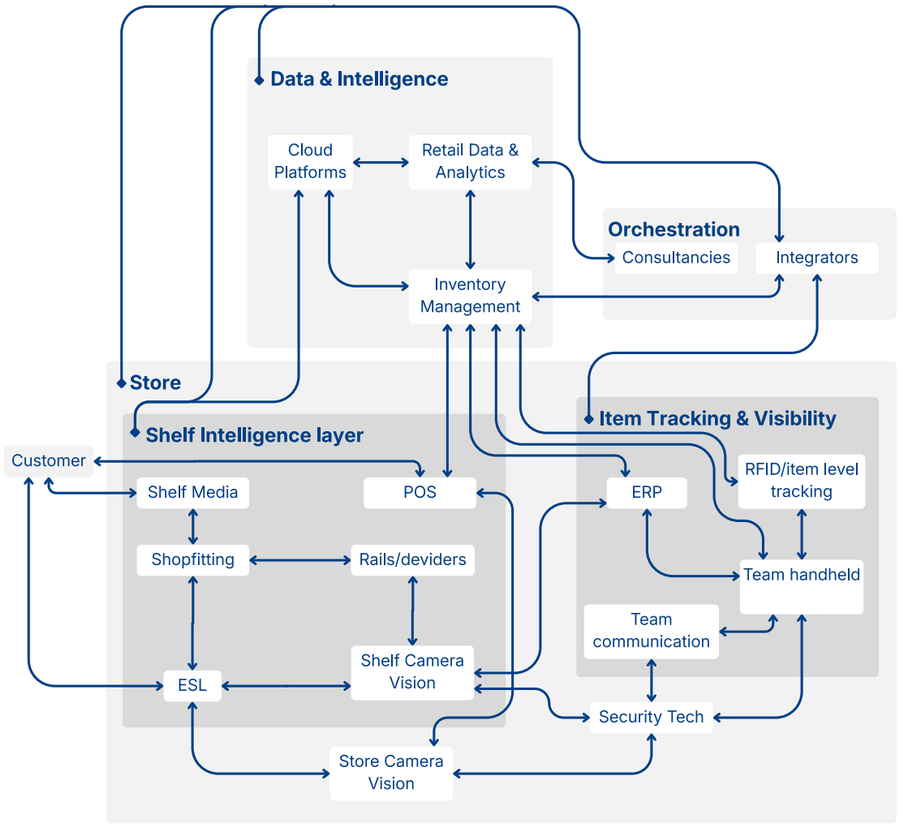

Figure 2.5a. Data flow and integration bottlenecks around shelf events.

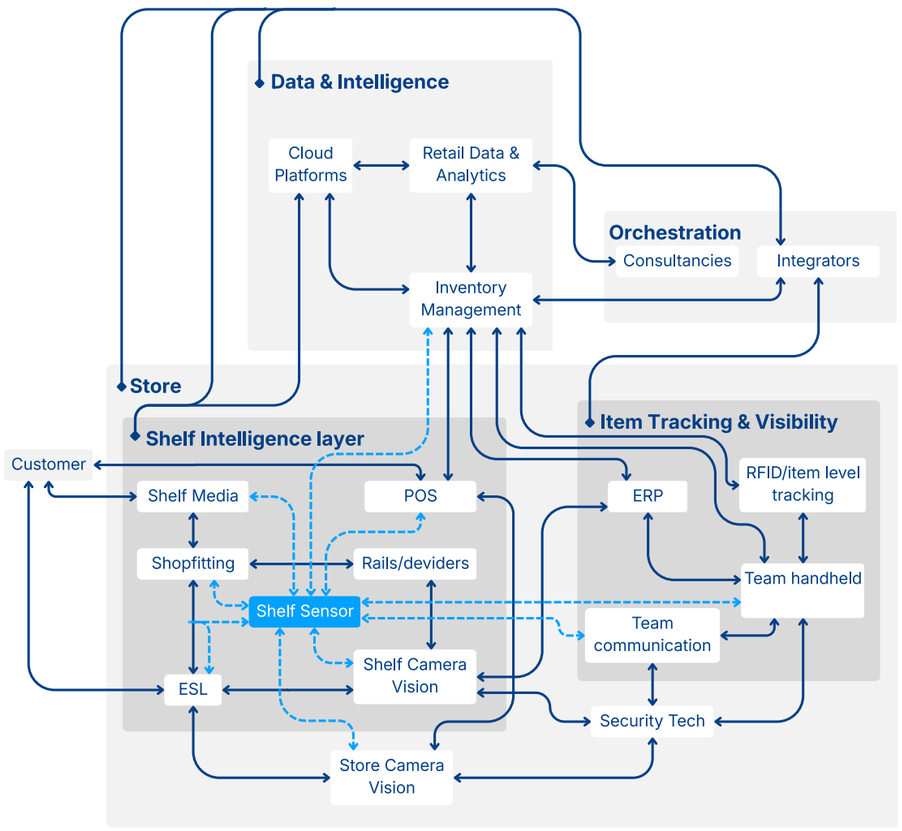

Figure 2.5b. Data flow and integration bottlenecks around shelf events, including shelf sensing.

2.5 Alternative solution routes around shelf-level visibility

This section narrows the context from actor categories to the alternative solution routes that address shelf-related problems. The analysis is not presented as a ranking of competitors. It is used to clarify how different technology categories create shelf-level visibility, which control points they rely on, and where their operational constraints differ.

The aim is not to identify one direct substitute. Retailers already use manual routines, shelf-edge systems, cameras, RFID, checkout data, loss-prevention tools, and analytics platforms to infer what happens at the shelf. From an operational point of view, the job around the shelf includes keeping products available, maintaining pricing integrity, reducing losses, and limiting workload for store staff. These tasks become harder when assortment complexity increases, labour is constrained, and inventory records diverge from shelf reality (McKinsey & EuroCommerce, 2022; DeHoratius & Raman, 2008).

For MOOS, competition should therefore be understood functionally rather than only technologically. The relevant question is not which company uses the same sensing principle, but which solution routes retailers may use to detect, interpret, and act on shelf problems.

A more detailed scan of specific companies and solution providers was used in the internal analysis but is not included in the public repository version.

2.5.1 Solution routes around shelf-level visibility

The alternative solutions around the shelf can be grouped by the mechanism through which they create shelf-related visibility. Some solutions digitise the shelf edge, some identify items, some observe behaviour, some use transaction data, and others organise operational response. This distinction is important because these routes may aim at similar outcomes, such as better availability or lower shrink, while relying on different control points in the store technology system.

Figure 2.6 summarises the main solution routes through which retailers can create visibility around shelf-level problems. The figure is not organised by company type, but by the mechanism through which each route helps retailers detect, interpret, or act on issues such as OOS situations, shrinkage, pricing errors, and store execution problems.

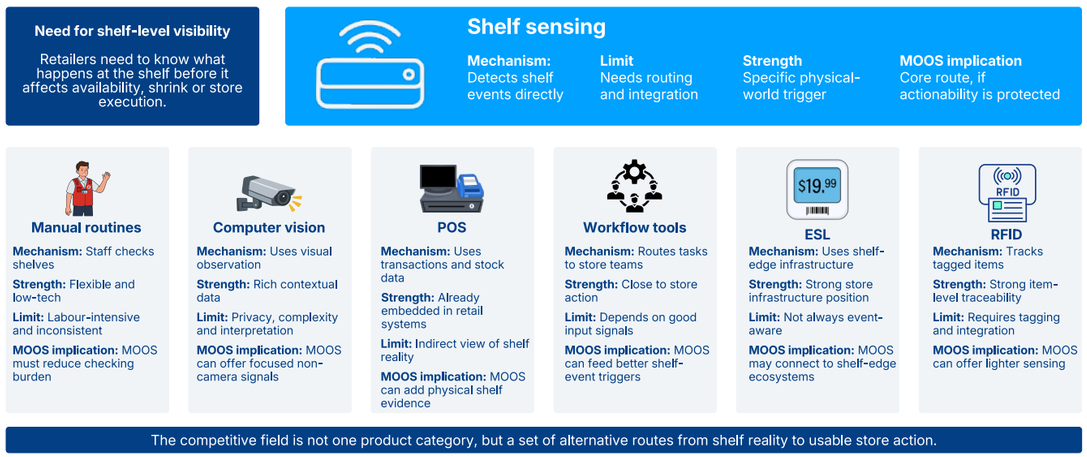

Figur2.6: Solution routes around shelf-level visibility. Source: own analysis based on the desk-research market scan.

The figure shows that alternatives around the shelf do not compete through one shared technical mechanism. Manual routines rely on staff observation, ESL systems digitise the shelf edge, RFID identifies tagged items, cameras observe behaviour or shelf states, checkout systems infer events from transactions, and analytics platforms organise interpretation and follow-up. Specialised shelf sensing occupies a more direct route because it aims to detect shelf events at the location where they occur. For MOOS, this means that the competitive field should be understood functionally: the relevant question is not only which technology is most similar, but which route retailers use to create usable shelf-level information.

Technology alternatives can be organised by the main mechanism through which they produce shelf-related signals. Shelf digitisation tools, such as ESLs, support pricing execution and can reduce manual work, but they do not automatically provide product-level availability or theft signals unless combined with other sensing layers (Deloitte, 2025). Checkout technologies, including self-checkout and checkout-less approaches, aim to reduce friction and may improve throughput, but they also introduce operational challenges that lead to continued iteration and complementary innovations (EY, 2023). Item identification approaches, such as RFID tags, can support automated checkout and stock management when combined with suitable infrastructure, which shows how sensing value depends on complementary components and store readiness (McKinsey & EuroCommerce, 2022). Camera-based approaches are often positioned as a way to observe behaviour and detect anomalies. Specialised shelf-event information can also function as a complement to existing operational or security systems where it helps focus follow-up action. This means that camera-based, transaction-based, and shelf-sensing routes should not only be viewed as substitutes, but also as potentially complementary categories within broader store operations.

Loss-prevention solutions form a separate but overlapping competitive field because many shelf technologies are evaluated partly through their impact on shrinkage. The scale of shrinkage and the reported rise in organised retail crime make it likely that retailers will continue to invest in tools that support prevention, investigation, and staff safety, including approaches that connect incidents to evidence and operational response routines (National Retail Federation, 2023). The National Retail Federation also shows that retailers are not only relying on traditional measures, such as CCTV, locking cases, EAS, and LP staffing, but are also researching or piloting technologies such as AI-enabled fraud detection, RFID, body-worn cameras, and facial recognition (National Retail Federation, 2023). In this context, MOOS' theft-prevention framing emphasises operational alerts and diagnostics that connect shelf events to POS data and video review workflows, which positions the product as a data source that can strengthen broader loss-prevention systems (MOOS, 2025).

The implication for MOOS is that the competitive field should be read as a set of partial substitutes and potential complements. ESL systems are close to the shelf, but they do not inherently solve shelf-level sensing. Camera platforms can observe behaviour, but they may need more precise shelf-event data to reduce interpretation errors. RFID can improve item visibility, but it depends on tagging infrastructure and category fit. Loss-prevention systems already organise evidence and response, but they may lack direct real-time shelf signals. This creates a positioning space in which MOOS can be framed less as a stand-alone replacement and more as a specialised shelf-event layer that can strengthen existing store technology routes.

2.5.2 Representative shelf-related competitors and adjacent providers

The alternative solution space was compared at the level of solution routes rather than individual companies. This is important for the public version of the thesis, because the aim is not to present a supplier ranking or a company-level competitor analysis. Instead, the comparison clarifies how different technology categories create shelf-level visibility, which control points they rely on, and where their operational limitations appear.

Table 2.3 compares the main solution routes that shape the ecosystem around actionable shelf insight. These routes include specialised shelf sensing, shelf-edge digitisation, camera-based analytics, RFID and item identification, POS and checkout systems, enterprise and replenishment systems, loss-prevention systems, workflow tools, shelf infrastructure, and system integration. Each route creates visibility or actionability from a different position in the store technology ecosystem.

The comparison shows that no single route fully owns the problem of actionable shelf insight. Some routes are close to the physical shelf, while others operate through transaction data, enterprise systems, security routines, workflow tools, or implementation architecture. This distinction matters because shelf-level value is not created only by detecting a shelf event. It also depends on whether the event can be interpreted, connected to relevant operational data, routed to the right actor or system, and followed up within existing store routines.

Solution route	*Main control point*	*Shelf visibility route*	*Main contribution*	*Main limitation at category level*	*Relevance for actionable shelf insight*
Specialised shelf sensing	Physical shelf surface and product movement	Direct shelf-event sensing	Provides shelf-level event information close to where the issue occurs	Usually needs connection to wider store systems and routines	Shows why sensing must be linked to interpretation, routing, and follow-up action
Shelf-edge digitisation	Shelf edge and digital label infrastructure	Digital shelf edge, pricing, and shelf communication	Creates a digital layer close to the product display	Does not automatically solve product movement, availability, or theft detection	May support shelf-level workflows, but often needs additional sensing or system links
Camera-based shelf analytics	Visual shelf or aisle observation	Image-based shelf monitoring and interpretation	Can provide broad visual context and shelf-state information	Can be affected by occlusion, calibration, privacy, and deployment complexity	Offers an alternative visibility route, especially where visual context is important
Store-level camera analytics	Store cameras and behavioural observation	Store-level visual analytics	Can support traffic, behaviour, and security-related insight	Often less specific to individual shelf events	May complement shelf-event data in broader operational or security routines
RFID and item identification	Tagged items and identification infrastructure	Item-level identification and inventory visibility	Strong where item-level tracking and tagging infrastructure are feasible	Depends on tagging economics, category fit, and infrastructure adoption	Relevant for inventory accuracy and traceability, but not always feasible across all categories
POS and checkout systems	Transaction and checkout data	Transaction-based evidence after product movement	Provides strong evidence of sales and checkout events	Observes shelf problems indirectly and often after the event	Important for comparing shelf reality with transaction records, but not sufficient alone
Enterprise retail and replenishment systems	Stock records, planning, replenishment, and master data	System-based planning and inventory logic	Shapes replenishment decisions, stock correction, and backend closure	Often far from physical shelf reality and dependent on data quality	Critical for turning availability-related signals into accountable process changes
Loss-prevention and security systems	Security routines, escalation, and evidence handling	Incident detection, deterrence, and investigation support	Supports shrink-related response and evidence workflows	May lack direct product-level shelf-event triggers	Relevant for theft-related actionability when shelf signals need safe escalation
Store workflow and communication tools	Staff task routing and operational follow-up	Alerts, tasks, and staff communication	Translates information into store-level action	Can create alert fatigue or extra workload if not well integrated	Important output layer for making shelf insight actionable in daily routines
Shopfitting and shelf infrastructure	Physical shelf, fixtures, and product presentation	Physical product placement and shelf configuration	Shapes installation, maintenance, and standardisation conditions	Does not create data or insight by itself	Relevant because physical deployment conditions affect scalability
System integration and implementation	Architecture, implementation, and change management	Connection between systems, data, and routines	Enables new data flows to fit existing retailer architectures	Can increase dependency load and implementation complexity	Important for scaling beyond pilots and connecting shelf insight to existing systems
Cloud, data, and analytics infrastructure	Data processing, storage, governance, and analytics	Data aggregation and interpretation	Enables shelf data to be combined with other operational data	Requires governance, cybersecurity, and integration discipline	Supports broader interpretation of shelf events, but does not itself create shelf action

Table 2.3. Category-level comparison of shelf-related solution routes.

Specialised shelf sensing is distinct because it can create event information close to the place where the shelf condition changes. However, this does not mean that it replaces the surrounding ecosystem. The category-level comparison shows that shelf sensing usually needs other system categories to become actionable. For example, availability-related use cases may depend on stock records, replenishment logic, or store tasks, while shrink-related use cases may depend on escalation routines, security processes, or evidence handling. These dependencies do not define a fixed partner route. They show why shelf sensing must be understood as one layer in a wider activity structure.

The table also shows why alternatives should not only be understood as direct substitutes. Shelf-edge systems, camera analytics, transaction systems, RFID, workflow tools, enterprise systems, and implementation actors can each function as substitutes, complements, or dependency points depending on the use case. This means that MOOS's positioning problem is not only a question of technical differentiation. It is also a question of ecosystem alignment: where shelf-event information should connect, which actor categories control the relevant links, and which routines can turn information into action.

For the thesis, the implication is that actionable shelf insight requires an ecosystem lens. The relevant positioning space lies between shelf reality and store action, but the value of that space depends on multiple surrounding control points. This conclusion narrows the research problem for the following chapters. MOOS does not only need to show that shelf sensing can detect events. It also needs to understand how those events can become reliable, interpretable, routable, and usable within existing retail operations.

2.5.3 Capability comparison and market gap for MOOS

Building on this category comparison, the market gap can be assessed through three positioning dimensions. These dimensions are not definitive market facts, but analytical comparisons based on the desk-research scan. They help clarify where direct shelf-event sensing may be stronger than indirect visibility routes, and where MOOS may remain dependent on other systems.

The first positioning comparison concerns the relation between sensing directness and operational actionability. Some routes create visibility indirectly, for example through transactions, periodic audits, staff interpretation, or aggregated analytics. Other routes are closer to the shelf event itself. This distinction matters because shelf-level problems such as OOS situations, suspicious product removal, incorrect facing, or replenishment failures lose value when they are detected too late or only inferred after the event. The more direct the signal, the more likely it is that the retailer can connect the insight to timely operational response.

The first comparison suggests that MOOS is most differentiated when shelf-level visibility must be both direct and actionable. Camera analytics may provide broader contextual information, RFID may provide strong item-level identification in suitable categories, and POS data remains essential for transaction evidence. The relevant point is narrower: MOOS appears to occupy a space where shelf events can be detected at the location where they occur and then translated into store-level action signals. This is different from routes that infer shelf problems from downstream data or depend mainly on human checks.

A second comparison concerns the trade-off between proximity to shelf reality and the burden of deployment or integration. Technologies close to the shelf may create richer information about what happens in the store, but they also need to fit into physical fixtures, installation routines, maintenance processes, data architectures, and store workflows. Technologies further away from the shelf may be easier to connect at system level, but they often provide less direct information about the shelf event itself. This trade-off matters for MOOS because its sensing advantage only becomes useful when it can be deployed and connected without creating too much operational complexity.

The second comparison nuances the market gap. MOOS may be close to the shelf and less visually intrusive than camera-based systems, but this does not automatically make deployment simple. Its position depends on how well the sensing layer can be installed, configured, maintained, and connected to existing store systems. This interpretation is consistent with the broader observation that retail technologies often create value only when architecture, operating models, and organisational routines support scaling beyond isolated pilots (Bick et al., 2022). For MOOS, the issue is therefore not only how well the sensor works, but how easily the signal can become part of the retailer's operating model.

A third comparison concerns privacy exposure and signal specificity. This is relevant because retailers increasingly need detailed operational data, while also facing concerns around visual surveillance, data ownership, security, and customer trust. Not all visibility routes create the same type of data. Camera-based systems may provide broad behavioural and contextual information, but they can also raise privacy and governance questions. POS data is less visually intrusive, but it only records what reaches the checkout. ESL systems have low privacy exposure, but their core function is not direct product-event detection. For MOOS, the relevant question is whether non-camera shelf sensing can provide specific operational signals without expanding visual observation in the store.

The third comparison indicates that the market gap is not only about precision, but also about the type of data that is used to create precision. Camera-based systems can be powerful where visual context matters, but they also depend on interpretation and may face privacy, occlusion, calibration, or false-positive challenges depending on the store setting. RFID can create strong item-level visibility, but only where tagging infrastructure and product-category economics make sense. MOOS appears to occupy a different position: it aims to generate specific shelf-event signals without using visual observation. This may make the route especially relevant for use cases where retailers need operational evidence at shelf level, but do not want to expand camera-based monitoring.

Taken together, the three comparisons indicate a specific market gap for MOOS. Many existing solutions provide broad store observation, transaction evidence, shelf-edge digitisation, item-identification infrastructure, or loss-prevention response. Fewer combine direct shelf-event detection, low privacy exposure, product-level relevance, and a route into operational triggers. MOOS therefore appears strongest when the use case depends on knowing what happened at the shelf, when it happened, and whether that event should trigger a store action. This is most relevant for high-risk removal events, targeted theft diagnostics, and availability situations where indirect data is too late or too incomplete. MOOS appears weaker when the retailer mainly needs broad store analytics, full enterprise planning, transaction reconciliation, item-level inventory control across the whole chain, or a complete workflow system. In those cases, cameras, RFID, POS systems, planning tools, or loss-prevention platforms may be stronger primary systems.

This white space should still be interpreted carefully. A strong sensing capability does not automatically create a strong market position. Retail technologies often depend on store architecture, operating models, staff routines, and integration capacity before they create value at scale (Bick et al., 2022). For MOOS, this means that shelf-event data only becomes valuable when it is connected to replenishment, shrink response, diagnostics, or decision-making. The market gap is therefore technical, butalso organisational and infrastructural.

2.5.4 Implications for the positioning problem

The category-level comparison leads to three implications for the positioning problem. First, specialised shelf sensing should not be positioned against one direct competitor category. Retailers can already create partial shelf-level visibility through several routes, including shelf-edge systems, camera-based analytics, RFID and item identification, transaction data, enterprise systems, workflow tools, security systems, and system integration. These routes do not replace each other in a simple way. They occupy different control points in the wider store technology ecosystem.

Second, the comparison shows that the value of shelf sensing depends on the use case and on the surrounding activity chain. A shelf event can only become valuable when it is reliable, interpretable, and connected to a possible follow-up action. This means that shelf sensing should not be assessed only by detection capability. It should also be assessed by whether the signal can be linked to existing routines, data categories, decision ownership, and operational response mechanisms.

Third, the comparison shows why actionable shelf insight is an ecosystem problem. Some solution routes are close to the physical shelf, while others are closer to transaction data, planning systems, security routines, workflow execution, or integration architecture. This creates a distributed positioning space. The question is therefore not only what a shelf-sensing component can detect, but where it can credibly contribute within a wider structure of actors, activities, links, and dependencies.

For MOOS as the focal case, this means that the positioning challenge should be framed at the level of ecosystem fit rather than direct technological substitution. The relevant issue is how shelf-event information can move from detection to interpretation, routing, and accountable action without requiring MOOS to own the full retail operating system. This conclusion prepares the theoretical framework in Chapter 3, where the ecosystem-as-structure lens is used to analyse activities, actors, positions, links, alignment conditions, and adoption risks.

Chapter 3

Literature Review

3.1 Theory background
3.2 Core building blocks from the “ecosystem-as-structure” and related literature
3.3 Operationalising the lens: risks, bottlenecks, and internal readiness
3.4 Integrated conceptual framework

Chapter 2 showed that actionable shelf insight is shaped by multiple technology categories, actor groups, and control points rather than by one direct competitor category. This points to an alignment problem rather than a simple adoption or product-comparison problem. Shelf-condition information only becomes valuable when it can move through systems, workflows, interfaces, ownership structures, and actor responsibilities in a way that supports accountable store action.

This chapter therefore develops the conceptual lens used to diagnose these dependencies. The unit of analysis is the ecosystem-alignment structure around actionable shelf insight, not MOOS as an isolated firm and not shelf sensing as an isolated technology. Ecosystem-as-structure is used as the core lens because it defines an ecosystem around a focal value proposition and examines the activities, actors, positions, and links that must align for that value proposition to be realised (Adner, 2017). The framework then combines this lens with risk and readiness concepts to make the links between shelf sensing, interpretation, routing, and store action diagnosable.

3.1 Theory background

Several theory families can inform parts of this problem, but they do not provide the full analytical fit required here. The resource-based view (RBV) helps explain firm resources and advantage, but its main unit of analysis is the focal firm rather than cross-actor alignment in a shared activity system (Barney, 1991). Real options logic is useful for understanding staged commitments under uncertainty, for example in pilots and rollouts, but it does not map the activities and dependencies needed to realise a focal value proposition (Dixit & Pindyck, 1994). Network theory offers useful language for position and centrality, yet it does not specify which activities must align for value creation in a specific ecosystem (Freeman, 1979). Strategic foresight supports future-oriented scanning, but this chapter requires a diagnostic structure for present ecosystem dependencies rather than a broad future-scanning lens (Rohrbeck, 2010).

Ecosystem-as-structure is selected as the core theoretical lens because it fits the central problem identified in Chapter 2: actionable shelf insight depends on a value proposition that can only be realised when several activities, actors, positions, and links align (Adner, 2017). This also fits broader ecosystem literature, which emphasises that value creation and value capture depend on complementarities, governance arrangements, and the coordination of interdependent actors rather than on one firm's technology alone (Jacobides et al., 2018). The lens is therefore useful because it starts from the focal value proposition and then asks which activity structure and actor configuration are required to make that value possible.

The supporting concepts are used only to make this ecosystem lens more operational. Co-innovation risk and adoption-chain risk help explain how value realisation may break down when required complements, interfaces, or adoption stages do not mature together (Adner & Kapoor, 2010). Diffusion theory is used only as background logic for staged progression under uncertainty, not as a separate adoption framework (Rogers, 2003). The five gates used in this thesis, pilot, rollout decision, store execution and routinised use, sustained use and renewal, and scale, are therefore an operational translation for enterprise retail deployment. They are not taken directly from Rogers.

Dynamic capabilities are used as a bounded internal readiness layer. They help assess whether MOOS can execute and adapt a chosen ecosystem position when partner conditions, interfaces, and complements change (Teece, 2007). This layer remains subordinate to ecosystem-as-structure. It does not replace the external ecosystem diagnosis, but adds a feasibility check for the positioning options developed later (Warner & Wäger, 2019).

The contribution of this chapter is not a new ecosystem theory. It is a focused operationalisation of ecosystem-as-structure for a retail execution value proposition. Conceptually, it translates the ecosystem grammar of focal value proposition, activities, actors, positions, links, and alignment into a diagnostic framework for actionable shelf insight (Adner, 2017). Methodologically, it turns that grammar into categories that can guide sampling, interview questions, and coding. Practically, it prepares a way to compare positioning options for a component provider whose value depends on actors and systems beyond its direct control.

The specific contribution lies in the link-level focus. Rather than treating alignment only as a general ecosystem condition, this thesis examines whether critical links between actor positions are technically, data-wise, operationally, and institutionally workable. Co-innovation risk and adoption-chain risk are attached to those same links, so that breakdowns can be located in complements, interfaces, gate decisions, store routines, or scaling conditions. Internal readiness is added only as a bounded feasibility layer.

3.2 Core building blocks from the "ecosystem-as-structure" and related literature

3.2.1 Focal value proposition

Having selected ecosystem-as-structure as the core lens, this section defines the building blocks used throughout the framework: focal value proposition, activities, actors, positions, links, and alignment (Adner, 2017). These concepts are not used as a general vocabulary list. They are used to explain how actionable shelf insight can be analysed as a value proposition that depends on several linked activities and actor roles. The section also introduces bottlenecks, complementarities, and substitution as diagnostic concepts that help explain where the ecosystem may fail or where alternative routes to the same value outcome may exist.

In an "ecosystem-as-structure" view, an ecosystem is defined relative to a focal value proposition (Adner, 2017). The focal value proposition specifies the value that is to be realised for an end beneficiary, and it acts as the reference point for identifying which activities and actors are structurally required (Adner, 2017). This implies that the same industry can contain multiple ecosystems, because different value propositions imply different activity structures, different dependencies, and different alignment problems (Adner, 2017).

For this thesis, the focal value proposition is actionable shelf insight. This refers to shelf-condition information that can trigger relevant store decisions or actions through reliable, timely, integrated, and accountable execution. Accuracy can strengthen actionability, but value realisation depends mainly on whether the information can be used in the relevant workflow. This value proposition is grounded in a known retail execution problem: information in retail systems can diverge from physical shelf conditions, which can increase both cost and OOS levels (Fleisch & Tellkamp, 2005). Because inventory records are often inaccurate in practice, actionable shelf insight targets an execution gap where better detection and response can improve availability and performance (Rekik et al., 2019).

Because the ecosystem configuration may differ by operational objective, this thesis treats actionable shelf insight as relevant to two use-case variants: OOS reduction and loss-prevention. These variants are not separate research projects, but they serve as boundary checks for whether the same ecosystem logic holds under different store execution goals.

3.2.2 Activity structure

Activities are the tasks that must be performed to realise the focal value proposition (Adner, 2017). The key analytical move in the ecosystem-as-structure lens is to describe an ecosystem as an activity configuration that must come together, rather than as a list of firms or a focal platform with a vague periphery (Adner, 2017). The activity structure functions as a blueprint, which specifies what needs to happen, in what sequence, and with what dependencies, for the value proposition to materialise (Adner, 2017). This blueprint makes it possible to ask a concrete question in later empirical work, namely, which activity assumptions hold in practice, and where they break down (Adner, 2017).

In the actionable shelf insight setting, the activity structure runs from shelf-event detection to interpretation, routing, decision ownership, store action, and learning. Detecting a shelf event is only the first activity. The signal must also be interpreted, linked to relevant reference data, routed to the right decision owner, and translated into an executable store action. This makes the store context structurally important, because value depends on closing the gap between information and action (Bottani et al., 2017). The activity structure therefore includes both technical activities and organisational activities, which means that better sensing alone is not enough to realise the value proposition (Adner, 2017).

3.2.3 Actors and roles

Actors are the organisations that perform activities within the activity structure (Adner, 2017). Ecosystem membership is not defined by affiliation, for example being part of retail, or being part of technology, but by structural necessity (Adner, 2017). An actor is part of the ecosystem to the extent that its contribution is required for the focal value proposition to be delivered (Adner,

2017). This framing is useful in a crowded retail technology domain, because it discourages inventorying the entire landscape and instead focuses attention on the minimal set of actors that must align.

Roles follow from what actors do in the activity structure and from the dependencies they create or control (Adner, 2017). In this thesis, actors are therefore not treated as a stakeholder list. They are analysed in terms of whether they perform, enable, block, own, or depend on activities needed for actionable shelf insight. A component provider may perform sensing or signal generation, a systems actor may enable integration into enterprise workflows, and a retailer function may own store execution and performance outcomes. These roles keep the analysis anchored in value creation rather than in vendor categories.

3.2.4 Positions and links

Positions show where actors sit in the activity structure and which activities they are responsible for (Adner, 2017). For this thesis, positions are defined at a level that is detailed enough to locate dependencies and possible failure points, but broad enough to compare alternative roles for MOOS.

Links show where positions depend on each other (Adner, 2017). These links can involve data flows, technical interfaces, approvals, commercial agreements, or operational handovers. They are central to the later analysis because actionable shelf insight can fail at these handover points, even when the sensing component works in isolation.

3.2.5 Alignment and misalignment

Alignment is the condition that the activities in the ecosystem can be performed as assumed in the blueprint, at a sufficient level of performance, and with coordinated expectations across actors (Adner, 2017).

Misalignment occurs when an assumed activity or dependency cannot be performed, or cannot be performed reliably, which prevents value realisation even if the focal technology itself is technically strong (Adner, 2017). This makes alignment a central construct for diagnosing why pilots stall, rollouts underdeliver, or solutions fail to sustain use in stores.

To support diagnosis in interviews and coding, alignment is decomposed into four related conditions that can be observed empirically. This decomposition is used as an analytic device to operationalise Adner's alignment construct in a retail technology context (Adner, 2017).

Technical alignment concerns whether systems and interfaces can connect and operate reliably across organisational boundaries.

Data alignment concerns whether data definitions, identifiers, and governance rules allow information to be interpreted and used consistently across actors.

Operational alignment concerns whether information outputs can be translated into routines, task ownership, and execution in the store context, which is often difficult when digital transformation requires changes in processes and organisational arrangements (Vial, 2019).

Incentive alignment concerns whether the actors who must invest effort or accept change also perceive value, and whether concerns such as privacy and shopper attitudes create constraints on adoption choices (Inman & Nikolova, 2017).

3.2.6 Bottlenecks, complementarities, and substitution

Bottlenecks are positions or links that constrain the ecosystem's ability to realise the focal value proposition (Adner, 2017). In this thesis, bottlenecks matter because actionable shelf insight can fail at specific handovers, for example when a signal is available but cannot be integrated, routed, owned, or acted upon. Bottlenecks therefore help identify where the action chain may break down.

Complementarities explain why MOOS may need partners. Even when the sensing component works, value may still depend on complementary systems, data access, workflow tools, store routines, or governance arrangements (Adner & Kapoor, 2010).

Substitution explains why alternative routes may solve part of the same shelf visibility problem through different activity structures, such as camera-based systems, ESLs, RFID, or POS based inference. These concepts help compare possible routes to actionable shelf insight without assuming in advance that MOOS should own the full solution.

3.3 Operationalising the lens: risks, bottlenecks, and internal readiness

Section 3.2 defined the core ecosystem grammar for this thesis. The next step is to make that grammar usable for empirical analysis. This section explains how alignment, co-innovation risk, adoption-chain risk, bottlenecks, points of influence, and internal readiness are used as diagnostic layers. These layers do not replace ecosystem-as-structure. They specify what must be examined when the framework is applied to interviews, coding, and later positioning synthesis.

3.3.1 Alignment conditions in enterprise retail settings

The focal value proposition of actionable shelf insight is operational in nature because it connects information about shelf conditions to store execution and performance outcomes. This has implications for what alignment means in practice. In enterprise retailers, technology change is often constrained by existing architectures, governance, and routines, which makes the problem less about technical feasibility in isolation and more about coordinated change across organisational units (Vial, 2019). Omnichannel settings further increase this complexity because they create dependencies between store operations, supply chain decision-making, and shared data infrastructure (Ishfaq et al., 2022).

Inventory inaccuracy and out OOS events illustrate why alignment cannot be reduced to one technical interface. In retail supply chains, discrepancies between physical inventory and system records can materially affect cost and availability, even when information systems are in place (Fleisch & Tellkamp, 2005). Evidence on inventory inaccuracy also suggests that performance losses are not explained by a single technical failure but by a wider set of process and data conditions that must hold consistently across the store system (Rekik et al., 2019). Field evidence also suggests that improving visibility through technologies such as RFID can reduce inventory record inaccuracy under certain conditions, which indicates that benefits depend on how data collection and operational routines are connected (Hardgrave et al., 2013). These findings motivate an alignment concept that spans both information flows and execution flows.

To make alignment observable in the empirical work, the four dimensions introduced in Section 3.2.5 are applied to each critical link in the ecosystem structure: **technical alignment, data alignment, operational alignment, and incentive** and **governance alignment**.

Together, these dimensions assess whether the required links can work across systems, data definitions, routines, ownership structures, and incentives. In the actionable shelf insight setting, this also includes constraints around acceptable sensing and data practices, because shopper attitudes and privacy concerns may shape what retailers are willing to deploy (Inman & Nikolova, 2017).

3.3.2 Co-innovation risk as complement maturity and integration readiness

Co-innovation risk refers to the risk that a focal offering cannot deliver value because required complements are not ready or not sufficiently integrated (Adner & Kapoor, 2010). In the context of actionable shelf insight, the focal offering is not a single product but a capability that depends on a set of complements, such as sensing hardware, connectivity, data pipelines, analytics, and the operational mechanisms that turn insights into action. The ecosystem lens implies that the focal offering can be technically sound while still failing, if the required complements lag behind or remain uncoordinated (Adner, 2017).

Two mechanisms are particularly relevant for operationalising co-innovation risk in this thesis.

Complement maturity captures whether the complementary components are stable, validated, and deployable under retailer constraints (Adner & Kapoor, 2010).

Integration readiness captures whether the interfaces and dependencies between components are sufficiently specified and supported to enable reliable end-to-end performance (Adner, 2017).

This distinction matters because a complement can be mature in isolation yet still create risk if the integration effort is underestimated, for example, due to data semantics, identity resolution, or operational edge cases that only emerge in stores (Vial, 2019).

In the empirical work, co-innovation risk will therefore be coded at the level of specific links in

the activity chain. This makes it possible to identify whether failure points are driven primarily by missing complements, immature complements, or integration challenges between otherwise mature components (Adner & Kapoor, 2010). It also supports a more concrete discussion of positioning options, because different ecosystem positions expose MOOS to different portfolios of complement dependencies (Adner, 2017).

3.3.3 Adoption-chain-risk as staged, multiple gate progression

Adoption-chain risk refers to the risk that value cannot be realised because several actors must adopt, approve, integrate, maintain, or use the solution before the focal value proposition can be delivered (Adner & Kapoor, 2010). This is not treated as a consumer diffusion problem. Diffusion theory is used only as background logic for understanding adoption as staged progression under uncertainty (Rogers, 2003). For this thesis, that logic is translated into five enterprise retail gates: pilot, rollout decision, store execution and routinised use, sustained use and renewal, and scale. These gates are an analytical device for this study, not a staging directly taken from Rogers.

In the empirical work, adoption-chain risk is analysed in relation to positions and links in the ecosystem structure. The aim is to identify where progression stalls, which actor role functions as gate owner, and which link-level misalignments explain failure to move from one gate to the next.

For operationalisation, the five gates are used to identify where progression stalls, which actor role functions as gate owner, and which link-level misalignments explain the breakdown. Adoption-chain risk is therefore analysed in relation to positions and links in the ecosystem structure. It is not treated as a separate adoption model.

3.3.4 Bottlenecks and leverage points through structure, centrality, and governance

Bottlenecks, within the ecosystem-as-structure view, are points where misalignment in activities, positions, or links prevents the ecosystem from delivering the focal value proposition (Adner, 2017). Bottlenecks are not defined only by where work happens, but by where dependencies concentrate and where failure has the largest downstream impact (Adner, 2017). This motivates the use of structural language to identify leverage points, meaning intervention points where improvements can unlock disproportionate ecosystem performance.

Centrality is a useful structural concept for this purpose, because it characterises which nodes occupy positions that connect many others and can therefore influence coordination outcomes (Freeman, 1979). In ecosystem terms, actors that are central in the relevant activity network often have the ability to set interfaces, standardise practices, or act as intermediaries between otherwise disconnected parties (Freeman, 1979). This does not imply that central actors always create value for all participants, but it does imply that their choices can affect the feasibility of alignment and scaling (Jacobides et al., 2018).

Governance complements this structural view by explaining how participation and value capture are shaped by rules, standards, and decision rights. Ecosystem strategy highlights that value creation and value capture depend on who controls critical interfaces and how complementarities are organised (Jacobides et al., 2018). Platform governance research similarly suggests that rules for access, coordination, and control can shape ecosystem outcomes by influencing partner behaviour and investment incentives (Springer et al., 2025). Together, these perspectives support a practical leverage point logic: leverage points are most credible when they sit at structurally central dependencies and when there is a plausible governance mechanism through which alignment can be enforced or incentivised (Adner, 2017).

Leverage points, for the empirical analysis, will be identified using three criteria. First, the point must be binding for actionable shelf insight, meaning that it constrains value realisation across multiple activities or links (Adner, 2017). Second, the point must have an identifiable actor or small set of actors that can influence it, which can often be approximated through structural centrality and interface ownership (Freeman, 1979). Third, the point must be actionable within realistic governance and capability constraints, which requires attention to incentives and control mechanisms (Jacobides et al., 2018).

3.3.5 Internal readiness as dynamic capabilities for ecosystem execution

Dynamic capabilities are used only as a feasibility overlay. The ecosystem lens explains the external dependencies around actionable shelf insight, while the internal readiness layer asks whether

MOOS can execute and adapt a chosen position within that structure. Dynamic capabilities theory describes this as sensing, seizing, and transforming resources and routines under changing conditions (Teece, 2007). In this thesis, these concepts are used to assess what different ecosystem positions would require from MOOS, not to explain the ecosystem as a whole.

Digital transformation research highlights that capability building is an ongoing process rather than a one-time project, especially when technologies must be integrated into existing routines and architectures (Warner & Wäger, 2019). For early-stage technology ventures, these capability demands can be amplified by resource constraints and by the need to build credibility through pilots and partner-delivered complements (Teixeira et al., 2021). The internal readiness question is therefore not whether MOOS has every capability in-house, but whether it can reliably orchestrate the capabilities required by its chosen ecosystem position (Teece, 2007).

In this thesis, dynamic capabilities are used only as an internal readiness layer. Sensing refers to MOOS' ability to identify relevant bottlenecks, gate owners, partner constraints, and possible points of influence (Teece, 2007). Seizing refers to its ability to turn this understanding into concrete choices about partnerships, integration depth, service delivery, and staged adoption paths (Teixeira et al., 2021). Transforming refers to its ability to adjust the operating model as learning accumulates, for example by standardising interfaces, improving deployment routines, or changing partner roles (Warner & Wäger, 2019). This layer does not explain the ecosystem by itself. It only assesses whether a possible position is executable for MOOS.

Together, these constructs define how the ecosystem lens will be used in the rest of the thesis. Alignment conditions specify what must hold across links. Risk concepts explain where value realisation may break down. Bottlenecks and leverage points help identify where intervention may matter most. Internal readiness specifies what MOOS must be able to execute. Section 3.4 does not introduce new theory. It integrates these elements into a version-one analytical framework for studying actionable shelf insight as an ecosystem-alignment problem.

3.4 Integrated conceptual framework

This section integrates the concepts introduced in Sections 3.1 to 3.3 into one diagnostic framework. The framework is literature-based and intentionally abstract. It does not present empirical findings, nor does it claim that specific bottlenecks, leverage points, or actor relationships have already been confirmed. Its purpose is to make ecosystem positioning conditions diagnosable by showing how actionable shelf insight depends on activity structure, actor positions, links, alignment conditions, ecosystem risks, and internal readiness (Adner, 2017).

The framework separates conceptual structure from research output. It defines how the ecosystem will be analysed, but it does not contain positioning recommendations. Positioning options are developed later from the empirical diagnosis, after activities, links, risks, ownership problems, and internal readiness constraints have been examined.

The framework has three roles in the thesis. First, it works as a diagnostic tool for mapping activities, positions, links, alignment conditions, and ecosystem risks around actionable shelf insight. Second, it provides sensitising concepts for the interview protocol and coding in Chapter 4. Third, it supports the later positioning synthesis in Chapter 6 by making alternative positions comparable in terms of dependencies, risks, leverage points, and internal readiness. Positioning options are therefore not built into the framework itself. They are derived after diagnosis, which helps avoid assuming the answer before the empirical work is conducted.

3.4.1 Framework overview and logic

Figure 3.1 summarises the full diagnostic framework used in this thesis, it should be read from left to right. It starts with the focal value proposition, which sets the ecosystem boundary. It then maps the activity structure needed to realise that value proposition, followed by the actor positions that perform, enable, own, or depend on those activities. The links between positions show where handovers, dependencies, or failure points may occur. Alignment conditions are then applied to these links to assess whether they can work in practice. Co-innovation risk and adoption-chain risk are added as overlays to identify where complements, interfaces, or deployment gates may block value realisation. The final layers identify

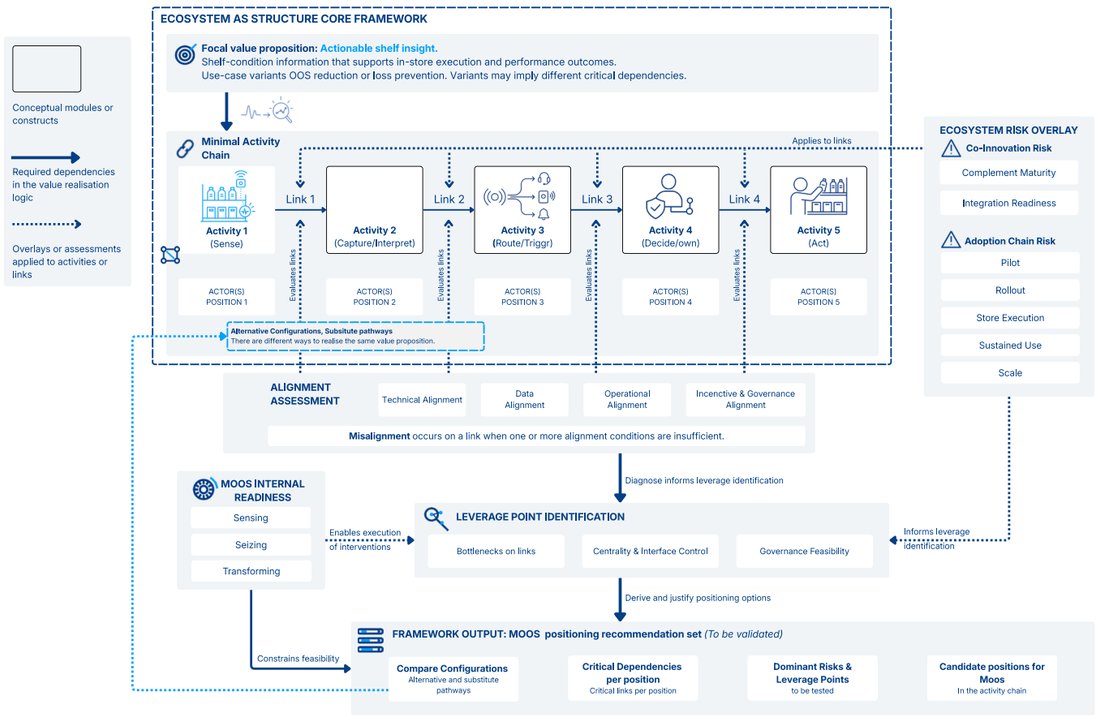

Figure 3.1. Integrated diagnostic framework for actionable shelf insight

bottlenecks, possible points of influence, and MOOS' internal readiness to act on the diagnosed ecosystem conditions.

The notation is deliberately simple. Boxes represent constructs or modules. Solid arrows represent required dependencies in the value-realisation logic. Dotted arrows represent influences or constraints rather than direct activity steps. Overlays represent evaluation lenses that are applied to activities and links. This distinction matters because the framework is not a linear process model. It is a diagnostic structure for analysing whether actionable shelf insight can be realised across actors, systems, interfaces, and routines.

3.4.2 Module 1: Focal value proposition as the boundary condition

The first module anchors the framework around actionable shelf insight: shelf-condition information that is reliable, timely, workflow-integrated, owned, accountable, and able to trigger store action. This keeps the framework focused on operational usefulness rather than sensing accuracy alone. The value proposition is applied to two use-case variants, OOS reduction and LP. These are not separate theoretical models, but boundary checks. OOS cases are likely to make replenishment routines, inventory data, and store task execution more important. LP-related use cases are likely to make escalation routines, evidence requirements, transaction-related information, and supporting security processes more important.

Figure 3.2 introduces the focal value proposition used in this thesis. It shows that shelf insight becomes relevant only when shelf information can support timely and accountable store action.

Focal value proposition: Actionable shelf insight.
Shelf-condition information that supports in-store execution and performance outcomes.
Use-case variants OOS reduction or loss prevention. Variants may imply different critical dependencies.

Figure 3.2. Focal value proposition for actionable shelf insight

3.4.3 Module 2: Activity structure, actors, positions, and links

The second module translates the value proposition into a minimal activity structure: sensing, signal capture, processing, decision support, and store action. The framework starts from these activities rather than from a list of companies, because the empirical analysis first needs to identify what must happen before it can assess who should occupy which position.

Generic actor positions are then mapped onto the activities. The links between these positions are the main unit of analysis. They show where one actor depends on another for data, integration, approval, operational handover, or action. This matters because actionable shelf insight may fail at the interface between actors, even when the individual sensing technology works.

Figure 3.3 translates the value proposition into a basic activity chain. It shows the minimum sequence of activities required to move from sensing to action. Added are the actor positions to the activity chain. The figure highlights that value depends not only on activities, but also on the links between the actors responsible for them.

3.4.4 Module 3: Alignment assessment across links

The third module applies alignment conditions to the links between positions. The four dimensions introduced earlier, technical alignment, data alignment, operational alignment, and incentive and governance alignment, are used as practical diagnostic categories. They ask whether systems can connect, whether information can be interpreted and shared, whether outputs fit store routines, and whether actors have the authority and motivation to make the link work. In Chapter 4, these dimensions become sensitising concepts for interview questions and coding.

Figure 3.4 shows the alignment conditions used to assess whether links between actors can work in practice. These conditions are technical, data-related, operational, and governance-related.

3.4.5 Module 4: Risk overlay with co-innovation risk and adoption-chain-risk

The fourth module adds ecosystem risk to the activity and link structure. Co-innovation risk is used to examine whether required complements are mature and sufficiently integrated for actionable shelf insight to work beyond a controlled pilot (Adner & Kapoor, 2010). Adoption-chain risk is used to examine whether progression from pilot to scale breaks down because different actors must approve, implement, use, renew, or expand the solution at different stages (Adner & Kapoor, 2010).

In this framework, both risks are treated as overlays on the same activity and link structure. They are not separate theories running beside the framework. Co-innovation risk focuses on complement maturity and integration readiness. Adoption-chain risk focuses on the five enterprise retail gates used in this thesis: pilot, rollout

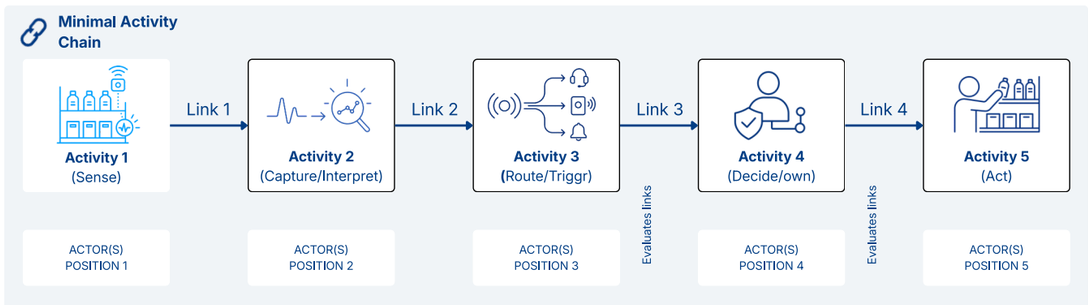

Figure 3.3. Basic activity chain from shelf sensing to store action and ctor positions and dependency links in the activity chain



Figure 3.4. Alignment conditions for actionable shelf insight

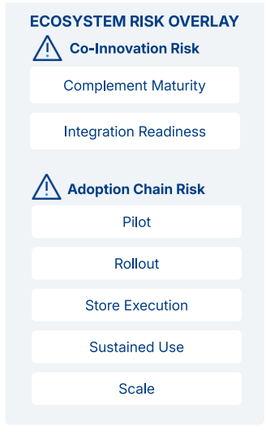

Figure 3.5. Ecosystem risks across complements and adoption gates

decision, store execution and routinised use, sustained use and renewal, and scale. This gate structure is an analytical device for this study, informed by diffusion logic, not a staging directly taken from Rogers (2003).

Figure 3.5 adds the two ecosystem risk lenses used in this thesis. Co-innovation risk concerns required complements, while adoption-chain risk concerns the actors and gates needed for use and scale.

3.4.6 Module 5: Bottlenecks and leverage points with governance feasibility

The fifth module uses the link diagnosis to identify bottlenecks and possible leverage points. A bottleneck is a position or link that constrains actionable shelf insight. A leverage point is a bottleneck or dependency where action could improve the wider ecosystem logic.

A possible leverage point is only treated as meaningful when three conditions are met. It must be binding for value realisation, it must have an identifiable actor or small group of actors that can influence it, and it must be addressable through plausible governance, partnership, or capability mechanisms. This prevents the later analysis from treating every issue as equally strategic.

Figure 3.7 shows how bottlenecks and leverage points are identified. A point of influence is only meaningful when it constrains value realisation and can plausibly be addressed by one or more actors.

The core theoretical lens of this study is ecosystem-as-structure. It is used to define the focal value proposition, map the activity structure, identify actor positions, and analyse the links that must align for actionable shelf insight to be realised (Adner, 2017). This lens provides the structural logic of the thesis. It explains where dependencies are located and why misalignment between actors can prevent value realisation, even when the sensing component performs adequately in isolation.

3.4.7 Module 6: Internal readiness through dynamic capabilities

The sixth module adds MOOS' internal readiness as a feasibility layer. This is included because ecosystem positioning is not only about where MOOS could sit in the activity structure. It is also about whether MOOS can execute and sustain that position. Dynamic capabilities are used in a limited way through sensing, seizing, and transforming (Teece, 2007).

In this framework, sensing refers to MOOS's ability to understand ecosystem bottlenecks, decision ownership, and external dependency conditions. Seizing refers to its ability to make choices about where and how to participate in the activity structure. Transforming refers to its ability to adapt its capabilities and routines as learning accumulates. Internal readiness is therefore not a second research question. It is a feasibility constraint that becomes important in Chapter 6, where positioning options are translated into roadmap logic.

Figure 3.7 adds MOOS' internal readiness. This layer is used to assess whether MOOS can sense ecosystem conditions, make positioning choices, and adapt its capabilities over time.

Figure 3.8 summarises the diagnostic outputs that are carried forward into the design phase. These outputs include candidate positions, critical links, risks, bottlenecks, and alternative pathways

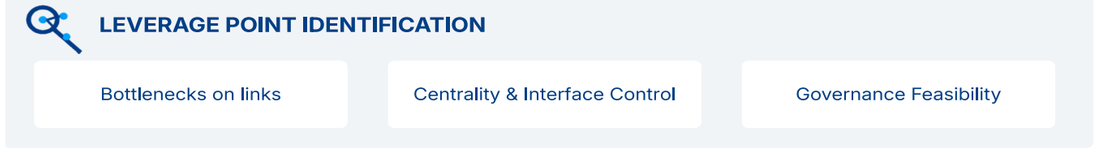

Figure 3.6. Bottlenecks and possible leverage points in the ecosystem structure



Figure 3.7. Internal readiness as a feasibility layer for MOOS

3.4.8 Using the framework for empirical work and positioning option synthesis

The framework guides the empirical work in three ways. In Chapter 4, the activity structure and actor positions inform sampling, because the interviews need to cover roles that are structurally relevant to actionable shelf insight. The alignment dimensions and risk overlays inform the interview guide and coding scheme, because each critical link can be examined in terms of technical, data, operational, incentive and governance alignment, co-innovation risk, and adoption-chain risk.

In Chapter 5, the same framework structures the empirical findings. Store-level evidence shows which actionability conditions matter in practice, while expert interviews show where link-level dependencies, bottlenecks, risk patterns, and ownership problems appear across actor positions. The framework is therefore not only a summary of Chapter 3. It is the analytical structure used to interpret the empirical material.

In Chapter 6, the framework supports positioning synthesis. It does not produce positioning options by itself. This separation matters because positioning should be an output of diagnosis, not an assumption built into the framework. Possible positions for MOOS are therefore derived after the empirical analysis by comparing roles in the activity structure, critical dependencies, alignment exposure, risk patterns, points of influence, and internal readiness. Chapter 4 explains how this framework is translated into the research design, interview protocol, and analysis approach.



Figure 3.8. Diagnostic outputs for positioning synthesis

Chapter 4

Methodology

4.1 Case context and use-case focus
4.2 Research process overview
4.3 Data collection Part 1: structured in-store visits with ranking exercise
4.4 Data collection Part 2: semi-structured expert interviews
4.5 Ethics, consent, and confidentiality
4.6 Data preparation
4.7 Data analysis
4.8 Synthesis of Part 1 and Part 2
4.9 AI-assisted analysis support
4.10 Co-creation as part of the design process
4.11 Validation logic
4.12 Methodological limitations

4. Methodology

This chapter explains how the research was carried out. It shows how the conceptual framework from Chapter 3 was used to guide data collection, analysis, synthesis, design translation, and validation. The chapter does not present the findings themselves. These are discussed in Chapter 5.

The methodology followed a qualitative and exploratory logic. This fit the research question because MOOS's positioning problem is not only about the technical performance of shelf sensing. The value of shelf sensing depends on whether shelf events can be turned into trusted, timely, workflow-connected and accountable store action. This requires insight into store routines, actor roles, system links, partner dependencies, and adoption barriers.

The empirical work was structured in two connected parts. Part 1 focused on store-level adoption thresholds through 41 structured in-store visits with a ranking exercise and short contextual conversations. Part 2 focused on ecosystem-level conditions through eight semi-structured expert interviews. The two datasets were first analysed separately and then brought together through the conceptual framework. This sequence helped keep point-of-use requirements separate from wider ecosystem conditions before comparing them.

4.1 Case context and use-case focus

MOOS develops shelf-sensing technology that detects shelf events and converts them into signals and insights for retail operations. In this thesis, the value proposition is examined through two main use cases. The first concerns on-shelf availability and OOS situations. The second concerns theft, LP, and shrink. These use cases are related because both start from shelf-level visibility, but they do not necessarily follow the same route from signal to action.

This distinction shaped the empirical design. Availability-related signals usually connect to routines around stock, replenishment, and store execution. Loss-prevention-related signals often involve different timing, accountability, and escalation conditions. For that reason, the research did not assume that one positioning route would fit both use cases equally well. Both use cases were kept visible during data collection and analysis, so that differences in timing, ownership, system dependency, and operational follow-up could be studied rather than hidden inside one broad shelf-sensing category.

The case context also affected the level of analysis. MOOS is not treated as a retailer, a full store operating system, or a general retail platform. It is treated as a component provider that may need to connect with other actors to make its value proposition work at scale. This makes the case useful for studying positioning from the perspective of a smaller technology firm that depends on other actors for adoption and value delivery.

4.2 Research process overview

The research process combined a literature-informed framework, two empirical parts, a synthesis step, a co-creation step, design development, and conducted validation. The literature review first led to a version-one diagnostic framework. This framework was not used as a fixed model that had to be confirmed. It was used as an analytical lens for the empirical work. It helped define what needed to be examined: the focal value proposition, the activity chain, actor positions, links between actors, alignment conditions, ecosystem risks, and internal readiness.

This research design followed a qualitative and exploratory logic. This was appropriate because MOOS's positioning problem is not only a product question. It is an ecosystem question. Shelf sensing only becomes valuable when signals lead to action through store routines, systems, partners, and accountable actors. A design that only studied the sensor, or only studied store users, would have missed these dependencies. The multi-actor design made it possible to compare the point-of-use logic in stores with the wider system logic needed to make store action possible at scale.

The empirical phase started with structured in-store visits. These visits were used to identify store-level adoption requirements through a ranking exercise and short contextual conversations. After that, semi-structured expert interviews were conducted to examine ecosystem-level conditions around feasibility, integration, accountability, adoption gates, and partner routes. The two datasets were first analysed separately. They were then brought together in synthesis to identify where store-level needs and ecosystem-level conditions reinforced each other, and where tensions appeared.

This sequence was chosen deliberately. Part 1 showed what store teams appeared to require before shelf sensing becomes useful in daily work. Part 2 examined what actors, systems, governance arrangements, and integration choices are needed to deliver those requirements across stores and partners. Keeping these two levels separate at first helped avoid an early collapse of store logic into ecosystem logic. Only after this separation were the two datasets compared through the conceptual framework.

The synthesis formed the bridge to the design phase. It did not yet produce the final roadmap. Instead, it translated the empirical findings into design inputs, including adoption thresholds, activity-chain gaps, actor dependencies, and positioning tensions. After this synthesis, a co-creation session was used as a research-based design intervention. This session helped discuss the implications of the findings and explore possible positioning directions. The final roadmap was then developed and later validated. The validation was conducted after the design phase and is described in Chapter 7.

Figure 4.1 shows how the methodology moved from the conceptual framework to empirical data collection, separate analysis, synthesis, design translation, and validation. This sequence matters because the thesis first diagnoses the ecosystem problem before developing and testing the roadmap.

4.3 Data collection Part 1: structured in-store visits with ranking exercise

The first empirical part consisted of structured in-store visits with 41 store operators and store managers. The visits were conducted over approximately two weeks in the Netherlands. The geographical scope included urban and less urban retail locations in the Netherlands. This created variation between urban and less urban retail settings, although the sample remained limited to stores that could be reached within the available time.

Part 1 was designed to identify store-level adoption requirements for shelf-sensing technology. The aim was practical and focused. It was not an ethnographic study and it was not a field experiment. A more accurate description is structured in-store visits with a ranking exercise and short contextual conversations. During the visits, attention was paid to visible shelf processes, store routines, and practical issues around theft, shrink, out-of-stock situations, and replenishment. These observations helped interpret the responses, but the ranking exercise was the primary structured data source.

Respondents were asked to rank eight possible adoption conditions from 1 to 8, where 1 meant most important and 8 meant least important. The eight conditions were: reliable and trustworthy signals, fast enough to act upon, no additional operational burden, works with existing store technology, low effort to keep running, fits current ways of working, clear ownership, and hardly visible in the store. Respondents did not select use cases in a formal coded task. They ranked adoption conditions and sometimes explained their choices through short examples. When theft or OOS situations were discussed, these were treated as contextual explanations rather than as closed-choice use-case data. The question list used during the store visits is provided in Appendix C.2.

The ranking exercise was recorded on prepared paper field sheets. These sheets also included space for store descriptors and short notes. The same field sheet was used across visits to keep the exercise comparable. The full instrument **will** be

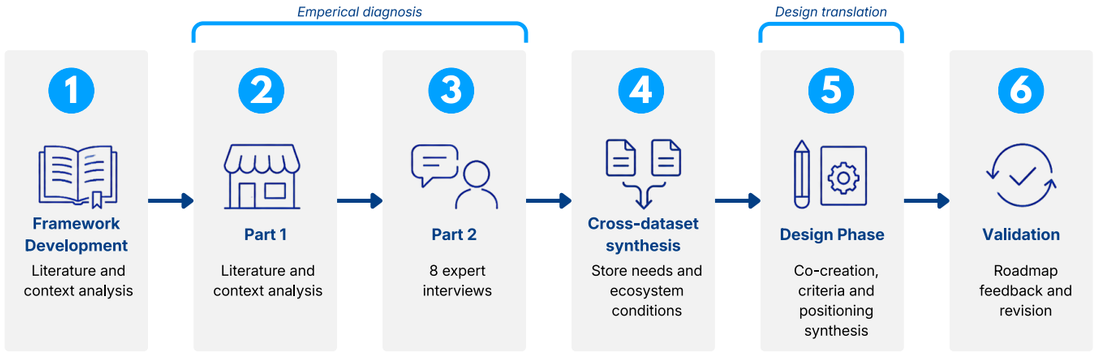

Figure 4.1. Research process overview from framework development to conducted validation.

included in the appendix, because the exact wording of the ranking labels matters for interpreting the results. Short comments were written down when respondents explained why they ranked certain conditions high or low. These comments were not analysed as a full qualitative interview dataset. They were used as contextual support for interpreting the ranking patterns.

The sampling strategy was purposive and partly opportunistic. Stores were approached directly, and the acceptance rate was roughly one in three or one in four stores. The aim was not statistical representation. The aim was to create useful variation across relevant store contexts and to identify practical thresholds for store-level adoption. The sample included supermarkets, drugstores, and health and beauty stores. It also included variation between branch-managed and franchise-like settings. A further distinction was made between lower-tech and higher-tech stores. This label was based on the researcher's in-store impression, the respondent's explanation, and the visible use of store technology. It should be read as an interpretive classification, not as a formal technology audit.

The value of Part 1 lies in its point-of-use perspective. The ranking exercise forced respondents to prioritise between adoption conditions instead of treating every condition as equally important. With 41 visits, the dataset was useful for exploratory store-level threshold identification. It helped show which conditions store teams appeared to treat as basic requirements, and which issues seemed less immediate in daily work. The dataset does not support statistical claims about all retail stores. It supports an exploratory diagnosis of what store teams may require before shelf sensing becomes acceptable in daily practice. The store visit protocol, question list, anonymised sample overview, and aggregated store visit result material are provided in Appendix C.

4.4 Data collection Part 2: semi-structured expert interviews

The second empirical part consisted of eight semi-structured expert interviews. These interviews were conducted over approximately one week. All interviews took place online through Microsoft Teams. The respondents were purposively selected to cover different positions in the retail and retail technology ecosystem around shelf sensing.

The sample was designed for analytical ecosystem diagnosis, not for statistical representation. Eight interviews were acceptable for this purpose because the aim was to compare role-based perspectives on the main dependencies around shelf sensing. The interviews were not used to measure how common a view is across the retail market. They were used to understand how different actors reason about feasibility, integration, accountability, adoption gates, and external dependency conditions. For this thesis, the spread of actor perspectives was more important than the number of interviews alone.

The interview set included perspectives from retail or head office, store innovation or Group IT, consultants or system integrators, shelf or display infrastructure, and technology or platform-side actors. These clusters were relevant because each controls or influences a different part of the route from shelf signal to store action. Retail actors can explain operational priorities and internal decision-making. IT and innovation actors can explain system fit, data governance, and security constraints. Integrators can explain rollout and implementation barriers. Shelf and display infrastructure actors can explain the physical store layer. Technology and platform actors can explain how shelf signals may connect to other store systems. The interview protocol is provided in Appendix D.

Included actor perspective	*Role in ecosystem diagnosis*	*Less visible or missing perspectives*	*Implication for limitations*
Retail or head-office perspective	Explains retail priorities, internal decision-making and adoption logic.	Store staff were covered mainly in Part 1, not in Part 2.	Part 2 gives a broader organisational view, not a second store-level dataset.
Store innovation or Group IT	Explains system fit, data governance, security and integration constraints.	Legal and privacy roles were not directly represented.	Privacy and compliance findings should be treated as expert reflections, not legal assessment.
Consultants or system integrators	Explains rollout barriers, integration work and implementation risk.	Large enterprise software vendors were less visible.	The analysis may underrepresent the logic of major platform owners.
Shelf, display or store infrastructure actors	Explains the physical store layer and deployment conditions.	Logistics and supply chain actors were less visible.	Backend replenishment and supply chain closure may need further validation.
Technology or platform-side actors	Explains possible connections between shelf signals and store systems.	Procurement roles were not directly represented.	Buying-centre logic is discussed, but not fully mapped from a procurement perspective.

Table 4.1. Actor perspectives included and less visible in the expert interview sample.

Table 4.1 gives an anonymised overview of the actor perspectives included in the expert interviews. The table is not meant to show a representative sample of the retail market. It shows which actor positions were covered and which positions were less visible in the dataset.

The sample also had limits. Procurement, legal, privacy, large enterprise software vendors, and some logistics or supply chain roles were less visible or absent. This means the interviews give analytical insight into the main dependencies around shelf sensing, but they do not cover the full buying centre of every retailer. This limitation is returned to in the discussion chapter.

Part 2 was used to understand ecosystem-level feasibility. While Part 1 focused on what stores require, Part 2 examined what is needed to make those requirements possible across multiple stores, organisational functions, and system categories. The interviews explored shelf-related retail problems, what makes shelf information actionable, the minimum activity chain from signal to action, adoption gates, accountability, ownership, integration depth, external dependencies, and differences between availability-related and loss-prevention-related use cases. The interviews also explored how scaling decisions are distributed across the retail technology ecosystem.

The interview guide was semi-structured. It was informed by the conceptual framework from Chapter 3, so the interviews were not fully inductive. Questions were linked to concepts such as ecosystem structure, activity alignment, adoption-chain risk, co-innovation risk, and internal readiness. At the same time, the guide left room for examples, disagreement, and unexpected patterns. This balance was important. The interviews were not meant to confirm the framework narrowly. They were meant to test, sharpen, and specify it through expert experience.

4.5 Ethics, consent, and confidentiality

Ethical care was part of both empirical parts, although the procedure differed between Part 1 and Part 2. In Part 1, a simplified form of consent was used. The store visits were short, practical, and often took place during store opening hours. For that setting, a shorter yes-or-no consent procedure was used. Respondents were informed about the purpose of the research and could refuse or stop at any moment.

Part 2 used a more extended consent form. This included consent for audio recording, transcription, anonymised analysis, and the use of anonymised quotes in the thesis. The interviews were longer and included more detailed professional reflections, so a more thorough consent procedure was appropriate.

In both parts, participation was voluntary. Names of respondents and identifying store details were not used in the analysis. Respondents were referred to through codes. The data were stored on a protected TU Delft drive, with access limited to the researcher. This was important because some interview material touched on partner relations, internal decision processes, and adoption barriers that could be sensitive in a business context.

The approved project brief and the Human Research Ethics Committee approval are included in Appendix A. The interview consent form template is included in Appendix D.3.

4.6 Data preparation

The Part 1 material was prepared as a structured ranking dataset. The paper field sheets were used to record the rankings, store descriptors, and short comments. After the visits, the data were transferred into a spreadsheet format. This made it possible to compare average ranking patterns, top-ranked conditions, bottom-ranked conditions, and segment differences. Store descriptors were used to compare patterns between, for example, supermarkets and drugstores, branch-managed and franchise-like stores, and lower-tech and higher-tech stores.

The open comments from Part 1 were handled carefully. They were not coded as a full qualitative dataset. Instead, they were used to support the interpretation of the ranking results. For example, when a respondent explained why reliable signals mattered, that comment helped clarify the meaning

behind the ranking. The comments were therefore used as contextual evidence, not as the main basis for theme development.

For Part 2, the interviews were recorded in Microsoft Teams. Transcripts were produced using a combination of Microsoft Teams transcription and Google transcription support. The transcripts were then reviewed and anonymised before analysis. They became the primary material for the expert interview analysis. Audio recordings were checked selectively when tone, hesitation, certainty, or emphasis seemed relevant. This was useful because some answers, especially around integration, ownership, or partner dependency, were not only factual. They also expressed uncertainty, caution, or strategic judgement. The data preparation process, cleaning rules, codebook and theme development material are provided in Appendix E.

4.7 Data analysis

The analysis moved from structured comparison to theme development. Part 1 and Part 2 were analysed separately first. This separation was important because the two parts answered different questions. Part 1 showed which adoption conditions store respondents prioritised and how they explained these choices in daily store terms. Part 2 showed what ecosystem actors saw as necessary for delivery, adoption, system alignment, and scaling.

For Part 1, the first step was descriptive analysis of the ranking data. The rankings across all 41 responses were compared to identify which adoption conditions were most and least important overall. After that, segment comparisons were made. These comparisons looked at differences between store formats and organisational types, such as supermarket versus drugstore, branch-managed versus franchise-like, and lower-tech versus higher-tech stores. The final step was interpretive. The ranking patterns were translated into store-level adoption requirements, practical thresholds, and possible rejection triggers.

The Part 1 analysis remained modest in scope. It did not claim to measure adoption statistically. The sample size and sampling method do not support that kind of claim. Instead, the ranking exercise was used to make trade-offs visible. Asking respondents to rank the conditions forced them to prioritise between requirements rather than simply agree that all conditions were important. The short comments were used to clarify the meaning behind the rankings, not to build a separate thematic model.

For Part 2, the analysis followed a step-by-step process. It started with familiarisation. The researcher read the transcripts, reviewed the interview notes, and marked recurring issues in the material. Initial codes were then created manually from the transcripts and notes. Some codes came directly from the data, such as reliability, timing, workload, ownership, integration, adoption gates, partner dependency, theft, and OOS. Other codes were informed by the framework in Chapter 3, such as activity chain, actor position, link, alignment, adoption-chain risk, co-innovation risk, and internal readiness.

The framework concepts were used as sensitising concepts. They helped direct attention to actor roles, activities, links, and dependencies, but they were not used as a fixed codebook. This was important because the interviews also needed to capture examples and tensions that were not fully expected before the analysis.

The first comparison was made question by question across all eight interviews. Answers to the same question were placed together and compared. This made it easier to see where respondents agreed, where they differed, and where differences seemed linked to their ecosystem role. After this comparison, the interviews were checked again as individual cases. This case-level check helped protect the internal logic of each respondent's explanation.

Initial codes were then grouped into broader categories, including signal reliability, workflow actionability, integration depth, accountability, adoption gates, partner routes, and use-case differences. These categories were refined into higher-level themes. Some early themes were split because they covered too many issues. Others were merged because they described the same pattern from different angles.

The final themes were checked against the original notes and transcripts before being used in Chapter 5. When a theme could not be supported clearly by the source material, it was narrowed, renamed, or removed. This helped keep the findings close to the data and reduced the risk of forcing the interviews into the framework.

4.8 Synthesis of Part 1 and Part 2

Part 1 and Part 2 were synthesised only after they had been analysed separately. This sequence was important. It prevented store-level requirements from being absorbed too quickly into ecosystem-level interpretation, and it kept point-of-use actionability analytically separate from scale feasibility. The synthesis then compared the two levels of evidence through the conceptual framework from Chapter 3.

The synthesis focused on two questions. First, where did the expert interviews help explain or support the store-level thresholds from Part 1? Second, where did the expert interviews show that a store-level requirement would be difficult to deliver across the wider ecosystem? This kept the synthesis close to the research question without turning it into a new data collection step.

The comparison was made through the conceptual framework from Chapter 3. Store-level requirements from Part 1 were compared with activity chains, actor roles, links, alignment conditions, adoption gates, and use-case differences from Part 2. This made it possible to translate point-of-use requirements into ecosystem-level conditions.

The synthesis was therefore not a simple combination of two datasets. It compared two levels of the same positioning problem. Part 1 showed what store teams appeared to require. Part 2 explained what the wider retail system would need to organise to deliver those requirements. Together, the two parts formed the empirical basis for the design phase by identifying adoption thresholds, bottlenecks, actor dependencies, and positioning tensions at an analytical level.

4.9 AI-assisted analysis support

NotebookLM was used as a supporting tool during the analysis. Its role was limited to organising material, summarising sources, comparing passages, and retrieving relevant excerpts. It was not used as an autonomous coder. It did not create the final code structure, define final themes, interpret the findings, select evidence, or decide design directions.

For Part 2, anonymised transcripts, the interview guide, and relevant analytical notes were uploaded to NotebookLM. The tool helped create structured summaries of each interview, including the respondent's role, main claims, examples, bottlenecks, integration logic, and references to theft or OOS. These summaries made the material easier to navigate, but they were checked against the transcripts before being used.

NotebookLM was also used to support cross-interview comparison. For each interview question, it helped retrieve relevant answers across respondents and place related passages next to each other. The researcher then compared these passages manually to identify patterns, differences, and role-based explanations. After that, the interviews were checked again as whole cases to avoid losing the internal logic of each respondent.

During theme development, NotebookLM helped retrieve supporting excerpts around topics such as shelf visibility, store execution, buying-centre fragmentation, system alignment, and external dependency conditions. These excerpts were not accepted automatically. They were checked against the transcripts and notes before being used. During synthesis, NotebookLM helped organise material related to Part 1 store priorities and Part 2 ecosystem conditions, but the synthesis choices were made by the researcher.

The researcher remained responsible for the analysis. The researcher created and refined the codes, grouped them into categories, named the themes, interpreted contradictions, selected evidence, and decided how the findings were synthesised. Outputs from NotebookLM were only used after checking them against the original notes and transcripts. This check was needed because AI-generated summaries can overstate patterns, miss context, or simplify disagreement. NotebookLM therefore supported organisation and retrieval, but the analysis itself remained researcher-led.

4.10 Co-creation as part of the design process

After the empirical synthesis, the project moved into a participatory design step. A small co-creation session was included before the final roadmap was developed. This step was used to discuss the implications of the findings with relevant stakeholders and to explore possible positioning directions for MOOS.

Co-creation was treated as part of the design process, not as a third empirical evidence layer. Its

purpose was to work with the findings, not to replace them. The session helped surface design trade-offs, role-boundary questions, and criteria for the roadmap. The final roadmap should therefore not be read as the direct outcome of the co-creation session alone. It was developed from the conceptual framework, Part 1, Part 2, the synthesis, co-creation input, and further design reasoning. A sanitised overview of the co-creation protocol and non-confidential design materials is provided in Appendix F.

4.11 Validation logic

Validation was conducted after the design phase. It was not used as a substitute for co-creation. Co-creation helped discuss and shape possible positioning directions before the roadmap was finalised. Validation was later used to assess the roadmap's clarity, relevance, credibility, and usefulness for the intended context.

This sequence fits the logic of the thesis. The empirical parts first generated evidence. The synthesis translated this evidence into design inputs. Co-creation helped discuss these inputs and sharpen the design direction. The design phase then developed the roadmap. After that, validation tested the roadmap logic and informed revisions.

The validation method and findings are described in Chapter 7. In this methodology chapter, the main point is that validation had a different role from the empirical research and co-creation. It did not generate the main diagnosis. It assessed the designed roadmap after it had been developed. A public-safe overview of the validation setup, criteria, and non-confidential revision logic is provided in Appendix H.

4.12 Methodological limitations

The empirical material should be read as exploratory and analytical rather than statistically representative. The 41 in-store visits helped identify store-level priorities and thresholds, but they do not represent all retail formats. The eight expert interviews gave insight into the main actor dependencies around shelf sensing, but they did not include every actor in the retail buying centre. Some perspectives, such as procurement, legal, privacy, large enterprise software vendors, and logistics, were less visible. These limits do not invalidate the findings, but they define how they should be used. The results are suitable for diagnosing patterns, tensions, and design requirements. They are not suitable for measuring how common these patterns are across the whole retail market. A fuller reflection on these limitations is provided in Chapter 8.6, where the empirical scope, design process, and validation approach are discussed in relation to the overall thesis contribution.

Chapter 5

Results

5.1 Store level findings
5.2 Ecosystem-level findings
5.3 Cross-dataset synthesis
5.4 Framework refinement

5. Results

This chapter presents the empirical results of the study. It first reports the store level findings from the 41 store visits. It then reports the ecosystem level findings from the eight expert interviews. The final sections compare both datasets and use them to refine the diagnostic framework from Chapter 3. The chapter therefore stays close to the evidence. It does not yet choose a final positioning direction for MOOS. The store visit protocol is included in Appendix C.1.

Figure 5.1 shows how the chapter moves from store-level evidence to ecosystem-level evidence, and then to synthesis and framework refinement. The figure is relevant because both datasets are used to build one empirical diagnosis of the positioning problem.

5.1 Store level findings

Before presenting the store-level findings, it is useful to clarify the empirical setting of Part 1. As can be seen in Table 5.1, the store visits covered 41 retail locations in the Netherlands across urban and smaller non-urban settings. The sample included supermarkets, drugstores, and beauty and health stores. It also included variation between branch-managed stores and more franchise-like settings, as well as stores with different levels of visible technology use. This variation was not meant to create a statistically representative sample. Instead, it was used to capture different store contexts in which shelf sensing could become relevant.

Table 5.1 shows the composition of the store visit sample. It is included to clarify the type of store-level evidence behind the ranking results and open responses.

The store visits focused on how store teams understood the possible value of shelf sensing in daily practice. Respondents were asked to rank eight adoption conditions and explain which store problem they most strongly associated with the technology. These rankings and open explanations were then read together. Together, the rankings and open explanations show which conditions store teams treated as most important for daily use.

The findings are presented in five steps. First, the ranking data are used to identify the strongest store-level conditions. Second, reliability, speed, and workload protection are discussed as the main requirements for practical use. Third, the open responses clarify which use cases were most urgent at store level. Fourth, segment differences are reviewed cautiously. Finally, the lower-ranked conditions are considered as possible implementation blind spots.

Table 5.2 shows the average ranking of the eight adoption attributes. Lower average ranks indicate higher priority. The table is used to identify which attributes store respondents treated as the most important conditions for daily use, not to make statistically generalisable claims.

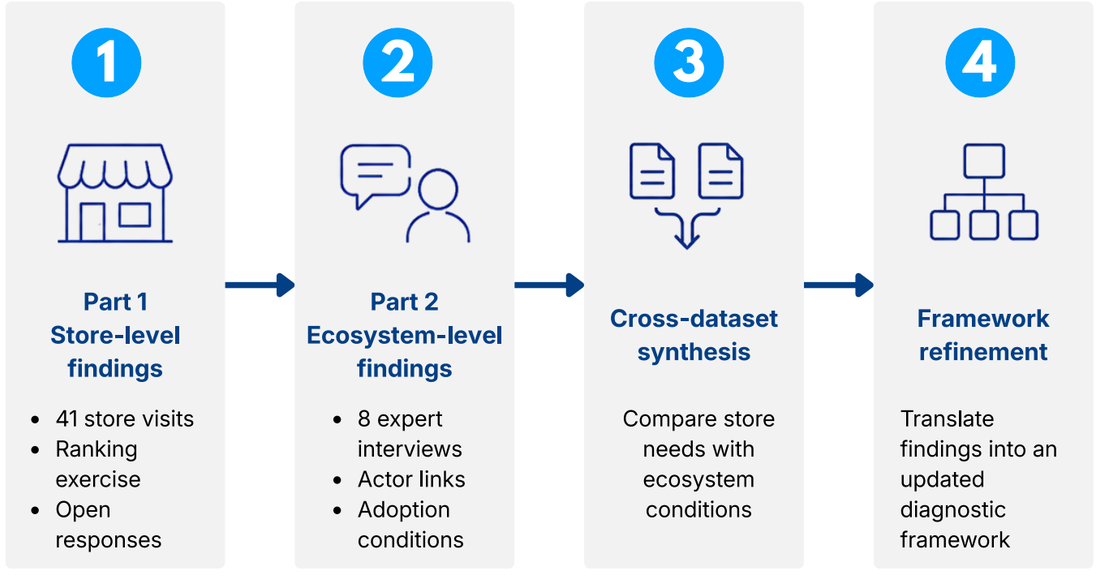

Figure 5.1, Flow diagram: Part 1 store visits → Part 2 expert interviews → cross dataset synthesis → framework refinement.

5.1.1 Store adoption follows threshold logic rather than feature preference

As shown in Table 5.2, reliable signals ranked first overall, with an average rank of 2.73. Fast enough to act ranked second, with an average rank of 3.98. Low effort to keep it running, no extra staff work, and works with existing store technology followed after that, with average ranks close to each other. By contrast, hard to notice in store and clear ownership ranked lowest. The exact differences between the middle-ranked attributes should not be overstated. The main pattern is that store teams first prioritised trust in the signal, then timely action, and then low burden in daily work.

The open explanations make this threshold structure easier to understand. Respondents repeatedly described the desired solution in practical terms: it should be trustworthy, quick enough to trigger a useful response, and invisible in terms of extra burden. When that basic test was not met, later benefits did not seem to matter. A system that might improve performance in theory was still viewed negatively if it produced false alerts, required babysitting, or forced staff into an extra workflow.

This pattern gives the first store-level explanation for why shelf sensing may fail in use. A technically useful signal still has to be trusted, timely, and easy to work with in the store routine.

5.1.2 Reliability functions as the main gatekeeper

Reliability is the clearest threshold in the store data. It ranked highest and also shaped how the other attributes should be read. Store staff first need to trust that the signal is correct. If they do not trust the signal, speed, integration, or other benefits become less relevant.

The comments support this strongly. Respondents warned that even minor errors or early “teething problems” could be enough to make the system irrelevant on the shop floor. The issue was not that stores expected perfection in an abstract sense. Rather, they expected the signal to be dependable enough to fit into a busy retail routine without constant doubt. If staff need to second-guess whether an alert is real, the system becomes one more distraction. It no longer reduces uncertainty. It adds it.

That also helps explain why some other attributes ranked lower than might be expected. Store operators seem willing to tolerate a certain amount of visible hardware, some organisational ambiguity, and even some change in process, but only after the reliability threshold has been met. In other words, reliability is not one preference among many. It sits first in the sequence. Other

Dimension	***Categories in sample***	***Count***
Store type	Supermarket	26
	Drugstore	10
	Beauty and health	5
Store size	Small	14
	Medium	19
	Large	8
Governance	Branch-managed	29
	Franchise-like	8
	Independent	4
Location type	Urban	32
	Non-urban / smaller-town	9

Table 5.1: overview of the store visit sample, showing store sizes, categories, location types, sample distinctions and tech level.

Overall order	*Charactersitic*	*Avg. Rank*
1	Reliable Signals	2,73
2	Fast enough to react	3,98
3	Low effort to keep it running	4,12
4	Works together with existing store technology	4,27
5	No extra work for staff	4,32
6	Fits in the current way of working	4,73
7	Hardly visible in store	5,88
8	Clear ownership	6,00

Table 5.2. Average ranking of store adoption attributes. Lower average ranks indicate higher priority.

concerns are often postponed until the signal itself has earned trust.

For the store respondents, reliability was therefore practical rather than abstract. The signal had to be dependable enough for staff to act without first checking whether the alert was real.

5.1.3 Speed and workload protection define practical actionability

Speed ranked second, but the comments suggest that stores did not mean generic real-time performance. What mattered was whether the signal arrived early enough for a meaningful response. In theft-related situations, that often meant while the event was still live or while there was still a realistic chance to intervene. In availability-related situations, the time logic was slightly different, but still practical. A late signal was often treated as a missed chance rather than useful information.

Workload protection formed the third part of the threshold structure. Both no extra staff work and low effort to keep it running ranked high, and the comments around them were very consistent. Respondents did not want to maintain another device, manage another app, or add another task to an already crowded shift. One store comment summarised this very directly: *"another app is a no go."* That remark captures the broader pattern well. The store team did not reject technology as such. It rejected additional burden that was not clearly absorbed into existing work.

This is also where integration first appears in the store material, but only in a narrow, practical sense. Works with existing store technology was valued because it reduced friction, not because stores were thinking in architectural terms. A system that fits current routines is easier to trust and easier to keep using. A system that depends on separate maintenance, separate devices, or visible supervision quickly starts to look like extra labour.

Figure 5.2 summarises the threshold pattern from the store visits. The figure matters because it shows that store teams did not judge shelf sensing only by the alert itself. They also judged whether the alert could be used without adding work.

5.1.4 The open responses make the use case hierarchy more concrete

The ranking data becomes more meaningful when read alongside the open responses about which problem the sensor would primarily address in practice. These were not closed-choice items. Respondents stated, in their own words, the main problem they associated with the technology, and these statements were later grouped into use case categories.

Figure 5.3 shows how the open responses were grouped into use case categories. The figure is relevant because it shows that theft, loss, and shrink were the most visible store-level reasons for sensor interest, while OOS formed a second demand cluster.

Loss, shrink, and theft formed the strongest cluster. Out of 41 respondents, 26 stated loss or shrinkage as the primary reason for needing the sensor, and another 6 mentioned it in combination with availability. This made shrinkage or theft the most frequently stated reason in the store responses. The comments on this cluster repeatedly mentioned rapid product removals in high-risk, easy-to-remove, or high-value product categories. The pain was described in two linked ways: direct financial loss and the empty shelf that remains for paying customers afterwards. In that sense, theft was not only a security problem. It also created a secondary availability problem.

Availability and OOS formed the second cluster. 8 respondents stated availability as the sole main reason for needing the sensor, and another 6 mentioned it together with shrink. Here, the pain looked different. The comments referred more often to missed sales, customer frustration, and manual checking routines. Several respondents

Table 5.2: Table 5.2 shows the average ranking of eight adoption attributes. Lower average ranks indicate higher priority.

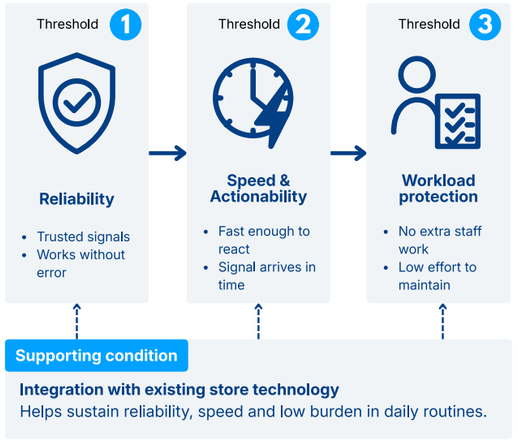

Figure 5.2: Store adoption threshold model, showing reliable signals first, fast enough to act second, and workload protection third.

described staff making repeated rounds to check whether shelves were still properly filled. Compared with theft, this cluster appeared less immediate in the store comments, but it still pointed to repeated operational work.

A third, smaller group of responses pointed to inventory accuracy. These responses described a mismatch between system records and shelf reality. This was less visible than theft or empty shelves, but it helps explain why the expert interviews later paid more attention to data, systems, and replenishment logic. The store responses already showed a practical difference between the two main use cases. Theft was linked to rapid response. Availability and OOS were linked more often to checking, correction, and replenishment.

5.1.5 Segment differences reveal different store logics

The store-level priorities were not identical across all settings. The sample was not designed for statistical comparison, so these differences should be read only as exploratory patterns. They are useful as context, but they should not be treated as firm segment conclusions.

As can be seen inFigure 5.4 in the page below, a possible difference appeared between branch managed settings and more franchise like or independently operated stores. Branch contexts seemed to place slightly more emphasis on fit with existing systems and operational consistency. Franchise and smaller-format settings seemed somewhat more sensitive to maintenance burden and direct usability. This pattern should be read

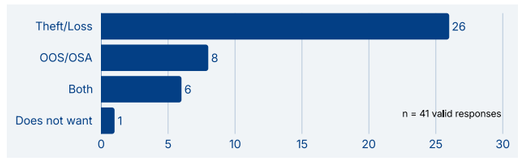

Figure 5.3, main reasons for interest in the sensor

cautiously, because the subgroup sizes were limited.

A similar exploratory pattern appeared between higher-tech and lower-tech stores could be seen in Figure 5.5. Higher-tech stores seemed somewhat more open to system fit, while lower-tech stores seemed to value direct use with little dependence on surrounding infrastructure. This pattern is suggestive only.

The segment differences are therefore mainly useful as a caution. Store context may shape which condition is emphasised most, but the overall pattern remains stable: reliability, timely action, and low burden remain the strongest store-level requirements.

5.1.6 Low-ranked attributes reveal implementation blind spots

The lowest-ranked attributes are also informative. Hard to notice in the store, and clear ownership was consistently placed near the bottom. A simple reading would be that these issues do not matter much. The store comments suggest a more careful interpretation.

Visibility appears to matter less because staff care more about whether the system works than whether it is aesthetically subtle. Ownership appears to matter less because respondents often assume that budgets, responsibilities, and maintenance arrangements will be organised elsewhere. That does not mean ownership is unimportant. It may mean that ownership is less visible to the store user until the system starts failing or creating extra work.

This is where the idea of a blind spot becomes useful. The ranking suggests that stores focus on what they directly experience during daily use. Hidden coordination conditions, such as who maintains configuration quality, who has authority to correct problems, or who is responsible when the system underperforms, are often treated as background issues. They only move into view when something breaks.

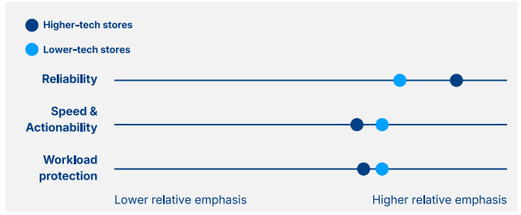

Figure 5.4: compares branch-managed and franchise-like stores on average ranks for the main adoption attributes.

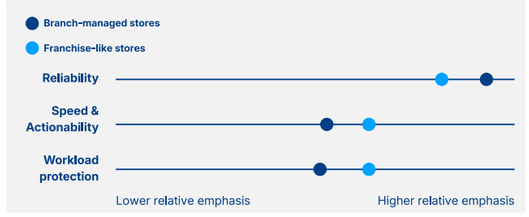

Figure 5.5: compares higher-tech and lower-tech stores on average ranks for the main adoption attributes.

Low priority at store level should therefore not be read as proof that ownership is unimportant. It may show that ownership sits outside the store's direct field of control. This issue returns in the expert interviews.

5.1.7 Conclusion

Part 1 shows that shelf sensing is judged at the store level through a narrow operational lens. Store teams do not first ask whether the proposition is strategically interesting. They first ask whether it is reliable, quick enough to matter, and easy to live with. The results also show that theft is the most visible local use case, while availability constitutes a second, still substantial, demand cluster.

At the same time, the store material leaves several conditions only partly visible. Store teams could describe what makes shelf sensing useful in daily work, but they were able to judge ownership, integration depth, and longer term governance arrangements. This reflects their position in the system. The next section therefore turns to the expert interviews. These interviews examine the wider conditions that must hold if the store-level requirements identified above are to work beyond a local pilot.

5.2 Ecosystem-level findings

The expert interviews made the conditions behind the store level findings more visible. While the store visits showed what store teams require from shelf sensing in daily use, the interviews helped explain what must be organised around those requirements before they can work beyond a pilot.

Four themes stood out across the interviews. First, respondents described a form of retail blindness, where the retailer's systems and the physical shelf do not always show the same reality. Second, they pointed to a fragile link between signal and action. Third, they described a fragmented buying centre, where different actors hold different decision rights. Fourth, they showed that integration is not one technical choice, but a layered trade off between fast local action and deeper system connection.

The expert interview findings are based on the eight semi-structured interviews described in Table 4.1. In this chapter, respondent codes are used to link findings and selected quotes to broad ecosystem perspectives while preserving anonymity. The purpose is not to compare individual respondents, but to show how different actor perspectives reveal recurring dependencies around actionable shelf insight.

5.2.1 Actionable shelf insight starts with retail blindness, but OOS and theft follow different action chains

Across the interviews, the starting point was usually not the sensor itself. Respondents first described a gap between what a retailer's systems assume and what is physically happening in the store. Participant P7 described this gap directly: *"A retailer does not really know what is lying in the store. Very simply put, they only know what is paid for at the checkout"* (translated, P7). This quote shows the basic problem behind actionable shelf insight. Retailers often have sales data, but not a live view of shelf reality.

For OOS, this gap appeared most clearly as phantom inventory. In this situation, the system still assumes that stock is available, while the shelf is empty. Participant P3 explained the effect as follows: *"When the system thinks there is still stock, but there is nothing on the shelf, the store will not receive new deliveries for that item"* (translated, P3). This makes OOS more than a local shelf issue. It also becomes a backend issue, because replenishment can continue to rely on incorrect information.

The theft pathway followed a different pattern. Respondents described sweeps and other loss events where products leave the shelf without passing through the checkout. Participant P5 gave the example of organised mobile groups that *"swipe the whole shelf"* and then move on to the next store (translated, P5). In this case, the issue is more immediate. The store loses products, the shelf becomes empty, and the stock record may remain wrong.

The two pathways therefore start from the same visibility problem, but they do not require the same response. Theft creates urgency and response risk. A signal only helps if it reaches the right person quickly and if the response is safe and realistic. OOS creates a different risk. Here the signal must support replenishment, ordering, or another backend routine. This distinction matters because it shows that actionable shelf insight is not one generic use case. After detection, OOS and theft require different forms of follow-up.

Figure 5.6 shows how one shared visibility problem branches into two different action chains. The figure is relevant because OOS and theft both start from poor shelf visibility, but require different forms of follow-up after detection.

5.2.2 The main operational failure point is the fragile link between shelf signal and store execution

A second pattern concerned the link between the shelf signal and the action that should follow. Respondents did not mainly question whether shelf signals could be useful. The concern was whether those signals would remain trusted and usable inside the daily rhythm of store work.

For OOS, the weak point often sat before the action. Several experts described planogram changes, shelf resets, assortment changes, and local variation as normal retail conditions. If the sensing logic does not stay aligned with those changes, the signal can lose validity. Participant P1 made this point by arguing that linked logic should refresh when shelf conditions or planograms change. This means that reliability is not only a sensor issue. It also depends on whether the shelf representation remains correct over time.

For theft, the weak point often sat after detection. The signal must be credible, routed well, and acted on by someone who can respond. Participant P7 linked this directly to trust: *"If you buy a car, you assume it is reliable. I cannot say that 10 out of 100 signals are wrong. Why would you still walk the eleventh time?"* (translated, P7). Participant P8

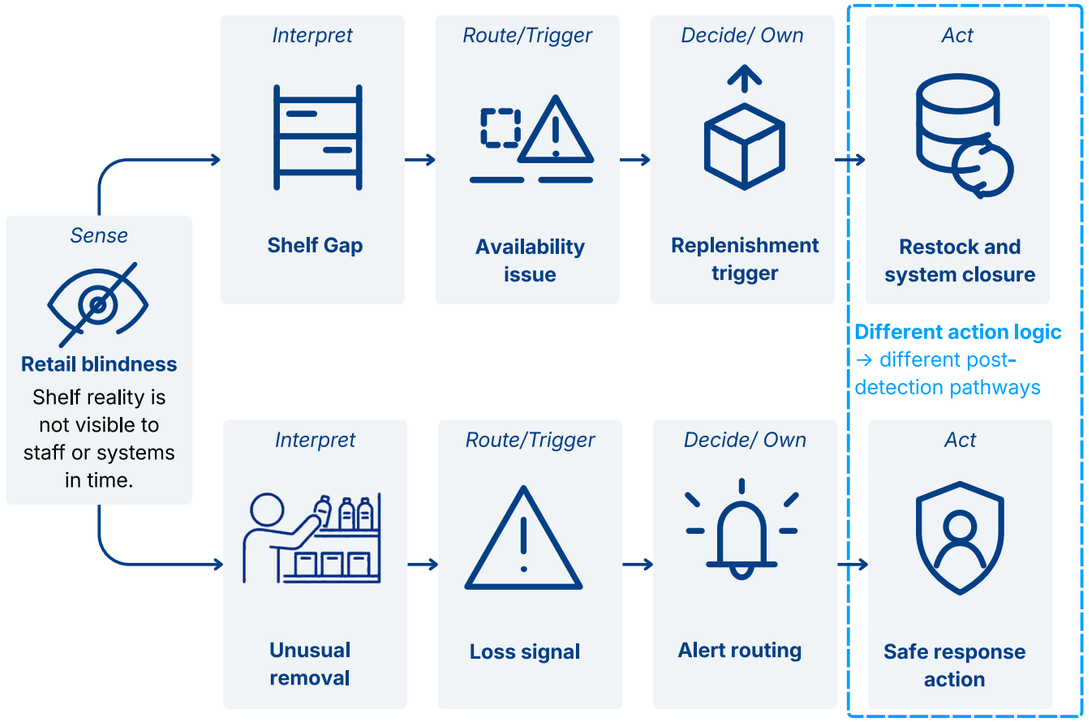

Figure 5.6. Retail blindness and the split between OOS and theft pathways.

described the same risk in daily store language: *"In a busy store, after a week of that beep beep beeping, there is a good chance that store staff are going to ignore it or it just gets lost in the background"* (P8). These quotes show why false or poorly routed alerts can quickly become background noise.

Routing also mattered. Participant P5 described an intervention model in which *"the employee gets the call, but also immediately receives the camera images on the PDA"* (translated, P5). The example shows that the signal alone is not enough. Staff also need context for the response.

The response role was more contested. Several respondents warned against making young or low-authority staff responsible for difficult theft interventions. Participant P6 explained that the output should be usable for young employees: *"It has to be converted into data that, for example, 16 year olds can also do something with, so that they can take quick action"* (translated, P6). Participant P3 gave a sharper safety boundary: *"No one is going to say to a 16 year old girl: a customer is now committing theft, go and stop them"* (translated, P3). These comments show that actionability depends on role, authority, and safety as well as speed.

Figure 5.7 shows where the path from shelf signal to store action can break. The figure is relevant because the interviews showed that detection is only one step. The signal must also be trusted, routed, acted on, and closed. The final learn step refers to feedback from use, such as false alerts, response time, missing action data, and closure data. Without that feedback, the trigger logic cannot easily improve over time.

5.2.3 Scaling depends on a fragmented buying centre rather than on a single decision maker

A third finding is that enterprise adoption is not controlled by one clear owner. The interviews described a fragmented buying centre in which different actors control different gates. Store operations may feel the pain, information technology may control system access, data or procurement roles may affect input quality, and sales or marketing may judge the customer experience impact.

The problem owner is often not the rollout owner. Several respondents described situations in which store operations or business teams recognised the shelf problem, but still depended on other actors to move further. Participant P3 described this as a financial gate, where a clear return on investment is needed before wider adoption can be justified. This shows that the value test at headquarters is different from the usefulness test in the store.

Information technology was one of the strongest recurring gates. Respondents described central IT teams as busy, selective, and often tied to larger system programmes. Participant P6 summarised the issue as a capacity problem rather than a willingness problem. Participant P7 made a similar point when discussing enterprise proof: *"If I can show together with a partner that I trigger this, and you can see in [specific enterprise environment] that the performance of my sensors increased my availability by 1 or 2 percent"* (translated, P7). This quote shows why the proof of value often has to appear inside the retailer's own system environment.

Master data added another layer. Participant P3 warned that incorrect input data, such as wrong product weights, can make the sensing logic fail even if the technology itself works. This matters because the actor who controls input quality is not always the actor who experiences the store-level pain.

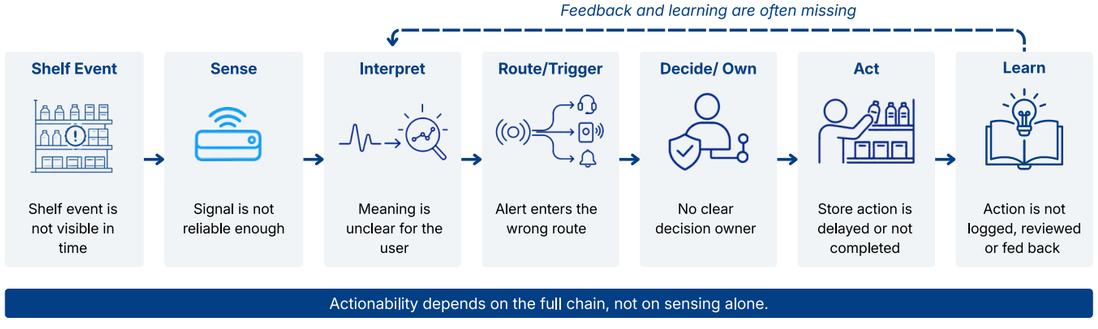

Figure 5.7. Signal-to-action failure map.

For theft related pathways, the veto structure looked different. Participant P5 described anti theft solutions that reduce theft but can also harm customer experience. In the example of protective cabinets, the repeated beeping made people look around and created a negative store experience. The respondent noted that such measures can reduce theft, but may also create a sales drop. This shows that LP is not the only decision logic. Customer experience can also become a gate for adoption.

Support after rollout formed another gate. Participant P8 pointed to the cost of field support: *"If you have to physically send somebody there, very quickly there is the travel time. Now I have got a six hour drive to the next store"* (P8). This quote shows that the operational burden does not end after installation. A solution that works in a pilot can still become difficult to maintain across many stores.

Figure 5.8 shows the actor groups that may influence shelf sensing adoption beyond the store team. The figure is relevant because it shows why a useful store-level signal can still be blocked by IT, procurement, data ownership, finance, customer experience, or service support requirements.

5.2.4 Integration differs by use case and route to scale

The fourth theme concerns integration and route to scale. Respondents did not describe integration as one technical choice. They described different layers, including backend connection, physical embedment, software openness, partner access, and support after rollout.

For OOS, respondents often argued that isolated alerts are not enough. To prove value, the signal must connect to replenishment, ordering, or management systems so that the correction loop can close. Participant P7 linked this to enterprise-system proof proof, where improved sensor performance should become visible in availability performance. Participant P3 made a similar point by stressing that the value must be shown in the retailer's own environment. In the interview material, OOS therefore appeared more dependent on backend closure than theft.

At the same time, respondents warned against building a custom solution for every client. Participant P5 used the image of Lego blocks to explain that technologies should fit together without forcing the retailer into one fixed route. Participant P6 also pointed toward open and repeatable system logic. The practical issue is that too much custom work can make each rollout expensive and hard to repeat.

The theft pathway looked different. Several respondents suggested that a lighter route can be enough when the main value lies in timely routing and response rather than full system closure. In this route, the binding link is not always the enterprise backend. It can also be the alert channel, the responder, and the proof that the intervention model works in the store.

A useful distinction appeared between physical integration and software integration. Respondents seemed more open to physical embedment when it fitted existing shelf infrastructure, store renewal, or installed systems. Software integration was treated more carefully, because it can create lock-in or reduce flexibility later.

Partner routes were also discussed as a possible way to scale. Participants described working

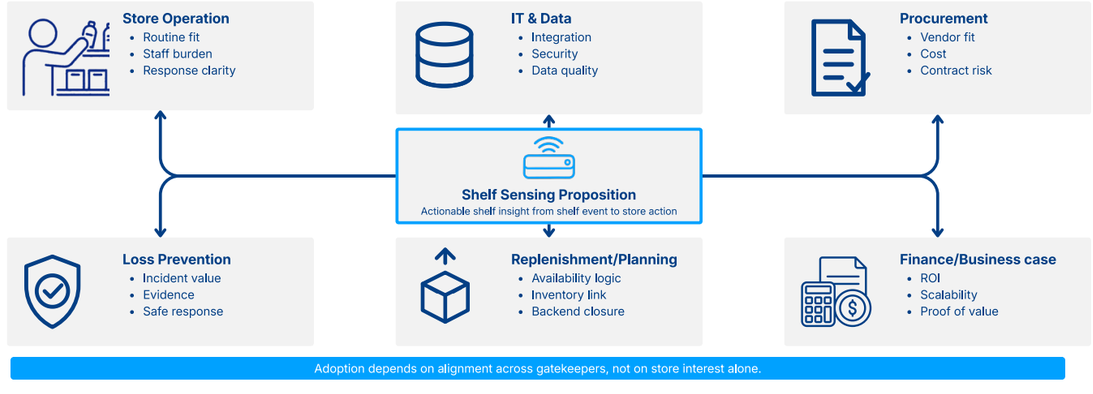

Figure 5.8. Fragmented buying centre around shelf sensing adoption

through existing store infrastructure, planned rollout moments, implementation partners, or adjacent technology categories as routes that may reduce disruption. Participant P4 suggested another option, where external brands help finance a rollout if the data helps their merchandisers: *"Supermarkets are very eager if they do not have to pay for it themselves. We often see that suppliers, large FMCG manufacturers, install it for supermarkets"* (translated, P4). This was not the dominant route in the interviews, but it shows that the financing and access route may vary by actor.

Partner routes also create dependency. Participant P8 warned that a large account can consume too much of a small provider's capacity: *"If you have got 100 leads, which ones do you decide to prioritise? If you are not careful, does a large supermarkt take up so much of the resource that these other medium fish lose momentum?"* (P8). This quote shows the tension clearly. A partner or large client can create access, but it can also create support pressure and dependence.

Figure 5.9 shows the difference between physical integration and software or data integration. The figure is relevant because theft-related action can often start from a narrow alert route, while OSA usually requires stronger backend closure through stock, replenishment, workflow, or planning systems.

5.2.5 Conclusion

The expert interviews show that shelf sensing depends on more than producing a useful signal. The signal must remain trusted, be routed into action, pass through organisational decision gates, and fit the required integration route. Four patterns were most visible: retail blindness, signal-to-action fragility, a fragmented buying centre, and layered integration.

These findings help explain why the store-level requirements from Part 1 cannot be assessed only from the store floor. The next section compares both datasets directly.

5.3 Cross-dataset synthesis

The value of the synthesis lies in comparison rather than repetition. Part 1 showed what store teams require from shelf sensing in daily use. Part 2 showed what the wider retail technology ecosystem must organise to deliver and maintain those requirements. Read separately, both datasets are informative. Read together, they show where actors agree, where their criteria differ, and where local store demands depend on conditions elsewhere in the chain.

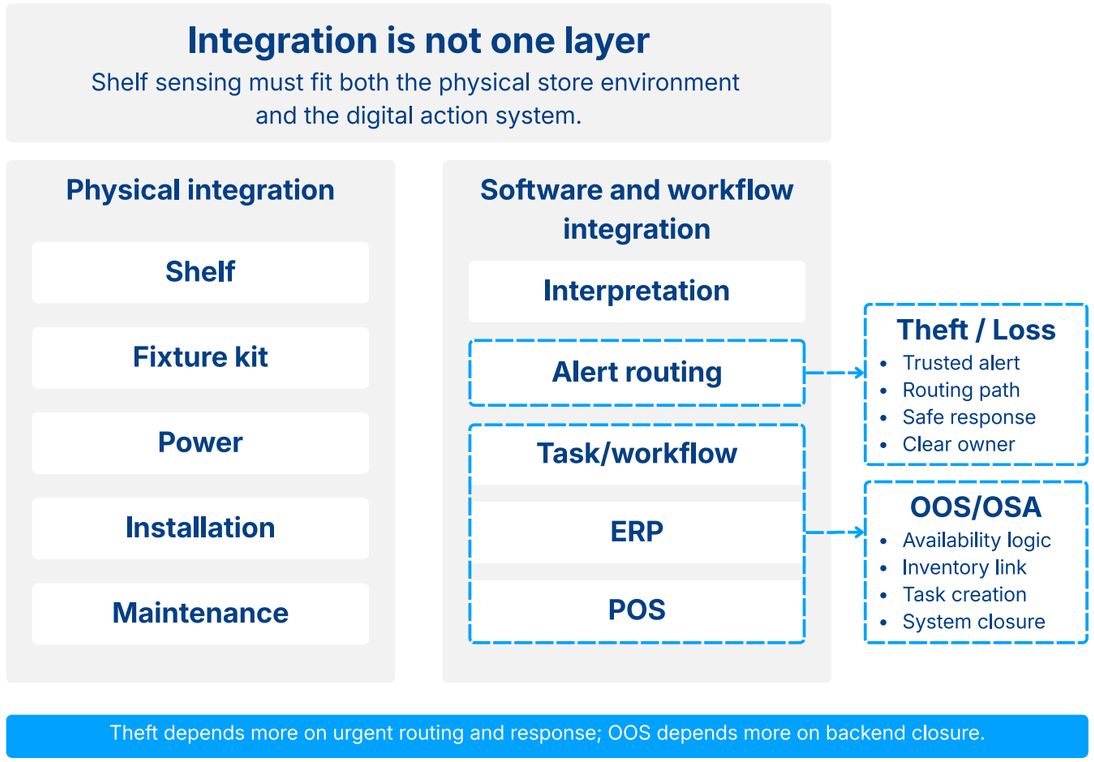

Figure 5.9. Two-layer integration view for theft and OSA.

5.3.1 What convergence reveals: the action chain is fragile at both ends

The clearest convergence is that both datasets point to fragility in the action chain. The store visits show this from the user side. Reliable signals, fast enough to act, and low extra workload define the threshold for adoption. The expert interviews show the same fragility from the system side. Trust can break through planogram drift, poor routing, master data errors, low alert quality, or unclear response ownership.

Both datasets describe the same chain from different sides. Store respondents describe the result they need: a signal they can trust. Experts describe why that trust breaks, for example through recalibration problems, poor routing, or alert fatigue. What stores experience as trust therefore depends on conditions they do not directly control.

A similar pattern appears around workload. Stores treat no extra staff work as a user requirement. Experts described the same issue as a wider execution problem. If the signal arrives through a separate app, requires manual checking, or depends on staff solving unclear exceptions, the burden shifts back into the store. What looks simple at store level can therefore depend on choices made elsewhere in the chain.

Speed also converges, but in a more pathway-specific way. Stores want the signal early enough to matter. Experts show that this means different things in different chains. For theft, it means credible routing into a live response. For OOS, it means timely correction through replenishment logic and routines. The store-level demand is the same in form, but the ecosystem burden behind it differs.

Figure 5.10 shows how the store-level and ecosystem-level findings connect. The figure is relevant because it makes clear that store requirements, such as reliability, speed, and low workload, depend on wider ecosystem conditions around data, ownership, routing, integration, and support.

5.3.2 What perspective gaps reveal: the same proposition is evaluated through different criteria

The strongest tension in the combined material was not disagreement about whether shelf sensing matters. Most actors treated the problem as real. The difference was how the proposition was judged.

Store teams evaluate the system against immediate workflow criteria. They ask whether it works reliably, whether they can act on it quickly, and whether it protects time and attention. Organisational gatekeepers use different tests. They ask whether the case can be proven financially, whether it fits the architecture, whether it disrupts customer experience, and whether it can be supported at chain-level. These are not small differences in emphasis. They are different evaluation logics.

That gap matters because a solution can pass the store-level threshold and still fail in rollout. The actor who feels the pain is often not the actor who controls the next gate. This was visible in the interviews around the ROI gate, the enterprise-system proof proof requirement, and the customer experience veto. It was also evident in the store data, where respondents often cared little about ownership, even though ownership later became central to sustained use.

The meaning of actionability also shifts across levels. For the store, actionability means that someone can do something useful now. For headquarters, it may mean that the signal can be translated into dashboards, replenishment logic, or enterprise proof. For loss-prevention actors, this may mean that the response model is legitimate and safe. The same proposition is therefore judged through different time horizons and risk concerns within the same retailer.

Table 5.3 compares the store-level and ecosystem-level evidence. The table shows that store requirements are often simple at the point of use, but depend on wider conditions around data, routing, ownership, integration, and support.

5.3.3 What is hidden: simple store demands depend on wider conditions

The synthesis also shows that simple store demands often depend on wider conditions. Store

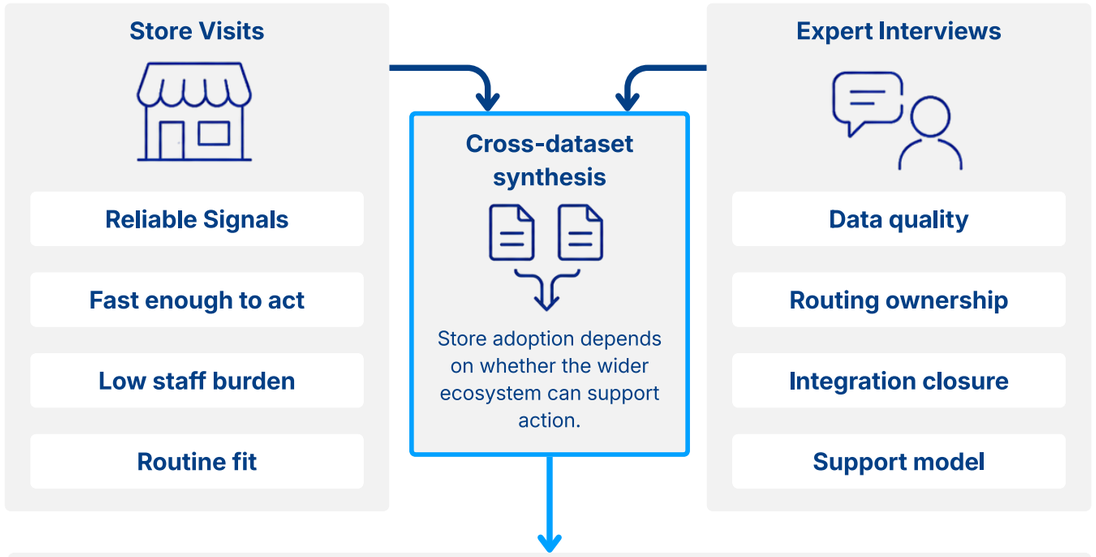

Figure 5.10. Cross-dataset synthesis chain.

This synthesis prepares the refined diagnostic framework in Figure 5.11.

teams usually name the outcome they need, not the full chain required to create that outcome.

Reliability is the clearest example. At store-level, it appears as a straightforward requirement: the signal must be trustworthy. The expert interviews show that this trust depends on a wider set of hidden conditions, such as stable configuration, accurate master data, credible routing, maintenance quality, and service response. The store does not usually name these mechanisms. It experiences their outcome.

The same applies to workload protection. Stores say the technology should not create extra work. Experts show what sits behind that demand: routing, ownership, remote updates, and an installation and maintenance model that fits existing infrastructure. In other words, low workload is not only a local preference. It depends on how the wider system is organised.

Works with existing technology showed the same pattern. Store respondents used this phrase to mean that the system should fit what is already there. Expert respondents showed that this can involve physical embedment, software openness, partner choice, and support after rollout.

5.3.4 What diverges after the signal: theft and OOS place pressure on different links

The synthesis confirms that theft and OOS both belong under actionable shelf insight, but they should not be treated as the same action chain. The store data hinted at this difference. The expert interviews made it clearer.

At store-level, theft was the dominant use case by volume. It is urgent, visible, and emotionally direct. OOS was also important, but it appeared less acute. At ecosystem-level, the balance shifts. Theft still matters, but OOS often carries the heavier integration burden because the value case depends on closing the replenishment and correction loop through other systems.

The sharpest divergence appears after the signal is generated. In the theft pathway, the weak point often sits at routing and response. The signal may be valid, but the response model can still fail due to

Perspective	*What counts as value*	*Main risk concern*	*What counts as proof*	*Possible rejection point*
Store users	Reliable signal, quick action, low extra workload	The system creates false alerts, takes time, or disrupts routines	It helps staff act in daily work	Staff do not trust it or stop using it
Store operations	Fits existing routines and staffing capacity	The task is not owned or cannot be handled during store work	It improves store execution without adding burden	The system creates more work than it removes
IT and architecture	Fits existing systems, data flows, and architecture	Integration effort is too high or technically hard to maintain	It can connect to current systems and be supported	It does not fit the architecture or requires too much IT capacity
Headquarters or business owner	Financial case, chain-level usefulness, repeatability	Pilot value does not translate into rollout value	It shows a clear business case across stores	The return on investment remains unclear
Loss prevention (LP)	Safe and legitimate response model	Alerts create unsafe or inappropriate staff responses	The signal supports credible escalation or intervention	The response model creates customer experience or safety concerns

Table 5.3: store users and organisational gatekeepers compared on what counts as value, risk, proof, and rejection.

low credibility, unsafe intervention assumptions, or customer experience concerns. In the OOS pathway, the weak point often sits in data, integration, and correction logic. The signal may detect the shelf issue, but the value is limited if the backend loop does not close.

These pathways also show different risk patterns. Theft appears to carry more response and adoption risk, because the signal must be trusted, routed, and acted on in a sensitive situation. OOS appears to carry more integration and backend-closure risk, because the signal must connect to data, replenishment, and correction routines. Both pathways matter, but they do not depend on the same links.

5.3.5 What the governance gap reveals: the structural problem neither side fully owns

The clearest combined finding is that no single actor clearly owns the full path from shelf signal to store action. In the store visits, clear ownership ranked last. On its own, this could suggest that ownership is not very important to store teams. The expert interviews showed a different pattern. Ownership became more important when respondents discussed rollout, maintenance, data quality, support, and decision making.

Together, the findings show that responsibility is spread across several actors. Store teams may need the signal, but information technology, operations, LP, procurement, data owners, finance, or external partners may control whether that signal can be used. Configuration quality may sit with one actor. Response logic may sit with another. Budget approval may sit somewhere else.

This creates a risk during rollout and maintenance. A pilot can appear useful in one store, while wider adoption still becomes difficult because no actor owns the full chain over time. For Chapter 5, the main result is therefore: product performance is not enough. Shelf sensing also needs clear responsibility for data, routing, response, support, and follow-up.

5.4 Framework refinement

The empirical findings did not replace the framework from Chapter 3. They made it more specific. The framework still treats actionable shelf insight as a problem of connected activities, actors, links, and alignment conditions. The results mainly sharpened three parts: the split between OOS and theft, the points where the action chain can break, and the role of ownership, integration, and internal readiness after a pilot

The framework should still be read from top to bottom and from left to right. It starts with actionable shelf insight as the focal value

proposition. It then shows the activity chain needed to move from shelf event to store action. The links in that chain are assessed through technical, data, operational, incentive, and governance alignment. The risk layer shows where the chain may break during pilot, rollout, store execution, sustained use, or scale. MOOS's internal readiness is kept separate because even a possible route may be difficult to deliver if MOOS cannot support it. Positioning is therefore not built into the framework. It follows after the diagnosis.

Figure 5.11 on the next page shows how the original framework from Chapter 3 was sharpened through the results. It keeps the same diagnostic logic, but makes three empirical refinements visible: the split between OOS and theft, the fragile links after detection, and the role of ownership, integration, and internal readiness in later positioning.

5.4.1 What became more specific through the findings

The main refinement concerns the activity chain. The original framework already described value realisation as a chain of activities, actors, positions, and links. The findings made this chain more concrete. Actionable shelf insight depends on detection, signal trust, routing, response ownership, and follow-up. Several problems found in the data were located between these steps rather than inside one step. Examples include false alerts, planogram changes, unclear routing, weak task ownership, and missing backend correction.

A second refinement concerns the difference between theft and OOS. Both pathways start from poor shelf visibility, but they do not fail in the same way. Theft places more pressure on urgency, alert routing, and realistic response. OOS places more pressure on data quality, system connection, and backend closure. The refined framework therefore separates these pathways after the shared visibility problem.

A third refinement concerns ownership and route to scale. The store visits showed that clear ownership was not a high store-level priority. The expert interviews showed that ownership becomes more important during rollout and sustained use. This means that some important conditions are partly hidden at store level. They become visible only when the signal has to be maintained, integrated, supported, and acted on across several actors.

5.4.2 Why the framework remains diagnostic

The refined framework does not choose a final position for MOOS. Its purpose is to show where actionable shelf insight becomes fragile, which links need attention, and where the two main pathways differ. This distinction matters because positioning should follow from the diagnosis. If a preferred position were built into the framework itself, the analysis would risk confirming its own assumption.

Chapter 5 therefore ends with a refined diagnosis, not a final design choice. The results have narrowed the design space, but they have not closed it. Chapter 6 starts from this smaller problem space and uses it to frame the design brief, compare route directions, and develop the roadmap.

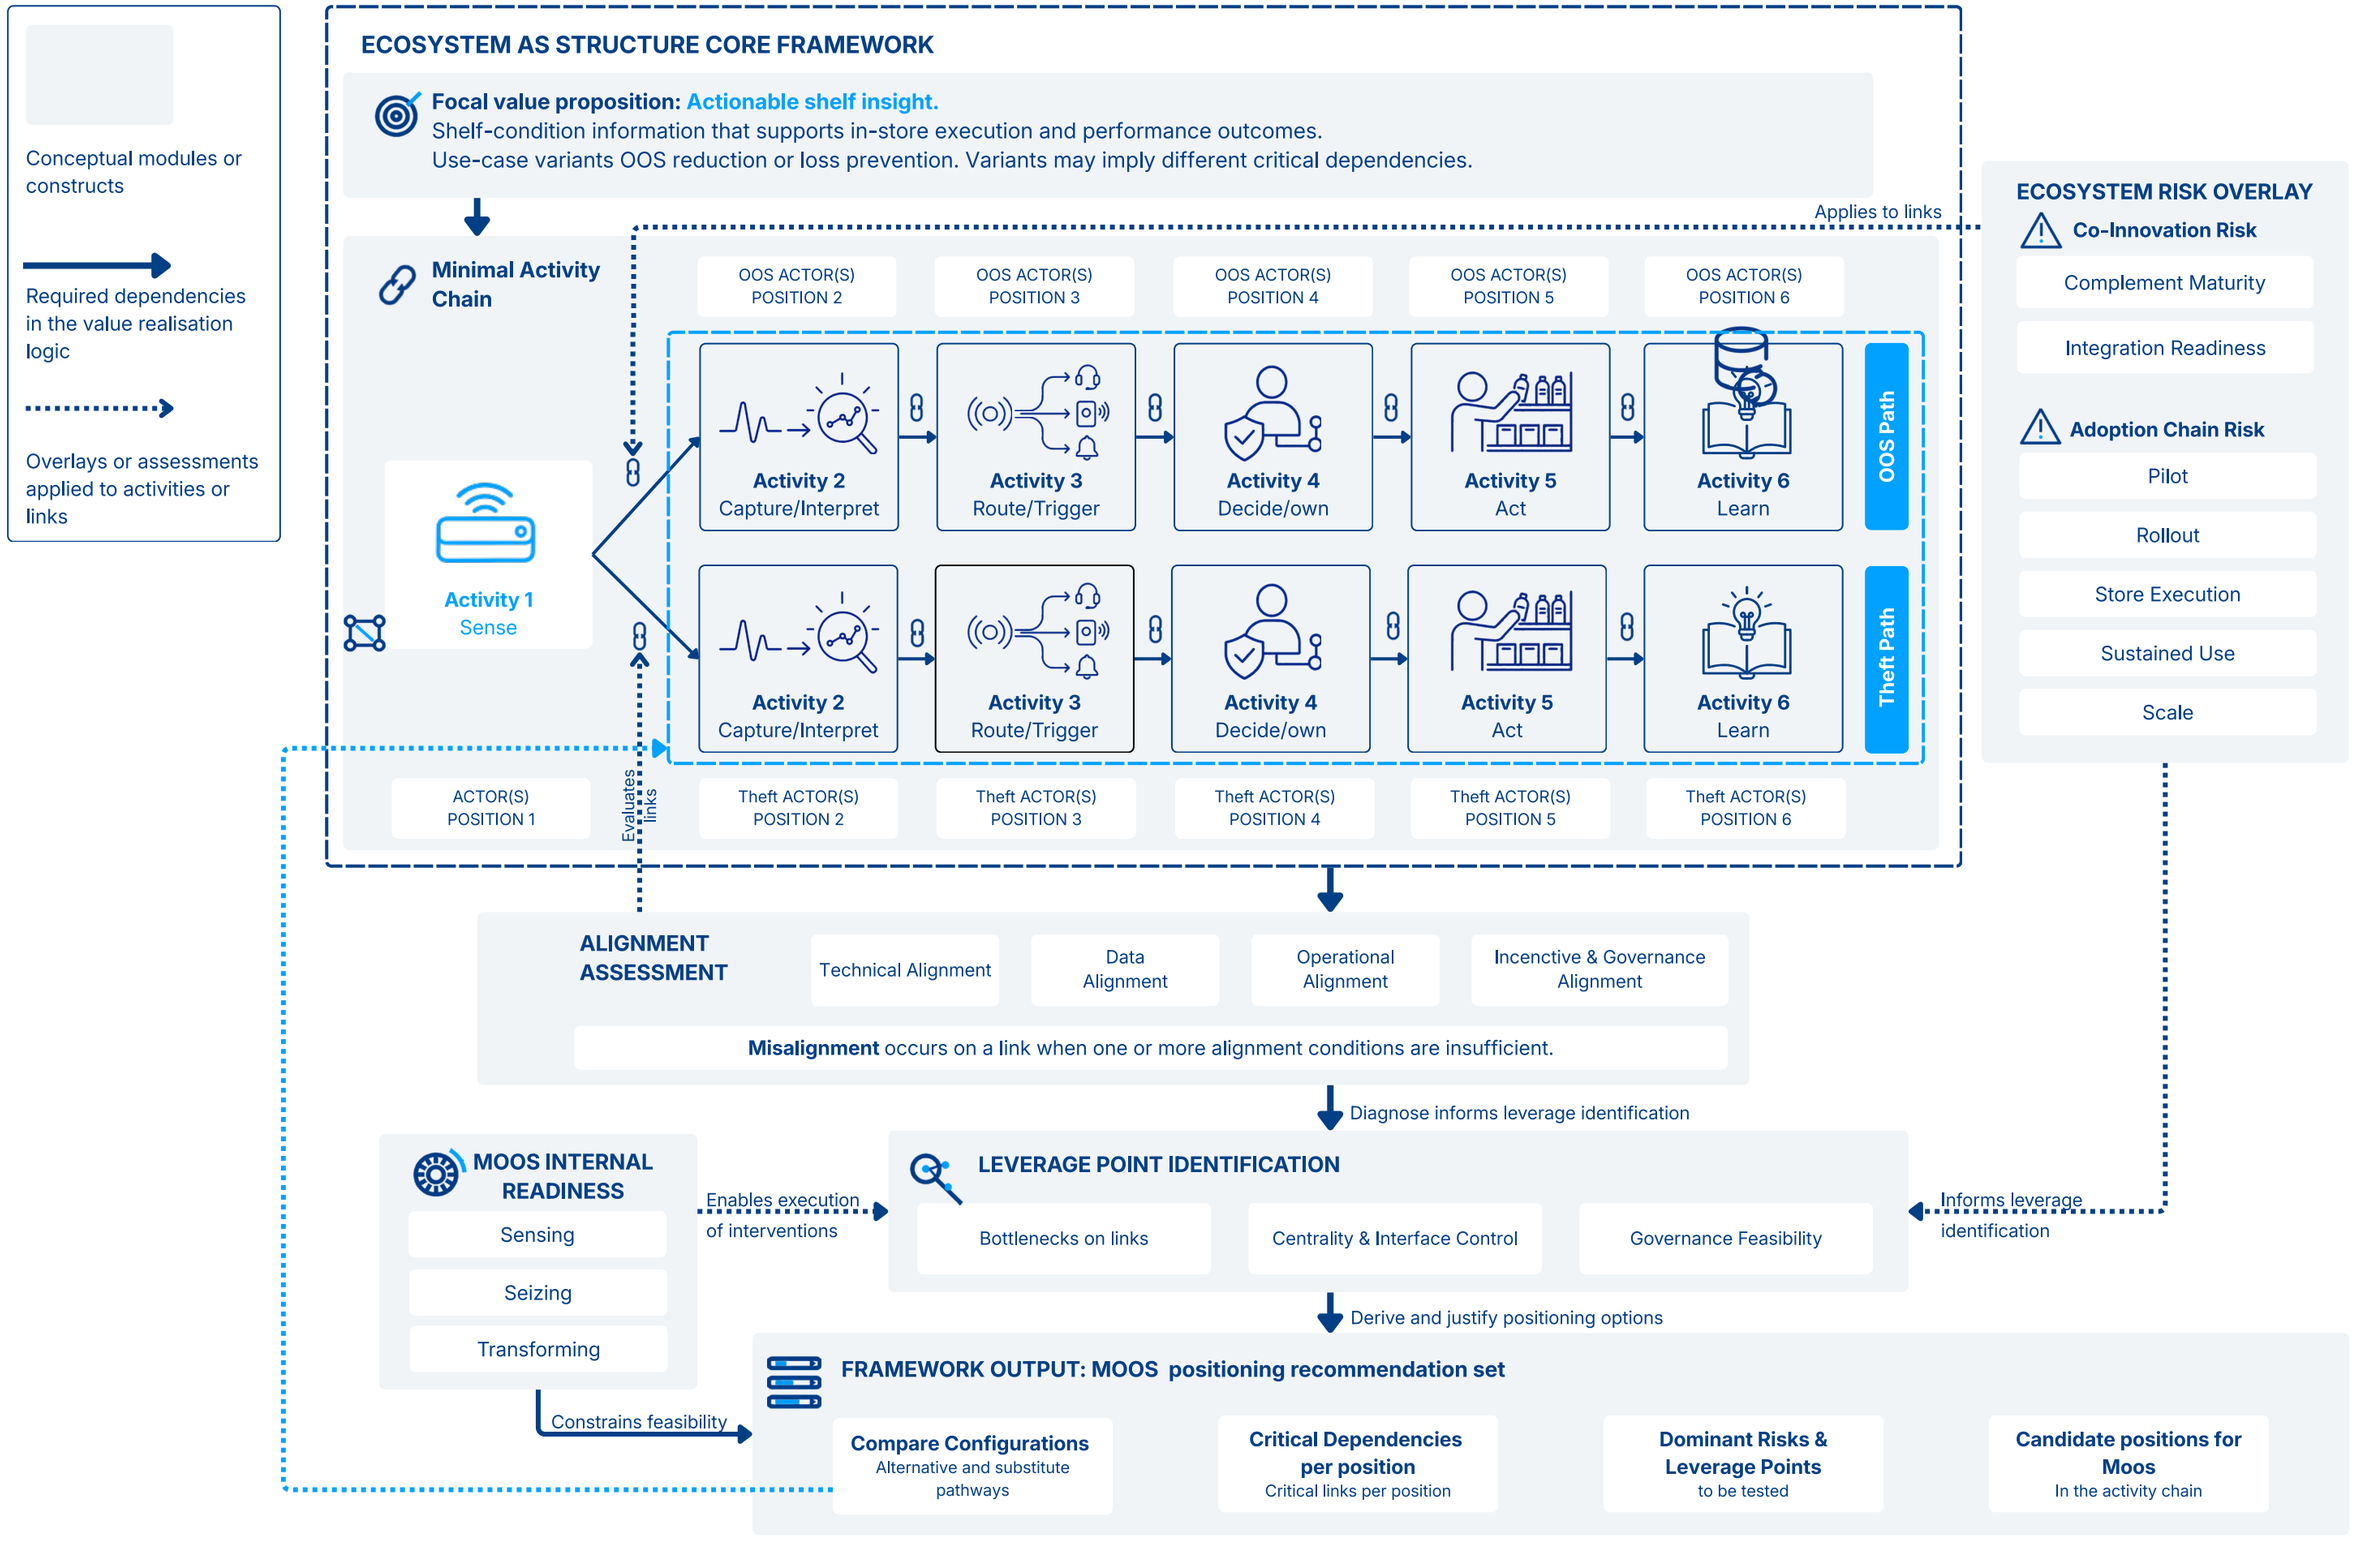

78 Figure 5.11. Refined diagnostic framework after the empirical findings.

Chapter 6

Design

6.1 From diagnosis to design
6.2 Design process and role of co-creation
6.3 Co-creation outputs as design input
6.4 Design criteria
6.5 Future vision
6.6 Ideation and development of positioning directions
6.7 Preferred direction
6.8 Public-level roadmap
6.9 Tactical level workstream
6.10 Conclusion

Chapter 5 narrows the design space, but it does not yet choose a positioning direction. The results show where actionable shelf insight becomes fragile, which ecosystem links carry most weight, and why MOOS cannot be positioned through product logic alone. Chapter 6 translates this diagnosis into design criteria, positioning elements, a preferred initial direction, and pre-validation roadmap logic.

The refined framework at the end of Chapter 5 forms the bridge into this chapter. It brings together the focal value proposition, the main pathway split between loss prevention and on-shelf availability, the most binding actor links, the dominant bottlenecks, and the internal readiness themes that matter for MOOS. From this point onward, the task shifts from diagnosis to design.

The design output in this chapter is deliberately pre-validation. It proposes a direction and roadmap logic, but it does not claim that the roadmap has already been proven. Chapter 7 then assesses whether this design output is clear, credible, useful, and relevant enough to be revised.

Figure 6.1 shows how the refined diagnostic framework from Chapter 5 is translated into the design phase. It connects the focal value proposition, pathway split, actor links, bottlenecks, internal readiness, design criteria, positioning direction, and roadmap logic. Customer references, specific segment choices, partner routes, and integration details are excluded from the public version.

6.1 From diagnosis to design

6.1.1 Design problem

The central design problem is to define how MOOS can move from a promising shelf-sensing technology toward a credible position in the retail technology ecosystem. This is first a positioning question, not a product question.

Chapter 5 showed that shelf sensing only becomes valuable when a shelf event reaches the right operational layer in a usable form. In this chapter, actionable shelf insight is therefore understood as shelf-related information that is reliable, specific, and timely enough to support store action. The design work stays tied to that action chain.

This means that the design phase cannot focus only on adding features. It must define where MOOS can credibly sit in the activity structure between shelf reality and store action. The positioning direction should explain what MOOS can own, what it should connect to, and where its role should remain bounded.

6.1.2 Scope of the design task

The applied framework narrows the scope. It shows that loss-prevention and on-shelf availability do not follow the same pathway. Loss-prevention use cases can start through lighter alert and response routing, while on-shelf availability depends more strongly on workflow links, planning routines, and system coordination. This difference means that MOOS cannot use one generic scaling route for all shelf-related use cases.

Element	*Input from earlier chapters*	*Meaning for the design phase*
Design challenge	Chapter 5 shows that actionable shelf insight depends on links between shelf events, actors, systems, routines, and store action.	Define how MOOS can move from shelf-sensing technology to a credible ecosystem position.
Core design focus	The focal value proposition remains actionable shelf insight.	Keep the design tied to store action, not sensing in isolation.
Main pathway split	Loss prevention and on-shelf availability follow different dependency patterns.	Avoid one generic route for all use cases.
Scope boundary	The findings do not support positioning MOOS as a full retailer platform or store operating system.	Define a bounded role between shelf event and store action.
Near-term logic	Early design focus requires clear value concentration, reliable trigger logic, and manageable response conditions.	Define early design logic through transferable conditions rather than named contexts or proof cases.
Longer-term logic	Heavier contexts require more system, workflow, and ownership alignment.	Treat system-dependent expansion as staged, not immediate.
Co-creation input	A confidential co-creation session supported the transition from diagnosis to design development.	Report co-creation only as process-level design support in the public version.
Intended output	The thesis moves from diagnosis to a designed positioning direction.	Produce a future vision, positioning logic, and public-safe roadmap abstraction.

Table 6.1. Design brief overview

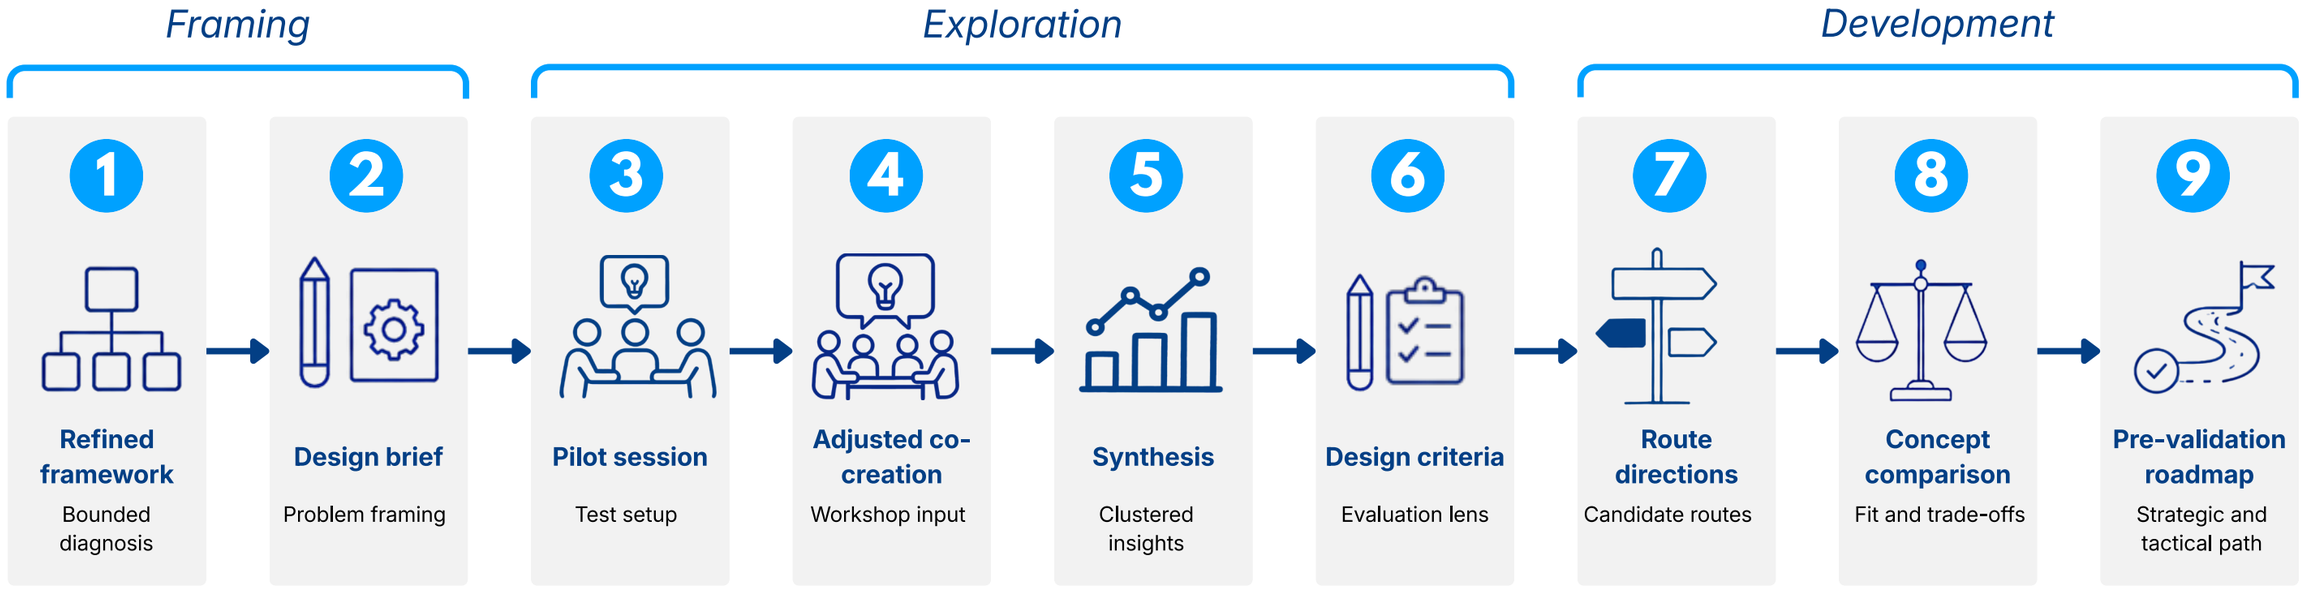

Figure 6.1. Design process from refined framework to roadmap.

The scope is also limited by MOOS's current role as a component provider. MOOS is not designed here as a full retailer platform, workflow suite, or store operating system. The design task is to define a position in which MOOS can own a useful part of the chain between shelf event and store action, while other actors remain responsible for wider retail systems and routines.

6.1.3 Intended design output

The chapter develops four connected outputs: a future vision, a set of positioning elements, a preferred initial direction, and two roadmap layers. The strategic roadmap explains how MOOS's position could develop over maturity stages. The tactical roadmap shows which workstreams need to be considered underneath that movement.

In this public version, the roadmap is presented at a higher level of abstraction. Detailed roadmap content, customer references, partner logic, integration routes, and internal strategic assumptions are excluded or generalised.

6.1.4 Design limits

Some limits are already visible before ideation begins. Not every retail context is equally feasible, and not every use case should be addressed at the same time. The findings in Chapter 5 suggest that contexts with clear shelf-event value, concentrated operational pain, and relatively direct response routes are easier to address early than contexts that require heavy integration with planning, replenishment, and wider operational systems.

These limits make the design brief more useful by preventing the chapter from treating every opportunity as equally realistic. The public version therefore reports the logic behind this selection, but not the detailed segment choices or internal prioritisation.

6.2 Design process and role of co-creation

After the design task had been bounded, the refined framework from Chapter 5 was translated into a design brief. That brief defined what the design phase had to address: the activities between shelf event and store action, the actor positions around those activities, the links that could break, the alignment issues that could block adoption, and the internal readiness conditions that MOOS had to consider.

Co-creation was then used as a design intervention. It did not function as a third empirical dataset or as validation. Its role was to help translate the diagnosis into route choices around role boundaries, partner responsibilities, sequencing, and what MOOS can credibly own. These are design questions, but they are also ecosystem questions.

The process moved from framework to design brief, from design brief to co-creation, and from co-creation output to synthesis, criteria, ideation, comparison, and roadmapping. The roadmap did not come directly out of the workshop. It was built through a sequence of translation, exploration, comparison, and synthesis.

6.3 Co-creation outputs as design input

Co-creation was used as part of the design process, not as a third empirical dataset and not as a validation step. Its role was to support the translation from empirical diagnosis to design development.

The content of the co-creation output is not included in this thesis. The session involved company-sensitive discussion about positioning, roadmap logic, strategic choices, possible routes, and internal assumptions. For that reason, detailed workshop output, participant input, visual materials, route discussions, and strategic conclusions are excluded from the public repository version.

In this chapter, co-creation is used to explain how the design phase was structured, not to disclose the content of the session. The design choices presented in the following sections are based on the combination of the theoretical framework, the empirical findings from Chapter 5, design reasoning, and a confidential co-creation process.

6.4 Design criteria

The design inputs needed one more step before concept development could begin. They had to be translated into criteria for comparing possible positioning directions. Without that step, the later selection would remain too intuitive.

The criteria are not general design requirements. They translate the diagnosis from Chapters 2 to 5 into design judgement. Each criterion connects back to a specific issue from the framework or empirical findings: the focal value proposition, store actionability, ecosystem links, repeatability, dependency load, internal readiness, and future option value.

Criterion	*Definition*	*Research logic behind the criterion*
Strategic fit	The direction should position MOOS as a shelf-sensing component provider with a clear ecosystem role.	The thesis question concerns positioning, not only product development.
Store actionability	The direction should help turn shelf events into reliable, timely, and usable store action.	Store value depends on signals that can actually be acted on.
Ecosystem fit	The direction should fit the activity chain, actor roles, and link dependencies around actionable shelf insight.	The framework shows that value depends on aligned actors and links.
Route to scale	The direction should be repeatable beyond isolated pilots or one-off deployments.	Chapter 5 shows that adoption can break across deployment, ownership, and routinisation.
Dependency load	The direction should avoid relying on too many uncontrolled external conditions at once.	Co-innovation and adoption-chain risks increase when value depends on external complements.
Internal capability fit	The direction should be feasible for MOOS's current and developing capabilities.	A credible position must be executable, not only attractive.
Future option value	The direction should keep later options open without making the near-term design logic too broad.	The roadmap should allow learning and adaptation over time.

Table 6.2. Public level design criteria

The design criteria in Table 6.2 are therefore presented as public-level criteria. They do not disclose the substantive output of the confidential co-creation process.

Table 6.2 shows how each criterion connects to the research logic. Strategic fit keeps the design tied to the thesis question. Store actionability keeps the design connected to the empirical finding that shelf information only becomes valuable when it can support store action. Ecosystem fit links the design to the framework's emphasis on actors, activities, positions, and links. Repeatability and dependency load address the risk that a technically promising proposition may fail when deployment, support, or external dependencies become too complex. Internal capability fit and future option value ensure that the direction remains feasible while keeping later learning paths open.

6.5 Future vision

Based on the empirical diagnosis, the design criteria in Table 6.2, and further design reasoning, the future vision is that MOOS develops toward a bounded shelf intelligence and trigger layer between shelf reality and store action.

The word bounded is important. It prevents the design direction from expanding into a full retailer operating system, workflow suite, or broad retail platform role. The term shelf intelligence indicates that the value is not only in raw sensing, but in interpreted shelf-event information. The term trigger layer indicates that MOOS's value lies in making shelf reality actionable, not only visible.

This future role therefore sits between two extremes. A narrow sensor-only role would keep MOOS too close to detection alone. A broad platform role would exceed what the findings support. The preferred future role sits between these positions. MOOS should focus on the part of the chain where shelf events are detected, interpreted, and prepared for follow-up action, while wider retail systems and operational ownership remain outside its core role.

Zone	*Public-safe content*
MOOS should own	Shelf-event sensing, event interpretation, trigger logic, signal quality
MOOS may connect to	Relevant actor or system categories where this improves actionability
MOOS should not own	Full retailer platform, full workflow suite, full enterprise architecture, all store operations

Table 6.3. Public-level future role boundary

6.6 Ideation and development of positioning directions

The ideation step opened the positioning space before a preferred direction was selected. The goal was not to create separate product concepts, but to identify possible positioning elements that could later be combined into a coherent direction.

The ideation was based on the refined framework from Chapter 5, the design criteria from Table 6.2, and further design reasoning. This helped keep the design development grounded in the empirical diagnosis rather than in speculative future scenarios.

Table 6.4 summarises the public-level ideation methods used in this step. The table reports the method logic, not the confidential content generated during the design process.

6.7 Preferred direction

The comparison of positioning elements led to one preferred direction. The public version does not include detailed numerical scores or internal prioritisation. Instead, it reports the qualitative reasoning behind the preferred direction.

The preferred direction can be described as follows:

MOOS should position itself as a bounded shelf intelligence and trigger layer between shelf reality and store action.

This direction combines several design elements. It keeps MOOS close to shelf-event sensing and interpretation, but it avoids positioning MOOS as the owner of the full retail operating system. It also recognises that shelf insight only becomes valuable when it can be interpreted, routed, and used in store operations.

Table 6.6 explains the public-level composition of the preferred direction.

This preferred direction fits the findings better than two more extreme alternatives. A sensor-only role would make the position too narrow and would not fully address the actionability problem found in Chapter 5. A broad platform role would make the position too wide and would create dependency and capability demands that go beyond the evidence. The preferred direction therefore defines a middle position: MOOS should focus on making shelf events reliable, interpretable, and usable, while leaving broader retail systems and operational ownership to other actors.

Method	*Purpose in the design phase*	*Public-safe output*
Design brief translation	Translate findings into design requirements.	Criteria and boundaries.
SCAMPER-style exploration	Open possible role and value directions.	Broad positioning elements.
Morphological exploration	Combine positioning variables and design boundaries.	Public-safe option space.
Thematic clustering	Group overlapping ideas into strategic themes.	Positioning themes.
Qualitative comparison	Compare themes against the design criteria.	Preferred direction logic.

 Table 6.4. Structured ideation methods

Positioning element	*Public-safe description*
Trusted shelf-event sensing	Reliable detection of relevant shelf changes.
Actionable trigger logic	Turning shelf events into signals that can prompt store action.
Bounded ecosystem role	Owning a focused part of the chain rather than the full store system.
Selective ecosystem connection	Connecting only where this improves actionability or repeatability.
Repeatable deployment logic	Reducing one-off implementation and support burden.
Use-case distinction	Recognising that different shelf-related problems follow different action chains.
Evidence-based expansion	Moving into new contexts only when comparable conditions are present.
Configurability	Making the offer easier to adapt without creating excessive custom work.

Table 6.5 Public-level positioning elements

Element of preferred direction	*Meaning*
Bounded	MOOS should not become the full retailer operating system or workflow owner.
Shelf intelligence	The value lies in interpreted shelf events, not only raw sensing.
Trigger layer	MOOS should help translate shelf reality into signals that can prompt follow-up action.
Between shelf reality and store action	The position is defined by the activity chain between detection, interpretation, routing, and response.

Table 6.6 Composition of the preferred positioning directio

6.8 Public-level roadmap

The preferred direction from Section 6.7 needs a maturity structure before it can guide strategic decision-making. MOOS cannot move directly from its current position to the future vision. The design logic needs to show how evidence, repeatability, readiness, and dependency load develop over time.

The roadmap in this public version is therefore not shown as a detailed implementation route. It is presented as abstract maturity logic. Specific customer references, proof cases, segment priorities, timing assumptions, partner routes, integration choices, and internal strategic decisions are excluded.

Table 6.7 presents the public-level roadmap logic. The table should be read from top to bottom as a maturity progression, not as a fixed calendar plan.

Table 6.7 changes the meaning of roadmap progression. The public version does not show a fixed sequence of markets, customers, years, partners, or integrations. Instead, it shows a movement from evidence to repeatability, from repeatability to selective deepening, and from selective deepening to justified ecosystem embedding.

This staged logic prevents the roadmap from becoming a broad expansion plan. It also prevents the future vision from being treated as an immediate operating model. The main function of the roadmap logic is to support decision-making about what should be developed now, what should be delayed, what requires further evidence, and what should remain outside the preferred scope.

6.9 Tactical level workstream

The tactical roadmap translates the preferred direction into workstreams. Its purpose is to show

which types of organisational and product development work need to support the strategic movement. In the internal version, this layer contains detailed assumptions about markets, use cases, integrations, deployment, external actors, proof points, and decision moments. These details are excluded from the public version because they are company-sensitive.

At a public level, the tactical logic can still be understood through generic workstreams. These workstreams help connect the preferred positioning direction to capability-building requirements without disclosing the detailed route.

Table 6.8 shows that the preferred position requires more than a market statement. It also requires capability development. Reliable trigger logic, repeatable deployment, manageable system links, and clear support roles are all part of making the position credible.

The tactical logic also makes the roadmap more useful. Instead of asking only which market or use case should come next, it asks which workstreams must mature before the position can scale responsibly. This connects the design output back to the internal readiness layer of the framework.

The preferred direction is that MOOS should develop as a bounded shelf intelligence and trigger layer. This position avoids two extremes. It does not reduce MOOS to a raw sensing component, but it also does not position MOOS as the owner of the full retailer operating system. Instead, it focuses MOOS's role on reliable shelf-event sensing, interpretation, trigger logic, and selective connection to the actors or systems needed for action.

The roadmap logic was developed as a pre-validation design output. In this public version, it is presented only as abstract maturity logic and generic tactical workstreams. Detailed customer references, proof cases, segment priorities, partner routes, integration choices, timing assumptions, and internal strategic details are excluded. Chapter 7 explains how the design output was validated and revised at a public level.

6.10 Conclusion

This chapter translated the empirical diagnosis into a design direction for MOOS. The design phase did not treat positioning as a branding exercise or as a product feature decision. Instead, positioning was developed as an ecosystem question: what role can MOOS credibly hold in the chain between shelf reality and store action?

Maturity stage	***Public-safe label***	***Strategic meaning***
Starting point	Current evidence base	Start from what is already credible and observable, without disclosing the proof case.
Stage 1	Prove repeatability	Strengthen reliability, deployment routines, support logic, and clear action routing.
Stage 2	Deepen selectively	Add interpretation depth and selected ecosystem links only where they improve actionability.
Stage 3	Embed where justified	Move toward broader ecosystem embedding only where evidence, readiness, and dependency load justify it.
Future ambition	Bounded shelf intelligence and trigger layer	Support accountable store action without becoming the full store operating system.

Table 6.7. Public-level roadmap maturity logic

Workstream	***Public-safe focus***
Market and customer logic	Assess fit through evidence strength, operational pain, and transferability.
Use-case development	Recognise that different shelf-related problems require different action chains.
Product development	Improve reliability, event interpretation, configuration, and actionability.
Ecosystem links	Connect selectively where links improve store action or repeatability.
Deployment and scale	Reduce one-off implementation through clearer routines and support roles.
Organisation and readiness	Define the internal capabilities needed to deliver and repeat the position.
Proof and decision logic	Move forward only when evidence, readiness, and dependency load are sufficient.

Table 6.8 Public-level tactical workstreams

Chapter 7

Validation & roadmap revision

7. Validation and roadmap revision
7.1 Object of validation
7.2 Participants and procedure
7.3 Validation criteria and analysis method
7.4 Validation findings
7.5 Roadmap revisions after validation
7.6 Final revised roadmap logic
7.7 Limitations of validation

Chapter 6 developed a pre-validation design output for MOOS. This output translated the empirical diagnosis, the preferred positioning direction, and the design reasoning into public-level roadmap logic. It positioned MOOS as a bounded shelf intelligence and trigger layer within the retail technology ecosystem.

The validation tested whether this pre-validation design output was clear, credible, useful, and practically relevant. The validation did not test market demand, technical feasibility, commercial performance, retailer acceptance, external actor willingness, or long-term implementation success. Its scope was narrower. The validation assessed whether the design output was strong enough as a strategic decision-support tool and whether it needed revision before being presented as the final design outcome of the thesis.

The public version reports the validation at a higher level of abstraction. Detailed validation feedback, customer references, proof cases, roadmap details, specific partner routes, integration choices, and internal strategic considerations are excluded or generalised. This keeps the validation process academically traceable without disclosing company-sensitive content.

The main outcome was that the broad positioning direction was not rejected. The idea of MOOS as a bounded shelf intelligence and trigger layer remained relevant. However, the roadmap logic needed to change. The validation showed that the pre-validation roadmap could still be read as too time-based, too broad, and too close to an expansion path. The revised version therefore shifts toward maturity stages, repeatability, readiness, dependency load, and decision criteria.

7.1 Object of validation

The validation focused on four connected parts of the design output.

First, the preferred positioning direction was tested. This concerned the idea that MOOS should not become the full operating system of the retailer. Instead, MOOS should focus on sensing shelf events, interpreting them, and turning them into signals that can support store action.

Second, the public-level strategic roadmap logic was tested. This roadmap logic showed the broad movement from current evidence toward the future ambition. The validation tested whether this movement was understandable, credible, and useful as a decision-support structure.

Third, the tactical workstream logic was tested. This layer explained what the strategic direction meant for use-case development, product development, ecosystem links, deployment, organisational readiness, and proof logic. The validation tested whether this level of abstraction helped reason about what to do, what to avoid, and what to prepare for.

Fourth, the use logic of the roadmap was tested. The validation did not only ask whether participants agreed with the direction. It also tested whether the roadmap logic could help assess new opportunities, requests, or dependency-heavy choices.

Together, these elements formed the object of validation. The session did not repeat the full design argument from Chapter 6. It tested whether the pre-validation design output was clear enough, specific enough, and practical enough to be revised into the final public-level design output.

7.2 Participants and procedure

The validation was conducted with three internal MOOS stakeholders. The participants are coded as V1, V2, and V3. Their roles covered different internal perspectives relevant to strategy, product development, and technical development. This gave the validation input from different parts of MOOS that would need to use or act on the roadmap logic.

This internal group was suitable because the roadmap first needed to work inside MOOS. The roadmap is meant to support choices about focus, sequence, product direction, dependency load, and internal readiness. These choices require knowledge of the company's current position, technical capacity, delivery limits, and strategic priorities.

The internal focus also creates a boundary for the findings. The session does not represent the full retail ecosystem. Retailers, implementation actors, technology providers, and external advisors may assess the roadmap differently. For this reason, the validation can support the revision of the roadmap for internal strategic use, but it cannot be used as external market evidence.

The session followed a guided review structure. First, the purpose and scope of the validation were explained. The participants then reviewed the pre-validation design output. The discussion focused on clarity, positioning, maturity logic, practical use, dependency load, and decision support. After the session, structured notes were organised in a validation matrix. This matrix was then used to translate the main feedback points into roadmap revisions.

7.3 Validation criteria and analysis method

The validation criteria used during the session are summarised in Table 7.1. These criteria were used to assess whether the pre-validation design output was understandable, useful, and specific enough to support strategic decision-making.

Table 7.1 shows that validation was not aimed at proving implementation success. The criteria focus on the quality of the roadmap as a decision-support tool. They assess whether the roadmap is understandable, bounded, useful, and explicit about assumptions and dependency conditions.

Feedback was classified into four types. Confirmation was used when a point supported the existing roadmap direction. Clarification needed was used when an idea was relevant, but the wording or visual explanation was not clear enough. Revision needed was used when the roadmap structure or content had to change. Limitation was used when a point was outside the scope of the validation or could not be solved by the roadmap itself.

Because the validation was documented through structured notes rather than a full transcript, this chapter does not use verbatim participant quotes. The validation matrix is the evidence base for the findings, but the public version only reports the non-confidential structure of the findings and revisions. The findings in Section 7.4 should therefore be read as interpreted validation findings. Their purpose is to show how structured feedback changed the design output.

The analysis followed three steps. First, feedback points were listed in the validation matrix. Second, related feedback points were grouped into themes. Third, these themes were translated into roadmap revisions. The classification and interpretation of the feedback were done by the researcher and checked against the structured notes before being translated into design responses. This keeps the link between validation feedback and roadmap revision traceable.

7.4 Validation findings

The findings below are based on the validation session. They are reported by theme, because several feedback points concerned the same underlying issue. The public version does not disclose customer references, proof cases, segment priorities, partner routes, integration routes, or confidential strategic assumptions.

7.4.1 The roadmap direction was clear, but the visual logic needed simplification

The first finding was that the general roadmap direction was understandable, but the visual form was too dense. Participants could follow the broad direction, but the pre-validation version contained too much text and too many parallel elements. This made it harder to use in a strategic discussion.

This feedback means that the issue was not only content, but also use. A roadmap that contains too many words and parallel elements can be difficult

Validation criteria	***Core question***
Understandability	Is the roadmap clear enough to understand without long explanation?
Positioning clarity	Is MOOS's proposed role clear and bounded?
Practical usefulness	Does the roadmap help MOOS make choices about focus and sequence?
Ecosystem credibility	Are the external dependency conditions believable and specific enough at a public level?
External value versus internal readiness distinction	Is it clear what creates external value and what improves MOOS's own readiness?
System-link abstraction	Is the system-link logic clear enough for different use cases and stages without disclosing specific integration routes?
Assumptions and risks	Are the main assumptions, risks, and decision points visible enough?

Table 7.1. Validation criteria used to assess the pre-validation design output

to use as a shared reference. It may show many relevant details, but still fail to guide choices quickly.

The roadmap revision was therefore visual and structural. The revised public-level version separates the strategic maturity logic from the tactical workstream logic. The strategic roadmap shows the main maturity path. The tactical roadmap contains the more detailed assessment rows.

7.4.2 The positioning direction remained relevant, but needed a clearer starting point

The validation supported the broad positioning direction. MOOS should still be positioned as a bounded shelf intelligence and trigger layer, not as a full store operating system. This confirmed that the core design direction from Chapter 6 remained usable.

However, the feedback also showed that the roadmap needed a clearer starting point. The pre-validation roadmap moved too quickly from current evidence toward future options. This made the roadmap look more like an expansion path than a staged movement from evidence to repeatability.

The interpretation was that the preferred positioning direction remained useful, but it needed to start from existing evidence. The revised roadmap therefore starts from an anonymised current evidence base. This does not prove that MOOS can expand everywhere. It shows that the next step should be to understand which conditions make the current evidence transferable and repeatable.

7.4.3 The roadmap needed to work as a decision tool, not only as a presentation of options

The validation showed that the roadmap had to help MOOS make choices. The pre-validation roadmap showed several possible directions, but did not always make the decision logic explicit enough. This made the roadmap broad, but not selective enough.

The interpretation was that the roadmap needed sharper decision logic. MOOS needs to assess new opportunities under limited time, capacity, and evidence. The roadmap therefore needs to help answer practical questions: does this opportunity fit now, should it wait, should it be developed with external support, should it be rejected, or is more learning needed first?

The roadmap revision was to add decision gates. These gates connect each stage to evidence, readiness, and repeatability. They prevent the roadmap from becoming a list of attractive options.

7.4.4 The roadmap logic was too time- and segment-led

The strongest structural finding was that the pre-validation roadmap could be read as too strongly based on time and segment sequence. It suggested more certainty about timing and market order than the evidence could support. This point changed the structure of the roadmap rather than only its wording.

The interpretation was that progress should not be defined mainly by years or retail segment labels. Progress depends on what MOOS has learned, what can be repeated, which dependency conditions are manageable, and how much complexity the next step adds. A segment label does not explain why an opportunity is suitable. The underlying conditions do.

This changed the roadmap logic. The revised roadmap no longer presents the stages as fixed time horizons. It uses maturity stages instead. These stages show what must be understood, repeated, tested, and prepared before moving forward.

7.4.5 Different shelf-related problems do not scale in the same way

The validation reinforced a distinction that was already visible in the empirical findings: different shelf-related problems do not scale in the same way. They can all be linked to shelf sensing, but they create different action needs and different dependency conditions.

Some problems are closer to direct trigger-based action when response routes are relatively clear. Other problems require more system alignment, ownership clarity, and backend closure before shelf insight becomes actionable. This distinction matters because a single roadmap path can hide important differences in complexity.

The roadmap revision was to keep use-case differences visible without disclosing specific route choices. The revised public-level roadmap therefore treats use-case development as a workstream and links progression to evidence, readiness, and dependency load.

7.4.6 External dependency roles needed clearer purpose

The validation showed that external dependency roles were too general in the pre-validation roadmap. The roadmap referred to external actors and system links, but it did not always explain what these dependencies were needed for.

The interpretation was that dependency purpose matters more than dependency labels. Some external links strengthen customer-facing actionability. Others improve MOOS's own delivery, support, monitoring, or repeatability. Mixing these purposes creates confusion.

The roadmap revision was to distinguish between external value and internal readiness. This makes it clearer whether an external link improves the retailer proposition, improves MOOS's ability to deliver, or creates additional dependency load.

7.4.7 System-link logic needed clearer abstraction

The validation showed that the earlier framing of system links was too simple. It helped as an initial contrast, but it did not explain what kind of system link was needed, why it was needed, or when it should be added.

The interpretation was that system links have several dimensions. One dimension is depth: how strongly the solution connects to another system. Another dimension is breadth: how many systems or components are involved. A third dimension is dependency load: how much the solution depends on other actors, systems, or routines to create value.

The roadmap revision was to replace simple integration language with a more abstract dependency logic. The public version does not disclose specific integration routes. It only shows that system links should be added selectively and only when the added complexity is justified by evidence and readiness.

7.4.8 The main risk was premature expansion

The main risk identified through the validation was premature expansion. The pre-validation roadmap could be read as encouraging MOOS to move too broadly before the current evidence pattern had been understood and tested.

The interpretation was that MOOS could be pulled into too many opportunities too early. This is risky because each new context, use case, external link, or system request may add custom work. It may also increase support burden and weaken the clarity of the company's role.

The roadmap revision was to slow down the logic of expansion. The revised public-level roadmap begins with current evidence, then asks MOOS to identify repeatability conditions, then tests adjacent possibilities, and only later moves toward deeper ecosystem embedding where justified.

7.5 Roadmap revisions after validation

The main roadmap revisions that followed from the validation are summarised in Table 7.2 on the next page. The table does not repeat the full validation matrix. It shows the main public-level changes that shaped the revised roadmap logic.

Table 7.2 shows that the largest change was not a new positioning direction, but a change in roadmap logic. The revised roadmap moves away from fixed time periods and broad expansion, and instead starts from evidence, repeatability, readiness, and decision gates.

7.6 Final revised roadmap logic

7.6.1 How the revised roadmap logic should be read

The revised roadmap logic consists of two connected layers. The strategic layer shows the maturity path. The tactical layer explains how each stage can be assessed in more detail.

The roadmaps should not be read as a fixed rollout plan. They do not state that MOOS must enter a specific segment in a specific year. They show how MOOS can move from current evidence toward future ambition through learning, repeatability, and readiness.

This reading responds directly to the validation. The pre-validation roadmap risked looking like a time-based expansion plan. The revised roadmap logic makes the sequence conditional. MOOS should move forward when the previous stage has produced enough evidence and when the next step does not create too much complexity too early.

Validation evidence	*Interpretation*	*Roadmap revision*
The roadmap direction was understandable, but the visual was too wordy.	The roadmap needed simpler communication and a clearer split between strategic logic and tactical detail.	The revised roadmap uses a simpler strategic figure and a separate tactical matrix.
The horizon logic was too strongly linked to time periods.	Progress should depend on evidence, learning, readiness, and repeatability instead of fixed calendar stages.	The roadmap shifts from time horizons to maturity stages: current evidence base, Now, Next, Later, and Future ambition.
The current evidence pattern needed to be made more explicit.	The roadmap needed to start from current evidence before discussing wider expansion.	An anonymised current evidence base becomes the starting point and anchor for the revised roadmap.
Segment sequence was too dominant.	Retail segments are less useful than the conditions that make a context repeatable.	Transferable conditions replace broad segment sequence as the main filter for early opportunities.
Different use-case pathways were treated too similarly in the roadmap.	Some use cases can fit a more direct trigger route, while others need more system links and operational ownership.	Use-case differences are kept visible in the tactical roadmap, while more system-dependent routes require stronger evidence and readiness before moving forward.
External dependency logic mixed customer-facing value with MOOS delivery needs.	External value and internal readiness require different choices.	The tactical roadmap splits customer-facing value from MOOS readiness.
External dependency and system-link categories were too general.	External dependencies and system links need to be linked to purpose, stage, and dependency load.	Dependency roles and system-link types are specified at an abstract level in the tactical roadmap.

Table 7.2. Main public-level roadmap revisions after validation

7.6.2 Revised strategic maturity logic

Figure 7.1 shows the revised strategic maturity logic after validation. It presents the roadmap as a maturity path with five stages: current evidence base, Now, Next, Later, and Future ambition. The visual should be read from left to right. Each stage represents a different level of evidence, learning, repeatability, and readiness. The figure should not be read as a fixed calendar plan or as a fixed segment expansion path. It omits customer names, proof cases, specific segments, dates, partner routes, integration routes, named systems, and internal priorities.

The strategic roadmap changes the meaning of progression. The pre-validation roadmap could be read as movement through time and segments. The revised version shows movement through maturity. This makes the roadmap more useful because it links each step to evidence and readiness.

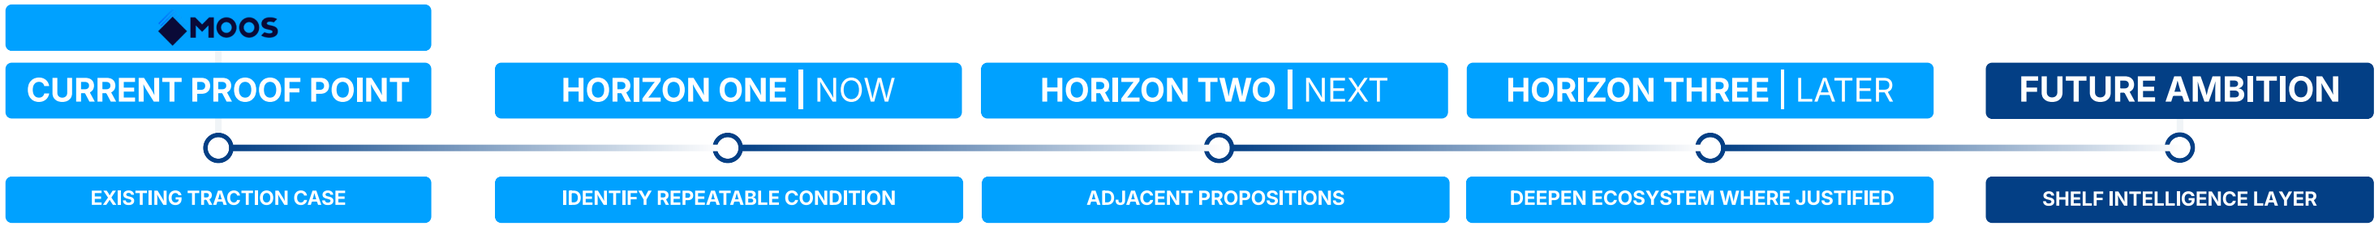

Figure 7.1. Revised strategic maturity logic after validation

Tactical workstream	*Roadmap revision*
Transferable conditions	The revised roadmap uses a simpler strategic figure and a separate tactical matrix.
Use-case and proposition focus	The roadmap shifts from time horizons to maturity stages: current evidence base, Now, Next, Later, and Future ambition.
Product and trigger logic	An anonymised current evidence base becomes the starting point and anchor for the revised roadmap.
External value	Transferable conditions replace broad segment sequence as the main filter for early opportunities.
Internal readiness	Use-case differences are kept visible in the tactical roadmap, while more system-dependent routes require stronger evidence and readiness before moving forward.
Deployment and scale	The tactical roadmap splits customer-facing value from MOOS readiness.

Table 7.3. Revised tactical workstream logic after validation

7.6.3 Revised tactical workstream logic

Table 7.4 shows the revised tactical workstream logic after validation. It is structured as a public-level matrix. The columns follow the same maturity stages as the strategic logic. The rows describe the practical questions that MOOS must answer in each stage. It presents the workstream logic without detailed sequencing, customer references, partner routes, integration assumptions, named systems, or segment choices.

The tactical workstream logic adds detail without changing the strategic direction. Its role is to test an opportunity against the maturity stage. This is why the matrix is important: it connects the validation feedback to practical decision-making without disclosing the confidential roadmap.

7.6.4 How MOOS can use the revised roadmap logic

The two roadmap layers have different roles. The strategic maturity logic locates the opportunity. It helps MOOS decide whether an opportunity belongs to the current evidence base, Now, Next, Later, or the future ambition. The tactical workstream logic then tests the opportunity in more detail.

This combined use is useful because opportunities can look similar at first, but differ in complexity. A new customer request may look close to current evidence, but may need a different response route. An external proposal may sound useful, but may mainly improve internal readiness rather than external value. A system-link request may improve the proposition, but may also add too much dependency too early.

For each new opportunity, use case, external proposal, or system-link request, MOOS can use the two roadmap layers to choose between five decision options. These options are summarised in Table 7.4.

Table 7.4 turns the roadmap into a decision-support tool. It helps MOOS keep future growth routes visible while avoiding premature expansion.

Decision option	***Meaning***
Do now	The opportunity fits the current stage and uses existing evidence or readiness.
Do later	The opportunity is relevant, but requires evidence, readiness, or external conditions that are not yet in place.
Do with external support	The opportunity fits the direction, but MOOS should not carry the full value delivery or delivery burden alone.
Do not do	The opportunity does not fit the bounded positioning or creates too much distraction.
Learn more first	The opportunity may fit, but the assumptions are not clear enough yet.

Table 7.4. Decision options for using the revised roadmap logic

7.7 Limitations of validation

The first limitation is empirical scope. The validation was conducted with three internal MOOS stakeholders. This was appropriate because the roadmap first needed to work as an internal strategic tool. However, it does not represent the full retail ecosystem. Retailers, implementation actors, technology providers, and external advisors may respond differently to the revised roadmap logic.

The second limitation is validation status. The session tested clarity, usefulness, credibility, and practical relevance. It did not test market demand, technical feasibility, commercial performance, retailer acceptance, external actor willingness, or long-term operational effectiveness. Because the validation was internal, the revised roadmap is validated for internal strategic use, not as externally confirmed market evidence.

The third limitation is implementation uncertainty. The revised roadmap depends on future pilots, technical choices, external conditions, and retailer priorities. These conditions may change. The roadmap should therefore be treated as a strategic design output that can guide decisions, not as a final answer for all future situations.

These limitations do not make the validation irrelevant. They define what the validation can and cannot claim. Within this scope, the revised roadmap logic can be used as the final design output of the thesis. The validation made the roadmap stronger by shifting it from a broad expansion path toward a maturity-based decision-support tool. The next chapter discusses what this output means for MOOS and for ecosystem-based positioning of component providers.

Chapter 8

Discussion

8.1 Theoretical contribution
8.2 Methodological and design contribution
8.3 Practical implications for MOOS
8.4 How the roadmap should be used
8.5 Broader practical relevance
8.6 Limitations
8.7 Future research
8.8 Closing reflection

8. Discussion

This chapter discusses what the findings and revised roadmap mean for theory, method, practice, and future research. The main argument is that shelf sensing becomes strategically relevant only when it is connected to action. Detection alone is not enough: a shelf event must become a trusted signal, reach the right actor or system, and lead to follow-up without creating too much extra work.

This shifts the positioning problem from product performance to ecosystem fit. The issue is not only whether the technology works, but whether the proposed role fits the systems, routines, external dependencies, and decision gates around retail execution. The revised roadmap matters because it turns this issue into a practical decision logic: each opportunity can be assessed against current evidence, repeatability, customer-facing value, internal readiness, and ecosystem fit. Chapter 9 gives the final answer to the research question.

8.1 Theoretical contribution

The theoretical contribution of this thesis is applied. It does not develop a new ecosystem theory. It operationalises ecosystem-as-structure for diagnosing the positioning problem of a component provider in store-level retail execution. Positioning is treated as an outcome of ecosystem diagnosis, not as a starting assumption.

Ecosystem-as-structure defines an ecosystem around a focal value proposition and examines the actors, activities, positions, and links needed to realise that proposition (Adner, 2017). In this thesis, the focal value proposition is actionable shelf insight. The framework shows that this value depends on the links between shelf events, data interpretation, operational ownership, external system and actor links, and store action.

This link-level view is the main conceptual contribution. A broad stakeholder map can show who is involved, but it does not show where value may break down. The framework helped identify where handovers, dependencies, and risks appear in the activity chain. For shelf sensing, the weak point is often not the sensor itself, but the movement from signal to action.

Co-innovation risk and adoption-chain risk were used to make this diagnosis more precise (Adner & Kapoor, 2010). Co-innovation risk helped show where MOOS depends on other actors, systems, or complements outside its direct control. Adoption-chain risk helped show where scale can break down across approval, implementation, store execution, routinised use, support, renewal, and wider adoption.

The internal readiness layer adds a design-oriented contribution. Ecosystem theory explains external dependencies, but MOOS's ability to deal with those dependencies also depends on its own capabilities. By adding internal readiness, the framework connects ecosystem position to deployment, support, documentation, system-link discipline, and evidence building. This makes the framework useful not only for analysis, but also for roadmap development.

The contribution remains bounded. The framework was developed through one case and one retail-technology context. It should not be treated as a validated model for all retail ecosystems. Its broader value lies in showing how ecosystem-as-structure can be used as a diagnostic and design-oriented approach for component-provider positioning.

8.2 Methodological and design contribution

The thesis also makes a methodological and design contribution. It shows how a store-level and ecosystem-level diagnosis can be translated into a strategic design output. The store visits identified practical adoption thresholds, while the expert interviews explained wider ecosystem conditions. The cross-dataset synthesis connected these levels before the design process translated the diagnosis into design criteria and roadmap logic.

This process matters because the roadmap was not created as a future vision alone. It was built from empirical findings about store-level adoption thresholds, expert insights about ecosystem dependencies, design reasoning, and validation feedback. The final roadmap therefore connects research findings to practical decision-making.

The validation further refined the design logic. It did not prove market success, but it did test whether the roadmap was clear, relevant, credible, and useful. Earlier roadmap thinking relied more on time periods and expansion logic. The revised version became more maturity-based. It starts from the current evidence base, then asks what can be repeated, what adjacent possibilities can be tested, and where selective ecosystem embedding becomes realistic. This shift made the roadmap more useful as a decision filter.

8.3 Practical implications for MOOS

The main practical implication is that positioning should be treated as an ecosystem decision. A promising context is not enough. MOOS should assess each opportunity through the conditions that make shelf insight actionable: customer-facing value, action route, dependency load, ecosystem fit, and internal readiness. This leads to five practical recommendations.

First, MOOS should start from current evidence and define the repeatable pattern behind it. A successful pilot or customer case does not automatically create a repeatable position. The company should document which conditions made the current traction work, such as use-case type, store context, hardware setup, trigger logic, support model, decision owner, and the route from signal to action. This prevents one evidence point from being treated as proof that the same route will work across all retail contexts.

Second, MOOS should use transferable conditions before entering new contexts. A new opportunity should not only be judged by market size or commercial interest. It should be compared with the conditions that already work. If the store context, use-case urgency, staff role, system-link need, or decision owner is very different, the opportunity should be treated as a learning case or not-yet option. This would help MOOS grow from evidence rather than from broad market ambition.

Third, MOOS should keep different use-case pathways separate in its scaling logic. Some shelf-related problems have clearer urgency and a more direct trigger-to-action path. Other problems require stronger links to operational systems, planning routines, data quality, ownership structures, and backend closure before they can become repeatable. They should therefore not be scaled in exactly the same way.

Fourth, MOOS should assess external actors and system links by purpose and dependency load. An external link may reduce deployment friction, provide infrastructure, improve actionability, support evidence quality, or open access to existing retailer systems. The question should not be whether more external connections are useful in general, but what a specific link is meant to solve. The timing and depth of each link should depend on purpose, system type, maturity stage, and dependency load.

Fifth, MOOS should build internal readiness before scaling the preferred position. The findings show that actionability depends on reliability, fast follow-up, and low additional workload. For MOOS, this means that deployment, configuration, alert filtering, documentation, support capacity, and evidence material should become more repeatable. Without these internal conditions, a promising opportunity can still become difficult to sustain.

These recommendations turn the revised roadmap into a practical decision filter. New opportunities should be tested against three questions: does the opportunity create customer-facing value, is MOOS ready to deliver and support it, and does the ecosystem around the opportunity allow actionability? If one of these conditions is weak, the opportunity should be treated as not yet ready rather than as a direct growth route.

8.4 How the roadmap should be used

The revised roadmap should be used as a decision-support tool, not as a fixed rollout plan, market forecast, technical proof, or proof of commercial success. Its value lies in helping MOOS make choices under uncertainty.

The roadmap supports three uses. First, it helps prioritise what should be proven now and what should wait. Second, it helps compare new opportunities through the same decision logic. Third, it supports internal alignment between sales, product development, technology, deployment, and leadership.

These uses keep the roadmap practical beyond its visual form. It should help MOOS discuss focus, timing, external dependencies, and readiness with the same logic, without treating the roadmap as a proven action plan for rollout.

8.5 Broader practical relevance

Although the thesis is centred on MOOS, the findings may also be useful for other early-stage retail technology providers with component-based offers. This transfer is limited. It is most relevant for firms that depend on retailer systems, store routines, data access, and investment decisions to create value.

For these firms, positioning should start from the activity chain required to deliver value. It is not enough to define a product category or list customer benefits. Firms should map who must act, which systems must connect, which external

actor links may be needed, and where adoption can stop.

This is especially relevant for technologies that enter physical stores. Stores combine labour pressure, legacy systems, local routines, and head-office decision-making. A roadmap for such technologies should therefore show dependencies, sequencing, and readiness conditions, not only future ambition.

8.6 Limitations

The first limitation concerns empirical scope and actor coverage. Part 1 consisted of 41 store visits and short conversations with a ranking exercise. This provided useful breadth and helped identify store-level adoption thresholds, but the visits were not designed as in-depth ethnographic research. Part 2 consisted of eight semi-structured expert interviews. The respondents were selected purposively to cover different ecosystem positions, which was suitable for analytical insight, but does not allow statistical generalisation. Some perspectives were also less visible, including store employees, procurement, legal and privacy roles, enterprise software providers, and some supply-chain roles.

The two datasets also differed in depth and purpose. Part 1 provided broader but shallower store-level input. Part 2 provided fewer but deeper explanations of ecosystem dependencies. They were therefore not treated as equal evidence sources. Instead, they were combined through analytical triangulation: Part 1 helped identify store-level adoption thresholds, while Part 2 helped explain the ecosystem mechanisms behind them.

The second limitation concerns design and validation status. The co-creation session supported the translation from findings to design criteria and roadmap logic, but it should not be treated as independent validation. The conducted validation helped assess clarity, relevance, credibility, and usefulness, and it informed the revised roadmap. However, it did not prove market success, technical performance, commercial viability, external actor willingness, or long-term operational effectiveness. Further external validation is still needed.

The third limitation concerns practical implementation. The revised roadmap is based on maturity, repeatability, and decision criteria rather than fixed year-based planning. This makes it more suitable for uncertainty, but it also leaves timing open. The roadmap has not been tested through long-term rollout or actual opportunity selection. It should therefore be treated as a strategic guide and decision filter, not as a detailed action plan.

The company-embedded nature of the project also affects interpretation. The research was conducted from late October to 1 June by an external graduation student working around MOOS. This position helped build sector understanding and gave access to company discussions, especially through contact with Marc and formal project moments. At the same time, company input could influence interpretation, because MOOS was already working on positioning, external links, and use-case focus. This risk was reduced by keeping the empirical findings, framework, and validation criteria leading, but it remains part of the research context.

A final limitation concerns confidentiality. Some detailed design and roadmap material cannot be fully reported because it contains company-sensitive information. This means that some of the design reasoning is presented at a higher level of abstraction than in the internal assessment version.

8.7 Future research

Future research should first test the revised roadmap with actors who were not involved in the co-creation or earlier validation process. Relevant participants include retail operations experts, retail technology experts, implementation specialists, potential external collaborators, and MOOS stakeholders who were not central in the design phase. This would show whether the design logic is understandable and credible beyond the immediate project context.

A second step should use pilot-based evidence. Useful indicators include signal reliability, false positives, response time, staff response, usefulness of diagnostic evidence, reduction in manual checks, operational response time, system-link effort, data accuracy, and fit with existing routines. This evidence would show whether the design assumptions hold in store practice.

A third step should test external feasibility. MOOS should examine whether relevant external actors recognise the proposed role of MOOS and see a feasible collaboration logic. This would test whether the ecosystem-compatible position is useful not only for MOOS, but also for the actors around MOOS.

Further research could apply the framework to other retail technology components. Possible cases include camera analytics, RFID infrastructures, smart carts, or autonomous checkout technologies. Comparing these cases would show whether the framework can transfer beyond shelf sensing and where it needs adjustment.

8.8 Closing reflection

This thesis started from a positioning problem. MOOS had promising technology, but its long-term role in the retail technology ecosystem was not self-evident. The research shows that this uncertainty cannot be solved by choosing a market label alone. It requires understanding how value is created through actors, interfaces, routines, and adoption gates.

The project also showed the difficulty of doing strategic design work as an individual graduation project in a small company context. MOOS provided valuable access to product knowledge, strategic discussions, and the practical reality of a start-up. At the same time, the project required enough distance to keep the research evidence and framework leading. For me, the main learning was that strategic design is not only about developing a strategic direction, but also about taking ownership of the process, asking for input at the right moments, and deciding what not to include.

These limits are important for how the thesis should be read. The revised roadmap clarifies a direction and decision logic, but the final answer remains bounded by the evidence gathered in this project.

Conclusion

9.1 Answer to the main research question
9.2 Answer to SQ1: structure and actionability
9.3 Answer to SQ2: dependencies and alignment
9.4 Answer to SQ3: risks and feasibility
9.5 Answer to SQ4: points of influence and readiness
9.6 Final remarks

9.1 Answer to the main research question

This thesis set out to answer the following research question:

How can MOOS, as a shelf-sensing component provider, position itself in the retail-technology ecosystem around actionable shelf insight?

MOOS can most credibly position itself as a bounded shelf intelligence and trigger layer between shelf reality and store action. This position means that MOOS should not be understood only as a sensor supplier, because the value of shelf sensing does not come from detection alone. At the same time, MOOS should not be positioned as a full retail operating system, because that would require ownership of workflows, systems, and decisions that sit outside its current role.

The proposed position is bounded because it focuses on the part of the value chain where MOOS can plausibly create distinctive value: reliable shelf-event sensing, interpretation, and the translation of shelf conditions into actionable signals. It is ecosystem-oriented because these signals only become useful when they can connect to the actors, routines, and systems that enable store action.

The final roadmap should therefore be read as a decision-support logic rather than as a fixed rollout plan. Its purpose is to help assess which opportunities fit the current evidence base, which opportunities require further proof, and where additional ecosystem alignment would be needed before broader adoption becomes realistic.

9.2 Answer to SQ1: structure and actionability

The first subquestion asked:

What activity structure, actor clusters, and actionability conditions are required to turn shelf-condition information into store action?

Actionable shelf insight requires an activity chain that starts at the shelf and continues into store action. A shelf event must be sensed, interpreted, filtered, routed, and connected to a decision or action. This means that shelf-condition information only becomes valuable when it is reliable, timely, understandable, and usable within existing store routines.

This chain is distributed across several actor groups. MOOS controls the shelf-sensing component and the related event logic, but retailers, store teams, information technology departments, loss-prevention teams, system providers, and implementation partners may control other parts of the route to action. The position should therefore be designed around the full activity chain, not around sensing alone.

The main actionability conditions identified in the study are reliable signals, fast enough actionability, and low additional workload. These conditions are practical rather than abstract. Store teams need to trust the signal, understand what action follows, and fit that action into daily work without creating an unrealistic extra burden.

9.3 Answer to SQ2: dependencies and alignment

The second subquestion asked:

Where do key dependencies, bottlenecks, and misalignments occur across ecosystem positions and links?

The main dependency is the link between shelf signal and store action. Shelf sensing can identify a change at the shelf, but value may break down if the signal cannot be interpreted, routed, owned, or acted on. This makes the route from event to action the central alignment problem.

The findings show that actionability depends on several connected conditions. Reliability can be affected by sensing quality, product context, store changes, installation quality, filtering logic, and surrounding data. Actionability depends on who receives the signal, who owns the response, and whether the response fits existing routines. These conditions are not controlled by one actor alone.

This means that alignment must be assessed at link level. The question is not only whether each actor or system works in isolation, but whether the handovers between actors, systems, and routines are clear enough to make shelf insight usable. The central bottleneck is therefore not simply detection, but the repeatable movement from shelf event to accountable action.

9.4 Answer to SQ3: risks and feasibility

The third subquestion asked:

How do co-innovation risk and adoption-chain risk affect the feasibility of different positioning options for MOOS?

Co-innovation risk appears when the value of shelf sensing depends on other technologies, systems, or partners. In this thesis, this risk is visible because shelf signals often need to connect to adjacent layers such as store systems, operational workflows, security routines, replenishment processes, or data environments. These connections can strengthen the value proposition, but they also increase dependency.

Adoption-chain risk appears because several actors may need to approve, implement, use, maintain, and scale the solution before value can be realised. Store teams may experience the operational problem, but other actors may control budgets, integrations, data access, security decisions, or rollout processes. As a result, a technically promising solution can still fail if the adoption chain is not aligned.

Different use cases create different risk profiles. Some depend more on direct and timely response, while others depend more on system integration, data quality, workflow ownership, and backend correction. This means that positioning options cannot be evaluated only by market attractiveness. They must also be evaluated by dependency load, proof requirements, implementation burden, and internal readiness.

9.5 Answer to SQ4: points of influence and readiness

The fourth subquestion asked:

Which points of influence and internal readiness conditions should guide MOOS's strategic positioning?

The main point of influence for MOOS lies between shelf reality and store action. MOOS can create value by improving the reliability, interpretation, and routing of shelf-event information. This is a stronger position than offering raw visibility alone, because it connects sensing to the conditions under which retailers can actually act.

However, this point of influence only becomes credible if MOOS can make its role repeatable. Internal readiness therefore matters. MOOS needs the ability to define clear use-case boundaries, support reliable deployment, manage partner and system dependencies, document repeatable implementation logic, and build evidence that can be transferred across contexts.

The study also shows that MOOS should avoid overextending its role. A broader position may appear attractive, but it would increase dependency and execution burden. A narrower sensor-only position would reduce complexity, but it would also leave too much of the value creation outside MOOS's influence. The preferred position sits between these extremes: bounded enough to remain feasible, but connected enough to make shelf insight actionable.

9.6 Final remarks

This thesis shows that shelf sensing becomes strategically relevant when it helps close the gap between what systems assume and what is actually happening at the shelf. The value of MOOS's technology therefore depends not only on technical detection, but on whether shelf events can become trusted, routable, and accountable signals for store action.

The main contribution of the thesis is the translation of this problem into an ecosystem-based positioning logic. By using ecosystem-as-structure, the thesis shows how actionable shelf insight depends on activities, actors, positions, links, alignment conditions, risks, and internal readiness. Positioning is therefore treated as an outcome of diagnosis, not as a predefined market label.

The final answer is deliberately bounded. The research does not prove market success, long-term rollout, or commercial performance. It provides a structured direction for how MOOS can think about its role in the retail-technology ecosystem. MOOS's most credible position is not to own the entire store technology architecture, but to make shelf reality usable for action through a focused shelf intelligence and trigger-layer role.

Appendix

Appendix A. Official project documentation

A.1 Approved project brief

A.2 Human Research Ethics Committee approval

Appendix B. Supplementary context analysis

B.1 Full DEPEST

Appendix C. Store visit research material

C.1 Store visit protocol

C.2 Store visit question list

Appendix D. Expert interview research material

D.1 Interview protocol

Excluded internal appendices

The following appendix materials are excluded from this repository version because they contain company-sensitive information, including internal roadmap assumptions, customer references, partner logic, confidential co-creation material, validation input, or detailed design choices:

Appendix E. Detailed analysis material
Contains detailed coding, raw synthesis, and internal interpretation material.

Appendix F. Design and co-creation material Contains confidential workshop material, route discussions, and design assumptions.

Appendix G. Roadmap development material Contains detailed roadmap logic, internal choices, and possible sequencing.

Appendix H. Validation material
Contains confidential validation feedback and roadmap revision details.

Appendix A. Official project documentation

A.1 Approved project brief

DESIGN FOR our future

TU Delft

IDE Master Graduation Project

Project team, procedural checks and Personal Project Brief

In this document the agreements made between student and supervisory team about the student's IDE Master Graduation Project are set out. This document may also include involvement of an external client, however does not cover any legal matters student and client (might) agree upon. Next to that, this document facilitates the required procedural checks:

- Student defines the team, what the student is going to do/deliver and how that will come about
- Chair of the supervisory team signs, to formally approve the project's setup / Project brief
- SSC E&SA (Shared Service Centre, Education & Student Affairs) report on the student's registration and study progress
- IDE's Board of Examiners confirms the proposed supervisory team on their eligibility, and whether the student is allowed to start the Graduation Project

STUDENT DATA & MASTER PROGRAMME

Complete all fields and indicate which master(s) you are in

Family name	de Meester	IDE master(s)	IPD ☐ Dfl ☐ SPD ✔
Initials	G. T.	2nd non-IDE master	
Given name	Gijs	Individual programme *(date of approval)*	
Student number	4876237	Medisign	☐
		HPM	☐

SUPERVISORY TEAM

Fill in he required information of supervisory team members. If applicable, company mentor is added as 2nd mentor

Chair	Shahrok Nikou	dept./section	DOS - RMCR
mentor	Nikolaos Kyriakopoulos	dept./section	DOS - RMCR
2nd mentor			
client:	Marc Reunis		
city:	Amsterdam, Delft	country:	Netherlands
optional comments	The chair and mentor represent different sub-groups and bring complementary expertise. One focuses on digital transformation and ecosystem strategy, the other on scenario-based design and strategic foresight. This combination is essential to cover the full scope of the project.		

! Ensure a heterogeneous team. In case you wish to include team members from the same section, explain why.

! Chair should request the IDE Board of Examiners for approval when a non-IDE mentor is proposed. Include CV and motivation letter.

! 2nd mentor only applies when a client is involved.

APPROVAL OF CHAIR on PROJECT PROPOSAL / PROJECT BRIEF -> to be filled in **by the Chair** of the supervisory team

Sign for approval (Chair)

Shahrokh Nikou — Digitally signed by Shahrokh Nikou Date: 2025.10.28 11:17:25 +01'00'

Name Shahrokh Nikou Date 20 Oct 2025 Signature

CHECK ON STUDY PROGRESS

To be filled in **by SSC E&SA** (Shared Service Centre, Education & Student Affairs), after approval of the project brief by the chair.
The study progress will be checked for a 2nd time just before the green light meeting.

Master electives no. of EC accumulated in total	EC
Of which, taking conditional requirements into account, can be part of the exam programme	EC

★	YES	all 1st year master courses passed
☐	NO	missing 1st year courses

Comments:

i.v.m. curriculumherziening een MSc 1 ruimte van 58EC

Sign for approval (SSC E&SA)

K. Veldman — Digitally signed by K. Veldman Date: 2025.10.30 15:19:26 +01'00'

Name K. Veldman Date 30 Oct 2025 Signature

APPROVAL OF BOARD OF EXAMINERS IDE on SUPERVISORY TEAM -> to be checked and filled in by IDE's Board of Examiners

Does the composition of the Supervisory Team comply with regulations?

YES	★	Supervisory Team approved
NO	☐	Supervisory Team not approved

Comments:

Based on study progress, students is ...

★	**ALLOWED** to start the graduation project
☐	**NOT** allowed to start the graduation project

Comments:

Sign for approval (BoEx)

Paul Mommers — Digitally signed by Paul Mommers Date: 2025.12.05 10:01:42 +01'00'

Name Paul Mommers Date 5 Dec 2025 Signature

DESIGN FOR our future

TU Delft

Personal Project Brief – IDE Master Graduation Project

Name student Gijs de Meester **Student number** 4,876,237

PROJECT TITLE, INTRODUCTION, PROBLEM DEFINITION and ASSIGNMENT
Complete all fields, keep information clear, specific and concise

Project title Strategic Positioning of MOOS in the Future Retail Ecosystem

Please state the title of your graduation project (above). Keep the title compact and simple. Do not use abbreviations. The remainder of this document allows you to define and clarify your graduation project.

Introduction

Describe the context of your project here; What is the domain in which your project takes place? Who are the main stakeholders and what interests are at stake? Describe the opportunities (and limitations) in this domain to better serve the stakeholder interests. (max 250 words)

Retail is undergoing a fundamental transformation, driven by digital technologies that change how value is created and delivered. Traditionally, in-store operations have relied on manual labor, leading to inefficiencies such as stockouts and theft, also known as shelf shrinkage. Rising labor costs and staff shortages have increased the urgency to adopt new digital solutions (Mostaghel, 2022).

In response, retailers have introduced technologies such as electronic shelf labels, RFID tagging, and point-of-sale analytics to enhance operational efficiency and customer experience (Inman & Nikolova, 2017). While these tools improve specific processes, they often function as isolated solutions and fail to address systemic challenges (Roggeveen et al., 2020). Recent developments point towards greater integration, where technologies such as camera vision and shelf-sensing are combined to enable formats like unmanned or scan-and-go stores (Thaichon, 2024).

This shift illustrates that competition in retail no longer takes place between individual firms but between ecosystems of interconnected actors and stakeholders. Ecosystem theory describes these constellations as networks of firms, technologies, and institutions that co-create value (Palmié et al., 2022). Retail ecosystems increasingly operate as “meta-ecosystems,” where physical and digital layers interact to deliver seamless shopping experiences. This trend creates opportunities, but also uncertainties, for smaller technology providers aiming to secure a role in these evolving systems.

The context of this project is a technology provider company known as MOOS. MOOS has developed a shelf-sensing technology that generates real-time product-level insights. This capability has the potential to improve inventory accuracy, operational efficiency, and customer satisfaction. However, to capture value from this innovation, MOOS must navigate the complexities of the retail-tech ecosystem and identify how to position itself strategically within it.

→ *space available for images / figures on next page*

DESIGN FOR our future

TU Delft

Personal Project Brief – IDE Master Graduation Project

Problem Definition

What problem do you want to solve in the context described in the introduction, and within the available time frame of 100 working days? (= Master Graduation Project of 30 EC). What opportunities do you see to create added value for the described stakeholders? Substantiate your choice.
(max 200 words)

Although MOOS provides a technologically advanced shelf-sensing solution, its long-term strategic position within the evolving retail-tech ecosystem remains unclear. Retailers are increasingly adopting a variety of digital technologies such as camera vision, RFID tagging, electronic shelf labels, and point-of-sale analytics to improve operational efficiency and customer experience. However, these technologies are often developed and implemented in isolation, resulting in a fragmented landscape where integration and interoperability remain limited.

Ecosystem research shows that innovation outcomes are shaped not only by the performance of individual technologies but also by how firms coordinate and align their roles within broader systems of collaboration (He et al., 2023). For MOOS, this complexity creates uncertainty about how its shelf-sensing technology should be positioned in relation to complementary or competing solutions. As multiple actors including large retailers, technology vendors, and system integrators are competing to define new standards and data architectures, MOOS faces the risk of being technologically capable yet strategically overlooked. The challenge lies in understanding how different technologies and players interact within the emerging retail-tech ecosystem and in identifying strategic opportunities for MOOS to create and capture value in this evolving environment.

Assignment

This is the most important part of the project brief because it will give a clear direction of what you are heading for. Formulate an assignment to yourself regarding what you expect to deliver as result at the end of your project. (1 sentence) As you graduate as an industrial design engineer, your assignment will start with a verb (Design/Investigate/Validate/Create), and you may use the green text format:

Design a strategic positioning framework and roadmap to guide MOOS in identifying and developing future business and partnership strategies within the evolving retail-tech ecosystem.

Then explain your project approach to carrying out your graduation project and what research and design methods you plan to use to generate your design solution (max 150 words)

This research combines three complementary theoretical perspectives to analyse and design MOOS’s strategic positioning within the evolving retail-tech ecosystem. Innovation Ecosystem Theory (Adner, 2017) forms the foundation for analysing the external environment, explaining how interdependent actors within the retail-tech ecosystem co-create and capture value through alignment and collaboration. Dynamic Capabilities Theory (Teece, 2007) focuses on the internal dimension, describing how MOOS can sense opportunities, seize them through strategic action, and reconfigure its resources and partnerships to remain adaptive in a changing context. Strategic Foresight Theory (Rohrbeck, 2012) provides the temporal bridge between today and tomorrow, enabling the exploration of technological and societal developments that may shape the future of retail and inform long-term strategic choices.

Together, these frameworks connect environmental analysis, organizational adaptability, and future orientation, supporting both the analytical and generative phases of the project. This integrated lens guides the creation of a strategic positioning framework and roadmap that enable MOOS to make informed, future-ready decisions.

Project planning and key moments

*To make visible how you plan to spend your time, you must make a planning for the full project. You are advised to use a Gantt chart format to show the different phases of your project, deliverables you have in mind, meetings and in-between deadlines. Keep in mind that all activities should fit within the given run time of 100 working days. Your planning should include a **kick-off meeting**, **mid-term evaluation meeting**, **green light meeting** and **graduation ceremony**. Please indicate periods of part-time activities and/or periods of not spending time on your graduation project, if any (for instance because of holidays or parallel course activities).*

Make sure to attach the full plan to this project brief.
The four key moment dates must be filled in below

Kick off meeting	14 Oct 2025
Mid-term evaluation	12 Jan 2026
Green light meeting	24 Mar 2026
Graduation ceremony	21 Apr 2026

In exceptional cases (part of) the Graduation Project may need to be scheduled part-time. Indicate here if such applies to your project

Part of project scheduled part-time	✔
For how many project weeks	25
Number of project days per week	4,0

Comments:
In this planning I work 4 days in the week on my graduation. I do have Christmas holiday, but no TU holiday between Q2 and Q3.

Motivation and personal ambitions

Explain why you wish to start this project, what competencies you want to prove or develop (e.g. competencies acquired in your MSc programme, electives, extra-curricular activities or other).

Optionally, describe whether you have some personal learning ambitions which you explicitly want to address in this project, on top of the learning objectives of the Graduation Project itself. You might think of e.g. acquiring in depth knowledge on a specific subject, broadening your competencies or experimenting with a specific tool or methodology. Personal learning ambitions are limited to a maximum number of five.
(200 words max)

What excites me most about this graduation project is the opportunity to bring together the three aspects of design that I value most: strategy, technology, and the ability to make sense of complex systems. During my master's in Strategic Product Design, I have learned how design can connect human needs, business goals, and technological developments. For me, design is not just about creating solutions, but also about shaping long-term direction and purpose.

Collaborating with MOOS feels like a natural fit. As a young and growing tech start-up, MOOS offers a dynamic and hands-on environment. I see it as a chance to be part of daily decisions and learn from real-life challenges. Something I have missed in the more academic parts of my studies. I'm especially interested in how design and strategy come together in this context, and I hope to gain a better understanding of what kind of team, organization, and culture suits me best.

I also want to grow in specific areas. One of my goals is to deepen my knowledge of ecosystem strategy and digital innovation. I am curious how smaller players like MOOS can find their position and voice within larger systems. At the same time, I want to improve my ability to translate complex insights into clear, visual frameworks that support better decisions. In this way, I hope to combine design thinking with real impact in the field.

A.2 Human Research Ethics Committee approval

Date 02-Feb-2026
Correspondence hrec@tudelft.nl

TU Delft Delft University of Technology

Human Research Ethics Committee TU Delft
(http://hrec.tudelft.nl)

Visiting address
Jaffalaan 5 (building 31)
2628 BX Delft

Postal address
P.O. Box 5015 2600 GA Delft
The Netherlands

Ethics Approval Application: Strategic Positioning of MOOS in the Future Retail Ecosysten
Applicant: Meester, Gijs de

Dear Gijs de Meester,

It is a pleasure to inform you that your application mentioned above has been approved.

Thanks very much for your submission to the HREC, which has been approved. We do additionally note/advise the following:
1. Please use NVIVO instead of Miro for the qualitative data analysis. Miro is not considered GDPR compliant.
2. You mention that the collected data can be sensitive to competitive advantage or brand reputation. Our suggestion is to ask the interviewees to read the transcript and, if they think the information in it can be shared, you are fine.

Good luck with your research!

Sincerely,

Dr. C. Shelley-Egan
Chair HREC
Faculty of Technology, Policy and Management

H: More on Informed Consent and Data Management

NOTE: You can find guidance and templates for preparing your Informed Consent materials) here

Your research involves human participants as Research Subjects if you are recruiting them or actively involving or influencing, manipulating or directing them in any way in your research activities. This means you must seek informed consent and agree/ implement appropriate safeguards regardless of whether you are collecting any PIRD.

Where you are also collecting PIRD, and using Informed Consent as the legal basis for your research, you need to also make sure that your IC materials are clear on any related risks and the mitigating measures you will take – including through responsible data management.

Got a comment on this checklist or the HREC process? You can leave your comments here

IV. Signature/s

Please note that by signing this checklist list as the sole, or Responsible, researcher you are providing approval of the completeness and quality of the submission, as well as confirming alignment between GDPR, Data Management and Informed Consent requirements.

Name of Corresponding Researcher (if different from the Responsible Researcher) (print)

Signature of Corresponding Researcher:

Date:

Name of Responsible Researcher (print)

Shahrokh Nikou
Signature (or upload consent by mail) Responsible Researcher:

Date: 10.11.2025

V. Completing your HREC application

Please use the following list to check that you have provided all relevant documentation

Required:

- **Always:** This completed HREC checklist
- **Always:** A data management plan (reviewed, where necessary, by a data-steward)
- **Usually:** A complete Informed Consent form (including Participant Information) and/or Opening Statement (for online consent)

Appendix B. Supplementary Context Analysis.

B.1 Full Depest

Domain	*TrendTitle*	*Description*	*APA7Reference*
DEMOGRAPHIC	Older shoppers shift demand to health, convenience and nearby formats	Ageing populations in Europe and North America increase demand for easy-to-shop formats, health-oriented offers, and stores closer to home.	McKinsey & Company, & EuroCommerce. (2022, October 25). Transforming the EU retail and wholesale sector. https://www.mckinsey.com/industries/retail/our-insights/transforming-the-eu-retail-and-wholesale-sector
DEMOGRAPHIC	Ageing retail workforce increases pressure on automation and reskilling	Retailers face an ageing workforce and skills shortages, pushing them to automate more work and invest in reskilling.	McKinsey & Company, & EuroCommerce. (2022, October 25). Transforming the EU retail and wholesale sector. https://www.mckinsey.com/industries/retail/our-insights/transforming-the-eu-retail-and-wholesale-sector
DEMOGRAPHIC	Time-poor shoppers increase demand for reduced-friction shopping	Retail pressure on time and convenience supports demand for lower-friction shopping and technologies that simplify store execution.	McKinsey & Company, & EuroCommerce. (2022, October 25). Transforming the EU retail and wholesale sector. https://www.mckinsey.com/industries/retail/our-insights/transforming-the-eu-retail-and-wholesale-sector
ECONOMIC	Margin pressure makes productivity and automation top priorities	Weak volume growth and rising labour, energy, and technology costs make productivity gains and automation central retail priorities.	McKinsey & Company, & EuroCommerce. (2022, October 25). Transforming the EU retail and wholesale sector. https://www.mckinsey.com/industries/retail/our-insights/transforming-the-eu-retail-and-wholesale-sector
ECONOMIC	Value-seeking consumers accelerate private label and discount growth	Consumers under financial pressure are trading down, comparing prices more aggressively, and shifting spend toward cheaper options such as private label and discount formats.	Deloitte. (2025, February 6). Retail trends 2025. https://www.deloitte.com/no/no/Industries/consumer/perspectives/retail-trends.html
ECONOMIC	Product availability remains a baseline trust signal and recurring loss driver	Availability is increasingly treated as a trust baseline: shoppers expect the right product at the right time, and OOS remains a recurring commercial loss driver.	Deloitte. (2025, November 19). The revolution and evolution of retail stores. https://www.deloitte.com/us/en/industries/consumer/articles/revolution-evolution-of-retail-stores.html
ECONOMIC	Shrink and organised retail crime become structural cost drivers	Retail shrink remains a material cost item and is strongly linked to organised retail crime, theft, and staff safety concerns.	National Retail Federation. (2023). National retail security survey 2023. https://nrf.com/research/national-retail-security-survey-2023
ECONOMIC	Organized retail crime becomes AI-augmented and cross-channel	ORC is no longer only a store problem: NRF flags generative AI, digital fraud, supply-chain coordination, and cross-channel attack patterns as rising risks.	National Retail Federation. (2026, January 7). NRF experts on what to watch in 2026. https://nrf.com/blog/nrf-experts-on-what-to-watch-in-2026
ECONOMIC	Retail media opens new profit pools beyond product margin	Retailers are actively looking for alternative revenue streams, with retail media standing out as a new top- and bottom-line growth engine.	Deloitte. (2025, February 6). Retail trends 2025. https://www.deloitte.com/no/no/Industries/consumer/perspectives/retail-trends.html
ECONOMIC	Faster AI-related ROI expectations shape retail tech investment choices	Retail leaders increasingly expect AI-led supply-chain and operational initiatives to show positive ROI within short time horizons, tightening investment discipline.	Deloitte. (2026, the company mentorh 11). 2026 global retail industry outlook. https://www.deloitte.com/mt/en/Industries/consumer/perspectives/global-retail-industry-outlook.html
POLITICAL	FSMA 204 pushes food retailers and suppliers toward end-to-end digital traceability	The FDA's Food Traceability Rule keeps pushing high-risk food chains toward digital recordkeeping and end-to-end traceability capabilities.	U.S. Food and Drug Administration. (2025, August 6). FDA proposes to extend compliance date for Food Traceability Rule and issues new FAQs and other resources. https://www.fda.gov/food/hfp-constituent-updates/fda-proposes-extend-compliance-date-food-traceability-rule-and-issues-new-faqs-and-other-resources
POLITICAL	Digital Product Passport shifts from policy concept to implementation architecture	The Digital Product Passport is moving from broad regulation into concrete implementation work around identifiers, access rights, registries, and product-level data exchange.	European Commission. (2025, April 9). Commission launches consultation on the Digital Product Passport. https://single-market-economy.ec.europa.eu/news/commission-launches-consultation-digital-product-passport-2025-04-09_en
POLITICAL	The AI Act increases the need for safe and trustworthy retail AI deployment	EU AI regulation makes AI governance, transparency, and risk management more important for retailers deploying AI in operations, analytics, and customer-facing flows.	European Commission. (n.d.). European approach to artificial intelligence. https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence
POLITICAL	The Data Act raises the importance of usable, shareable and well-governed industrial data	As connected devices generate more operational data, the EU Data Act raises the strategic importance of clear data access, sharing, and governance rules.	European Commission. (2025, December 15). Data Act explained. https://digital-strategy.ec.europa.eu/en/factpages/data-act-explained
POLITICAL	Policy focus on retail crime and staff safety drives tech-based loss prevention	Retail crime and violence are increasingly treated as policy and governance issues, which strengthens the case for tech-enabled loss prevention.	National Retail Federation. (2023). National retail security survey 2023. https://nrf.com/research/national-retail-security-survey-2023
ECOLOGICAL	Sustainability, digitalisation and skills form a triple transition for EU retail	European retail is under simultaneous pressure to invest in sustainability, digitalisation, and workforce capabilities rather than treating them as separate agendas.	McKinsey & Company, & EuroCommerce. (2022, October 25). Transforming the EU retail and wholesale sector. https://www.mckinsey.com/industries/retail/our-insights/transforming-the-eu-retail-and-wholesale-sector
ECOLOGICAL	Energy-efficient stores and DCs become core decarbonisation levers	Energy efficiency in stores, refrigeration, and distribution centres is becoming a core path to reduce both emissions and operating costs.	McKinsey & Company, & EuroCommerce. (2022, October 25). Transforming the EU retail and wholesale sector. https://www.mckinsey.com/industries/retail/our-insights/transforming-the-eu-retail-and-wholesale-sector
ECOLOGICAL	Food waste reduction becomes a strategic operating KPI	Retailers are under rising pressure to cut food waste through better forecasting, labelling, rotation, and operational processes.	European Commission. (n.d.). Food waste. https://food.ec.europa.eu/safety/food-waste_en
ECOLOGICAL	2D product data carriers help reduce waste and improve expiry management	GS1's retail 2D barcode guidance shows how richer product-level data can support expiry control, automatic markdowns, traceability, and waste reduction.	GS1. (n.d.). 2D barcodes at retail point-of-sale implementation guideline. https://ref.gs1.org/guidelines/2d-in-retail/
SOCIALCULTURAL	Consumers expect seamless, engaging and low-friction shopping journeys	Retailers are expected to deliver easier, more invisible and more engaging journeys across physical and digital touchpoints.	EY. (2023, February 17). How retailers' value propositions need to evolve for success. https://www.ey.com/en_gl/insights/consumer-products/defining-the-future-of-retail
SOCIALCULTURAL	Stores are being reimagined as service, fulfilment, brand and community ecosystems	The physical store is not dienterprise-system proofpearing; it is taking on broader roles in fulfilment, service, community-building, and brand experience.	Deloitte. (2025, November 19). The revolution and evolution of retail stores. https://www.deloitte.com/us/en/industries/consumer/articles/revolution-evolution-of-retail-stores.html
SOCIALCULTURAL	Demand for healthy, transparent and trustworthy products grows	Consumers increasingly reward clear labels, healthier products, and trustworthy sustainability and ingredient information.	EY. (2023, February 17). How retailers' value propositions need to evolve for success. https://www.ey.com/en_gl/insights/consumer-products/defining-the-future-of-retail
SOCIALCULTURAL	Retail media shifts toward closed-loop, omnichannel measurement	Retail media is expanding from impression-selling toward closed-loop, omni-channel measurement that links exposure to in-club, online, and repeat-purchase outcomes.	Sam's Club MAP. (n.d.). The new definition of retail media: Omnichannel, full-funnel, closed-loop. https://map.samsclub.com/resources/articles/new-retail-media
TECHNOLOGICAL	Retailers overhaul tech architecture and operating models to unlock store tech value	McKinsey frames tech transformation in retail as an architecture and operating-model challenge, not just a tooling challenge.	McKinsey & Company. (2022, May 20). The tech transformation imperative in retail. https://www.mckinsey.com/industries/retail/our-insights/the-tech-transformation-imperative-in-retail
TECHNOLOGICAL	Software-driven operating models become a competitive capability for retailers	Retailers are under pressure to behave more like software organizations, with stronger software cultures, operating models, and customer-centric innovation capabilities.	McKinsey & Company. (2023, September 6). Turning consumer and retail companies into software-driven innovators. https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/turning-consumer-and-retail-companies-into-software-driven-innovators
TECHNOLOGICAL	AI adoption shifts from experimentation to operational use in retail and supply chains	Retail AI is moving from broad experimentation toward practical operational deployment, especially in supply-chain visibility, optimization, and decision support.	Deloitte. (2026, January 8). 2026 retail industry outlook. https://www.deloitte.com/us/en/insights/industry/retail-distribution/retail-distribution-industry-outlook.html
TECHNOLOGICAL	In-store automation expands through RFID, smart carts, sensors and connected workflows	Retail technology continues to broaden through RFID, smart carts, traffic sensors, and other building blocks for more instrumented store execution.	McKinsey & Company, & EuroCommerce. (2022, October 25). Transforming the EU retail and wholesale sector. https://www.mckinsey.com/industries/retail/our-insights/transforming-the-eu-retail-and-wholesale-sector
TECHNOLOGICAL	Electronic shelf labels turn the shelf into a managed digital surface	Shelf infrastructure such as ESL is increasingly part of the digital store stack, enabling more central control of pricing and promotion execution.	Deloitte. (2025). Global powers of retailing 2025. https://www.deloitte.com/global/en/industries/retail/research/global-powers-of-retailing.html
TECHNOLOGICAL	Agentic commerce becomes a new decision layer before checkout	AI agents increasingly influence product discovery and decision-making before purchase, creating a new layer above retailer interfaces.	McKinsey & Company. (2025, October 17). The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-agentic-commerce-opportunity-how-ai-agents-are-ushering-in-a-new-era-for-consumers-and-merchants
TECHNOLOGICAL	Merchandising shifts from dashboards to agentic AI workflows	Retail merchandising is starting to move from reporting support toward agentic AI workflows that reshape roles, decisions, and ways of working.	McKinsey & Company. (2026, January 9). Merchants unleashed: How agentic AI transforms retail merchandising. https://www.mckinsey.com/industries/retail/our-insights/merchants-unleashed-how-agentic-ai-transforms-retail-merchandising
TECHNOLOGICAL	Physical AI moves closer to practical deployment in stores and supply chains	Deloitte describes physical AI and robotics as moving toward practical enterprise use as reliability, economics, and hardware maturity improve.	Deloitte. (2025, December 10). Physical AI and humanoid robots. https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2026/physical-ai-humanoid-robots.html
TECHNOLOGICAL	QR codes powered by GS1 move from pilots to live supermarket infrastructure	Next-generation barcodes are moving from trials into live supermarket infrastructure, making richer product, traceability, and consumer data more operationally usable.	GS1 UK. (2026, April 17). a large supermarkt becomes the first UK supermarket to transition an entire product range to QR codes powered by GS1. https://www.gs1uk.org/insights/news/a large supermarkt-becomes-the-first-UK-supermarket-to-transition-an-entire-product-range-to-QR-codes-powered-by-GS1
TECHNOLOGICAL	Retailers reduce tech debt to unlock faster AI-led innovation	Retail leaders are explicitly reducing technology debt so they can innovate faster, integrate AI more effectively, and scale operational improvements.	Deloitte. (2026, the company mentorh 11). 2026 global retail industry outlook. https://www.deloitte.com/mt/en/Industries/consumer/perspectives/global-retail-industry-outlook.html
TECHNOLOGICAL	Retail AI governance becomes a capability, not just a compliance topic	Retailers are treating responsible deployment, governance, and risk management as integral capabilities for scaling AI rather than after-the-fact controls.	National Retail Federation. (2025, December 17). Retail AI Trends 2025. https://nrf.com/research/retail-ai-trends-2025

C.Store Visit Research Material

C1. Store Visit Protocol

Part 1. Intervieew Protocol – MOOS (EN)

1. Goal and Stance

The purpose is to understand store-operator perspectives on:

a) To identify which adoption attributes matter most to store operators when implementing shelf sensing
b) To capture short explanations and concrete examples for the top priorities
c) To generate comparable data across many stores, with light qualitative colour

Interviewer stance and quality requirements:

- Neutral and curious, not selling and not defending the technology
- Focus on store reality and trade-offs, not ideals
- Ask for examples and last week situations, not opinions only
- Avoid leading language like better, best, should, obviously

1. Checklist to Bring

- Printed one-page info and consent sheet, short version
- Printed ranking sheet or phone version with the 8 attributes en de open vragen
- Pen plus backup pen
- Phone timer
- Small notebook for notes
- TU Delft student card if you have one
- Small visual of the sensor, only if they ask what it looks like

1.

2. Checklist for start

- Confirm they have 10 to 15 minutes
- Confirm role and whether they know store routines
- Explain anonymity, no store name, no person name
- Ask permission to take notes
- If you want audio, only do it if it is truly feasible in-store. Otherwise do notes only
- Assign participant code

2.

3. Introductory script

Hi, my name is Gijs. I am doing my master thesis at TU Delft.

I am collecting short inputs from store teams to understand what matters most in daily work if a shelf sensing solution is introduced. This is not a sales conversation. There are no right or wrong answers.

Your answers are anonymous. I will not write down your name or your store name. You can stop at any moment.

I will ask you to rank 8 attributes and then ask one or two short follow-up questions. It takes about 10 to 15 minutes. Is that okay?

To ensure we talk about the same concept, I will first show a short MOOS introduction video (about 1 minute). After the video.

[Show Video]

The context of this research is at the MOOS. MOOS is in developing shelf-level sensing technology to measure product presence and movement and turn this into real-time shelf event signals and insights. The technology helps stores better understand what is happening on the shelf and to support timely action by store teams.

In practical terms, MOOS aims to help with two operational domains:

- On-Shelf Availability (OSA): detecting when products are missing from the shelf and supporting replenishment actions. Out-of-Stock (OOS) is when there is nothing available on the shelf.
- Loss/shrink/theft: detecting shelf situations that may indicate loss risk and supporting appropriate in-store responses.

MOOS can operate as a stand-alone component, but the value for store teams may increase when MOOS' technology is combined with other technologies already used in retail. Examples of such combinations (for context only) include:

[examples were discussed at category level only]

6. Checklist for closure

- Thank them
- Confirm anonymity again
- Ask if I may contact them later for one clarification, optional
- Immediately after, write my own notes on the why explanation and the examples

C2. Store Visit Question list

1. Interview question guide

Note: Ask the core questions in order. Use probes selectively based on the participant's role and the depth needed. Total target: 9–12 core questions plus closing.

1. Agreement:

- ☐ I agree to participate in this 10–15 minute anonymous interview for TU Delft research, where the researcher will take notes; I understand I can skip questions or stop at any time.
- ☐ I do not agree.

2. What is your role in this store?

Sector:
- ☐ Grocery
- ☐ Convenience
- ☐ Health and beauty

Store format
- ☐ Small
- ☐ Medium
- ☐ Large

Governance
- ☐ Branch
- ☐ Franchise
- ☐ Independent

Location type
- ☐ Urban
- ☐ Rural

3. Below are eight features of a shelf sensing solution. Could you rank them from 1 (most important) to 8 (least important), based on how important you think these features are for integration and use in this shop?

1. 1 means most important for your store team to make this workable, 8 means least important
2. ___ No extra staff work
3. ___ Hard to notice in store
4. ___ Reliable signals
5. ___ Fast enough to act
6. ___ Low effort to keep it running
7. ___ Fits how the store already works
8. ___ Works with existing store technology
9. ___ Clear ownership

4. You ranked the attributes. Explain the most important:

_____

_____

_____

- Can you give a concrete example from last week
- What happens today when this goes wrong
- Who is affected most, customers, staff, manager

5. If you must pick one main reason to improve the shelf in your store, what is the priority today, and why?

- ○ Availability and out of stock
- ○ Loss and shrink
- ○ Other:_____

6. What is the biggest practical barrier that would stop your store from using a shelf sensing solution successfully?

_____

_____

_____

- Is it time, staffing, trust, process fit, support, or cost
- What would need to change to remove this barrier

7. Is there anything important I did not ask that I should know?

_____

_____

_____

8. Could you give me a reference to someone in your network who also would be valuable for me to interview?

_____

_____

_____

Store type

40 antwoorden

Diagram kopiëren

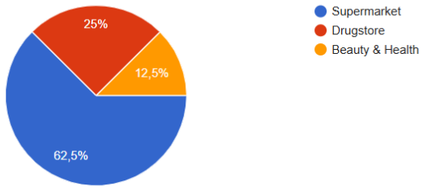

Size

41 antwoorden

Diagram kopiëren

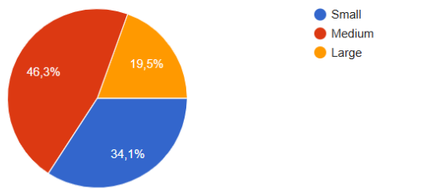

Governance

41 antwoorden

Diagram kopiëren

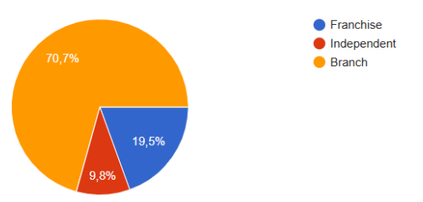

Location

41 antwoorden

Diagram kopiëren

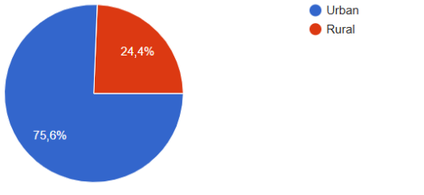

Specialty

19 antwoorden

Diagram kopiëren

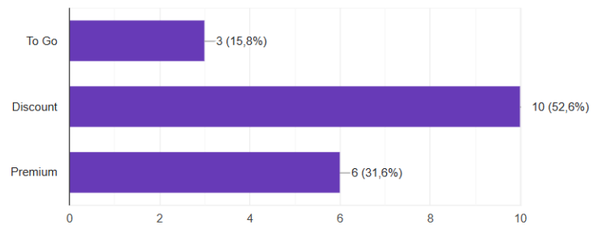

Tech level

41 antwoorden

Diagram kopiëren

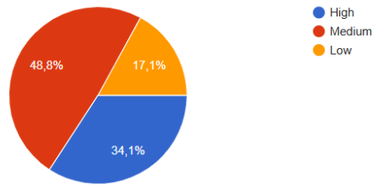

What is the main reason for your store to choose this sensor?

40 antwoorden

Diagram kopiëren

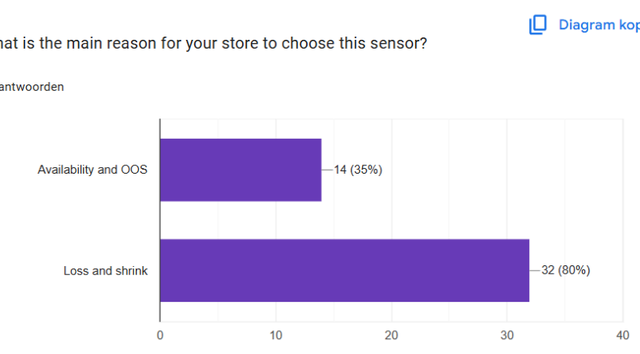

D. Expert Interview Research Material

D.1 Interview Protocol

Semi-Structured Interview Protocol – MOOS

1.Purpose and interviewer stance

Purpose

- Understand what makes shelf sensing actionable in practice, not just technically possible.
- Map the minimum chain of activities and dependencies needed to capture value at scale.
- Identify adoption gates and the conditions where deep integration wins versus light integration.
- Translate insights into positioning options for MOOS in the retail ecosystem.

Interviewer stance

1. I am not selling a product. I am learning how real rollout decisions are made and why scaling succeeds or fails.
2. I will ask for concrete examples from your experience.
3. I will challenge assumptions gently to understand trade-offs.
4. There are no right or wrong answers. I am interested in your logic and lived reality.

2.Checklist for start

- Confirm participant identity/role and allocate anonymised code (e.g., P1, P2, …).
- Confirm informed consent (incl. recording) is signed.
- Explain confidentiality and voluntary participation; confirm understanding.
- Test audio recording (and backup if available).
- Confirm duration and allow time for questions before starting.
- Note context variables (segment, format, governance, urban/rural) without identifying details.

3.Introductory script

Thank you for taking the time to participate. My name is Gijs and I am conducting interviews for my master thesis.

Before we start, did you have a chance to read the short pre-interview information document I sent?

If yes

Great, thank you. Then I will keep the context short.

If no

No problem. I will give a short neutral overview now so you know what to expect.

MOOS develops shelf-level sensing technology that can detect shelf events and turn them into signals and insights. The broader topic of this interview is not product evaluation. It is about what makes shelf information actionable at scale, and what this means for integration choices, adoption gates, and ecosystem positioning.

In this conversation I will focus on two operational domains where shelf sensing is often discussed. The first is on shelf availability and out-of-stock situations. The second is loss, shrink, and theft risk. I may also touch on related areas like inventory accuracy or planogram changes, depending on your experience.

I am especially interested in the trade-off between two integration approaches. One is deep integration with one or a few partners. The other is light integration with many partners using APIs and project-based integrations. I want to understand when each approach works best and why.

Everything you share will be treated confidentially. In my thesis I will not use your name or identifying details, and quotes will be anonymised. Your participation is voluntary. You can skip any question and you may stop at any time. There are no right or wrong answers. I am interested in your experience and concrete examples.

With your permission, I would like to audio record this conversation so I do not miss details. The recording will only be used for analysis. Is that okay?

4.Interview question guide

1. **To start, could you briefly describe your role and how you are involved in decisions about store technology, store execution, or partnerships?**
    - Which types of retailers or store formats are you most involved with? Grocery, convenience, health and beauty, or other.
    - What does success look like in your role?

1. **From your perspective, what are the most important shelf-related problems today, and why?**
    - Which problems create the most operational pain? Availability and out-of-stock, shrink and theft, inventory accuracy?
    - Which of these problems receive the most budget and attention? Which stays unsolved?

1. **When is shelf information actionable in day-to-day operations? What must be true for information to lead to action?**
    - Who needs to act, and within what time window?
    - What false-alarm rate is acceptable before people start ignoring alerts (alert fatigue)?
    - In practice, what is more valuable: high accuracy later, or 'good-enough' accuracy now?
    - What is the smallest unit of action that matters, shelf-level, section-level or item-level?

Before speaking with ecosystem stakeholders, I conducted a short field study with 41 store operators and store managers across different retail formats.

They were asked to rank eight potential adoption conditions for shelf-level sensing solutions.

Across formats, three attributes consistently ranked highest:

- *Reliable and trustworthy signals*
- *Fast enough to act upon*
- *No additional operational burden*

Attributes such as invisibility of hardware and clear ownership were ranked significantly lower.

1. **From your perspective in the broader retail ecosystem, how do you interpret these store-level priorities?**

    Recognition & alignment

    - Do these results match what you see in practice?
    - Where do you see alignment or misalignment between store-level priorities and head-office or IT-level priorities?

    Strategic tension

    - Are there attributes that store teams underestimate but that are critical for scaling?
    - Conversely, are there attributes that store teams overemphasize?

    Ecosystem implications

    - What would need to happen at system or partner level to truly deliver on reliability and actionability?
    - Which actor in the ecosystem is structurally responsible for guaranteeing reliability?
    - Does delivering on these top priorities push MOOS toward deeper integration, or can they be achieved through lighter integration?

    Governance reflection

    - If ownership was ranked low, does that reflect reality, or is ownership simply invisible until something breaks?

1. **Imagine this sensor which can detect a shelf event. What are the minimum steps that must happen next to create value at scale?**
    - What should happen operationally? Who does what?
    - What should happen technically? Which systems must receive the signal?
    - What should happen organizationally? Who owns the process and the outcome?
    - Separate the chain into installation, daily operation, and maintenance. Where does it usually break?

1. **Under what conditions is deep integration the best path, and under what conditions is light integration the best path?**
    - What deep integration means in practice in your context? Think about device management, dashboards, shared data model, embedded workflows.
    - What light integration means in practice in your context? Think about APIs, project based integrations, multiple partners, faster market coverage.
    - What is the tipping point? Store count, complexity, labor model, legacy stack, security requirements, partner leverage.
    - What is lost in each option? Speed, brand ownership, margin, flexibility, ability to scale.
    - Looking five years ahead, how do you expect shelf-level sensing to be integrated into the retail technology stack? (As part of camera systems, As part of shelf hardware, As part of retail data platforms, As a standalone infrastructure layer)

1. **Where do you most often see misalignment where the scaling of store technology fail?**
    - Technical alignment. Interfaces, reliability, network, device management.
    - Data alignment. Data ownership, definitions, source of truth, access rights.
    - Operational alignment. Workflows, task acceptance, store routines, training.
    - Incentive alignment. Who pays, who benefits, who carries workload.
    - Can you share one concrete example of a rollout that failed and why?

1. **What conditions determine whether a retailer says 'yes' or 'no' to this type of technology, and which preconditions are non-negotiable?**
    - Privacy and compliance. What is acceptable in store, what triggers a hard stop?

- Security and IT governance. Vendor risk, data residency, access control.
- Commercia. Procurement rules, payback thresholds, budget owner.
- Operational. Store workload, alert fatigue, maintenance burden.
- If you had to pick the top three conditions that decide go or no go, what are they?

1. **In the retail ecosystem, who actually determines what happens on the store floor, and who can realistically drive change?**
    - Within a retailer, who can truly drive adoption? Operations, IT, loss prevention, procurement, or category.
    - Outside the retailer, who shapes rollouts? Platform vendor, system integrator, ESL provider, camera platform, enterprise software provider.
    - If a company wants to scale a shelf sensing solution across hundreds or thousands of stores, which actor in the ecosystem ultimately controls that scaling? (Retail operations, IT architecture, Shelf hardware provider, Platform vendors (ESL, camera, analytics), System integrators)

1. **If you were advising MOOS on focus for the coming year, what would you prioritize? Why?**
    - •Would you focus on working deeply with one or two partners, or connect more lightly with several?
    - •If a startup like MOOS wanted to scale quickly, which type of partner would give them the strongest leverage? (Shelf infrastructure companies, ESL platforms, Camera / analytics platforms, System integrators, Retailers themselves)
    - •Which use case would be the most realistic entry point to scale this type of technology? (Shrink detection, On-shelf availability, Inventory accuracy, Store execution automation)
    - •What would you need to see before saying: “this can scale”?
    - •What would you consciously *not* focus on yet?

1. **What did I not ask that I should have asked to understand scaling and positioning correctly?**
    - Who else should I speak to, and why them specifically?
    - Can I follow up if I need to clarify details?

5.Checklist for closure

1. Explain next steps: transcription, thematic analysis, anonymised reporting; possible short follow-up for clarification.
2. Thank the participant, ask for the consent form and stop the recording.
3. Write immediate field notes/memos (non-identifying) right after the interview.