Adaptive Orchestration for Autonomous Airside Operations at Schiphol Airport

Design an adaptive orchestration model and maturity-aligned transition plan to facilitate the implementation of autonomous airside operations at Schiphol Airport's docking operations across multiple levels of autonomy.

By

Tessel I. Pluym

Student Number: 5058392

In partial fulfilment of the requirements for the degree of:

Master of Science

in Strategic Product Design

at the **Delft University of Technology**,

to be defended publicly on

10th of July, 2025

Thesis Committee:
Prof. mr. dr. ir. S.C. Santema, TU Delft
A. Ianniello, TU Delft
Rosina Kotey, Schiphol

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

Abstract

The aviation industry faces increasing pressure from sustainability concerns, labour shortages, and operational inefficiencies, driving airports toward automation solutions. This research addresses the challenge of orchestrating semi-autonomous and fully autonomous systems within airport operations, focusing on Schiphol Airport's vision for autonomous airside operations by 2050. While automation technologies are looking promising in isolation, the development of an integrated system remains limited. As a result, there is a risk of creating isolated solutions that can obstruct coordinated efficiency gains if not addressed.

This study uses a foresight-driven methodology combining backcasting and forecasting approaches, integrated with a specifically constructed six-layer maturity model which spans from manual operations to full autonomy. Using aircraft docking operations as a use case, the research develops a strategic roadmap for implementing remote monitoring and interventions functions that enables centralised orchestration of autonomous systems while maintaining human oversight.

Findings reveal that successful automation integration requires a phased approach prioritising human-machine collaboration over an approach that solely focuses on technology implementation without sufficient organisational and ecosystem readiness. Current airport control operations are not sufficient for managing fleets of autonomous vehicles and processes, requiring new orchestration frameworks that balance automated efficiency with human decision-making capabilities. The proposed roadmap establishes three strategic horizons necessary for reaching the future vision: *Foundation* (controlled operations with basic automation support, *Collaboration* (intelligent operations with AI-assisted decision-making, and *Autonomy* (remote and resilient operations with minimal human intervention).

To achieve this, Schiphol must adopt maturity-layered deployment strategies, leverage AI-growth strategically, and integrate co-development with change management. This research contributes a practical framework for Schiphol's transition toward autonomous operations while maintaining operational safety, workforce trust, and regulatory compliance. The findings have broader implications for complex socio-technical systems which require coordinated automation integration.

Keywords: Autonomous operations, Airport orchestration, Human-machine collaboration, Remote monitoring, AI decision-making, Schiphol Airport, Maturity framework

Preface

Digital transformation has always sparked my interest, particularly the intersection of future-forward thinking with practical implementation challenges. The opportunity to get on board with Schiphol Airport's autonomous airside operations vision represents exactly the kind of project that excites me, one with significant impact potential that requires careful alignment of *technology, people, and organisational capabilities*.

My interest in this research was sparked by an article about Schiphol's vision for autonomous airside operations. The idea of rethinking how a major airport could operate, by integrating advanced technologies while preserving the critical role of human oversight, immediately spoke to me. What stood out was not just the innovation itself, but the broader challenge of managing an intricate transition that could influence the future strategic direction of airport operations.

What drew me further into this challenge was its intricacy. The transition to autonomous operations is not just a technical problem, but rather a multi-faceted challenge requiring deep understanding of organisational dynamics, operational realities, human factors, and technological capabilities. The opportunity to contribute to solving this challenge and being able to step behind the scenes at Schiphol and explore its daily operations firsthand has been an exciting, enriching, and enjoyable experience.

The research process itself has been just as valuable as the results. Being part of Schiphol's innovation environment gave me a unique chance to see how airport operations work in practice. It quickly became clear that there's often a gap between academic models and day-to-day reality, which highlighted the importance of keeping the research grounded and practical.

Beyond the academic scope, this thesis has been a chance to meaningfully engage with a sector evolving along their transition. The outcomes are shaped by a combination of practical insights, strategic thinking, and ongoing conversations with those directly involved in airport operations. Throughout, the focus has been on producing something that not only contributes to academic understanding but also supports real-world decision-making.

Acknowledgements

I am grateful for the support and guidance I received throughout this project, and I would like to acknowledge the important role my supervisory team played in shaping the development of this research and its outcomes.

I would like to thank Sicco Santema, chair of my supervisory team at TU Delft, for his guidance throughout the process. His experience in both the aviation sector and academic research provided insightful and useful context for navigating the strategic aspects of the topic. His thoughtful and creative perspective consistently challenged assumptions, ensuring the work remained sharp and relevant. This approach also helped maintain academic standards, enhanced the overall quality and pushed me to think critically and outside the box. His constructive feedback and consistent encouragement were essential in developing my ideas and refining the direction of the research. His guidance ensured I approached both challenges and opportunities with an open, well-informed mindset.

I am thankful to Alessandro Ianiello, my mentor at TU Delft, whose expertise in human-technology-organisation systems and research methodology played an important role in shaping the study. His experience with research and projects on human-machine collaboration, along with his ability to navigate the broader trajectory of a project, brought both structure and clarity to the process. On top of that, his approach brought me a lot of inspiration throughout all parts of the process, from exploring ideas more deeply, considering multiple perspectives, and thinking about concepts in more than one way. His input consistently challenged me to approach the research with both critical thought and creative openness.

At Schiphol, I am grateful to Rosina Kotey, who served as my company supervisor. Her in-depth understanding of Schiphol's operations and her leading role in the autonomous airside operations vision provided a strong foundation for aligning this research with real-world innovation efforts. Her emphasis on grounding ideas in operational reality challenged me to move beyond academic reasoning and to consider the practical implications and constraints that come with implementation. This helped strengthen the relevance and applicability of the research. Rosina's role in driving Schiphol's autonomous airside operations vision made it especially inspiring to have her as a supervisor. Her insights offered valuable direction and consistently challenged me to stay grounded in the operational realities of the airport.

In addition to my supervisory team, I would like to thank the entire Innovation Hub team at Schiphol for their support and collaboration throughout this project. Their openness, willingness to share expertise, and engagement added valuable depth to the research. The opportunity to discuss ideas, test assumptions, and learn from their perspectives greatly deepened my understanding of the broader context at the airport.

Conducting this research within a setting that combined academic supervision with direct industry involvement proved to be both enriching and valuable. The combination helped ensure that the work remained grounded in practice while contributing to broader academic discussions on automation and organisational change.

Finally, I want to acknowledge the many people, formally and informally, who contributed their time, ideas, and reflections throughout this project. Their input has been important to developing the findings presented in this thesis.

Abbreviations

AI – Artificial Intelligence
APBB – Autonomous Passenger Boarding Bridge
API – Application Programming Interface
APOC - Airport Operations Control
APU – Auxiliary Power Unit
ATC – Air Traffic Control
CISS – Central Information System Schiphol
DST – Decision-Support Tool
EFDS – Electronic Flight Data Systems
FOD – Foreign Object Debris
FAA – Federal Aviation Administration
GSE - Ground Service Equipment
GPU – Ground Power Unit
IATA – International Air Transport Association
ICAO – International Civil Aviation Organization
IoT – Internet of Things
OCCR – Operationeel Controle Centrum Rail
OTP - On Time Perfomance
PBB – Passenger Boarding Bridge
VDGS – Visual Docking Guidance Systems

Executive Summary

The global aviation industry is undergoing significant transformation, driving by sustainability concerns, persistent and expected labour shortages, regulatory pressures, and operational efficiencies. Airports worldwide must adapt by integrating technological advancements to improve efficiency, reduce environmental impact, and mitigate the impact of labour shortages. Schiphol Airport, as one of Europe's busiest airports, exemplifies these challenges while leading the transition toward autonomous airside operations by 2050.

This vision arises from multiple aligning pressures. Environmental concerns have worsened following reports highlighting health risks associated with airport emissions and noise pollution for workers and surrounding communities. At the same time, structural workforces challenges and continued reliance on physically demanding, repetitive tasks further emphasise the need for intelligent automated solutions.

This research addresses a gap in airport automation strategy: while various pilots of individual autonomous technologies prove effective, their orchestration and interoperability within existing operations is underexplored. Accordingly, this study focuses specifically on designing an orchestration system and its implementation for semi-automated and autonomous processes during aircraft docking and ramp operations at Schiphol Airport.

The research question centres on how Schiphol can effectively design a control area system to manage the transition from manual to autonomous operations while ensuring effective implementation and integration. Using aircraft docking operations as a use case, the integration challenges between humans, autonomous systems, and orchestration tools is explored. Ultimately proposing solutions that balance technological advancement with essential human readiness and input.

To come to such solutions, a comprehensive foresight-driven methodology was used, combining multiple approaches and methods. The core methodology integrates backcasting (working backward from the 2050 vision) with forecasting techniques to create a balanced and holistic perspective on both necessary endpoints and realistic transition pathways.

A six-layer maturity model was developed based on various existing autonomy models, but tailored to airport operations, spanning from manual operations through controlled, intelligent, remote, and resilient operations to full autonomy. This model was aligned with the forecasting framework to create a structured development transition. Additional methods included enablers and blockers analysis, futures wheel exploration, and cross-impact matrix assessment to understand systemic interdependencies and potential acceleration or deceleration factors.

Field research at Schiphol provided operational context through stakeholder observations, informal interviews, and process mapping. This approach revealed how current operations rely heavily on human routines, verbal communication, and system workarounds, providing essential insights fro designing realistic transition strategies towards a new orchestration model.

Several critical insights emerged about automation integration in multi-faceted operational environments: The current state, orchestration gaps, human-machine collaboration requirements, and system interdependencies.

Schiphol's current aircraft docking process is heavily reliant on manual coordination, verbal communication, and fragmented tools. The existing system operates with a high degree of human involvement, leading to inefficiencies in escalation paths and response times especially with labour shortages. The lack of centralised coordination and reliance on physical presence creates a system that is not equipped to accommodate the integration of autonomous processes. This manual structure does not allow for seamless integration with autonomous systems, which could lead to challenges in scaling automation and improving operational efficiency in the future.

Accordingly, the traditional Schiphol structure are not designed to manage autonomous operations. The risk being newly developed autonomous processes to act in isolation, creating operational siloes, communication gaps, and ultimately increased inefficiencies.

In stating so, successful automation integration requires careful attention to changing human roles rather than simple technology substitution. Ground handlers possess vital practical influence over implementation success. Their cooperation and trust are essential for

effective automation adoption.

The cross-impact analysis revealed intricate relationships between automation deployment, human role evolution, system integration, and decision-making structures. These interdependencies need to be addressed in unison rather than sequential.

The research comes together in a strategic roadmap structured across three main horizons leading up to the future vision. Horizon 1 (Foundation - Controlled Operations) establishes basic automation support systems while maintaining primarily human control in decision-making and intervention. Key initiatives include implementing smart monitoring systems along with a remote dashboard and beginning stakeholder engagement processes. This phase focuses on building trust and demonstrating automation value while maintaining operational stability. Horizon 2 (Collaboration - Intelligent Operations) introduces AI-assisted decision-making and advanced system integration while preserving human oversight for final decisions. Features include predictive analytics implementation, enhanced intervention functions, and structured human-machine collaboration protocols. This phase emphasises learning and collaboration while building toward more autonomous operations. Horizon 3 (Autonomy - Remote and Resilient Operations) enables automated sequencing and scalable resilience with minimal human intervention for routine operations. Introduces autonomous sequenced processes, autonomous conflict resolution systems, and adaptive workflow management. Human operators transition to exception handling and strategic oversight roles. Future Vision (Autonomous Operations) represents the ultimate goal of fully autonomous operations with human involvement primarily limited to exception handling and strategic decision-making. This phase features comprehensive system integration, advanced AI orchestration, and fully resilient operational capabilities.

Several strategic recommendations and implications are outlined for Schiphol aiming to implement orchestration systems for autonomous operations: Adopting maturity-layered deployment, leveraging AI growth strategically, integrating co-development and change management, and treating the roadmap as an adaptive tool.

Thus, implementing automation in structured phases that build upon previous capabilities rather than attempting comprehensive transformation simultaneously. Especially technological capabilities must be aligned with organisational and human readiness. This approaches mitigates risks while building stakeholder confidence and system reliability.

Then, acknowledging that AI advancements will likely outpace the development of autonomous processes and ecosystem readiness. Ensuring that AI capabilities are leveraged in a way that aligns with the pace of process automation, maintaining balance between rapid technological innovation and the need for operational stability, ecosystem integration, and building stakeholder trust.

To build this trust, the operational stakeholders in system design and implementation processes must be involved. Addressing cultural and procedural changes alongside technological deployment to ensure successful adoption.

Finally, due to the dynamic nature of maturity development, it is crucial to maintain flexibility in implementation timelines and approaches while preserving the strategic direction. Regular reassessment and adjustment capabilities are essential for long-term success.

The overall findings offer valuable insights for understanding the transformation of complex socio-technical systems, extending beyond the context of the scope. These insights are particularly relevant for the total automation vision of all airside operations that also involve multiple stakeholders. The human-centred approach to automation orchestration, with a focus on gradual, trust-building processes, provides a framework that can be applied to the bigger system that requires coordinated technological and organisational change.

Furthermore, the study contributes to the growing body of knowledge on responsible AI implementation. It highlights how technological advancements in AI can be balanced with maintaining human-in-the-loop principles and ensuring organisational sustainability. The emphasis on gradual integration and stakeholder trust-building provides valuable lessons for other sectors undergoing similar transformation challenges, illustrating how to implement innovative technologies while facilitating a supportive and adaptive ecosystem.

Contents

1 Introduction

The aviation industry is undergoing a period of significant transformation, driven by sustainability concerns, regulatory pressures, labour shortages, evolving passenger demand, and operational inefficiencies [62, 23, 81, 49]. Airports must adapt to these challenges by integrating technological advancements to improve efficiency, reduce environmental impact, and maintain competitiveness. Among these, automation and digitalisation are playing an increasingly central role. The benefits of automation include reduced fuel consumption and emissions, less reliance on physical labour, and faster turnaround times. For example, optimised taxiing and automated ground handling can significantly reduce fuel consumption, operational costs, and environmental impact [55]. Additionally, these advancements create a healthier working environment and improve safety conditions for both employees and nearby communities by addressing emissions-related challenges [27]. Furthermore, digitalized ground handling reduces delays and improves airline efficiency [61]. The implementation of AI- and IoT-driven "smart airport" management systems is a key strategy for enhancing operational efficiency, as these technologies enable real-time tracking of processes, passengers, and aircraft, improving security and process coordination [8, 20].

As one of Europe's busiest airports, Schiphol Airport is at the forefront of this transition, aiming for fully autonomous airside operations by 2050 [71, 66]. This ambition is not only driven by innovation and efficiency goals, but also by growing public and regulatory scrutiny due to its environmental footprint, including noise pollution and air quality concerns. Investigative reports such as Zembla's documentary "Ziek van Schiphol" have brought attention to the health risks associated with emissions and noise exposure for both airport workers and surrounding communities, increasing pressure to seek low-emission solutions [81]. At the same time, Schiphol faces structural challenges such as workforce shortages and a continued reliance on physical labour for repetitive and demanding tasks, further reinforcing the need for smart, automated systems that can maintain service quality while reducing the reliance on on-site crew.

At the same time, irregular passenger volumes and operational inefficiencies, have highlighted the need for smarter, automated processes in areas such as baggage handling, fuelling, and aircraft turnaround [16, 72]. In 2024, Dutch airports handled 76 million passengers, 6% fewer than in 2019, demonstrating the shift in travel patterns and an ongoing recovery process [23]. This highlights the need for innovative airport management strategies to handle these fluctuations efficiently, particularly in terms of airport labour force optimisation. As passenger volumes continue to grow while the workforce decreases, a clear challenge emerges, highlighting the need to decrease the reliance on physical labour [13].

To address these challenges, airports worldwide, including Schiphol, are investing in automation and AI-driven operations to reduce emissions, improve efficiency, reduce reliance on physical labour, and reduce costs. However, the transition to automation presents critical challenges in orchestration, coordination, and integration of autonomous systems within airport environments. This research explores how Schiphol can effectively design a control area framework to manage semi-automated and fully autonomous airside operations, ensuring clear responsibility distribution and a smooth transition toward an efficient, intelligent, and sustainable airport ecosystem.

1.1 Schiphol's transition to autonomous airside operations

In 2020, Schiphol officially launched its vision for fully autonomous airside operations by 2050, marking a strategic shift toward innovation, sustainability, and operational resilience [66]. The vision outlines a gradual transformation of airside operations, facilitated by an interconnected system of self-driving vehicles and automated equipment capable of executing tasks (such as baggage handling, aircraft towing, refuelling, and ground servicing) without direct human intervention. This long-term ambition is part of Schiphol's broader response to sector-wide challenges, ranging from environmental pressure to labour shortages, and reflects its aim to become one of the world's most sustainable and technologically advanced airports.

Schiphol's transition to autonomy is already in its early stages and proves its vision through their early implementation efforts. The airport is gradually introducing and exploring automation through trials and pilot projects, focusing on initiatives such as the deployment of autonomous ground vehicles for airside logistics and servicing [28]; sustainable taxiing using a semi-autonomous tow vehicle or “taxibot” to move aircrafts [66]; automated Foreign-Object-Debris (FOD) detection and removal on runways to enhance safety, among other automating initiatives.

These initiatives support Schiphol’s vision for 2050, which is centred on achieving fully autonomous airside operations enabled by an interconnected fleet of autonomous vehicles and associated automated processes [66]. However, as these technologies are being explored, a new challenge arises, not just in their individual implementation, but in how they collectively function within a broader operational framework. The effectiveness of automation depends not only on the capabilities of each system but also on their ability to work together. Without proper orchestration, these advancements risk operating in isolation rather than as a cohesive network, limiting their full potential.

1.2 Problem Definition

As Schiphol and other airports advance their automation technologies, the challenge shifts from simply implementing autonomous and semi-autonomous systems to ensuring their coordination within existing operations. Traditional airport control structures, such as the Air Traffic Control tower and various ground handling coordination teams, were not designed to manage fleets of autonomous vehicles or intelligent systems beyond aircraft movement [30, 43]. Consequently, each new automated vehicle or process tends to operate in isolation, leading to operational silos and communication gaps between systems [65]. Without a dedicated, integrated control mechanism, the potential efficiency gains of automation may be undermined by a lack of coordination. For example, an autonomous baggage tug, an automated docking guidance system, and a robotic FOD detector might all function well on their own, but without central oversight, they could conflict with each other’s activities or with human-operated procedures, causing delays or safety risks [74].

This thesis addresses that problem, focusing on airplane docking operations as a use-case. Airplane docking (the process of guiding an arriving aircraft to the gate and handling its turnaround services) involves many interdependent tasks performed by ground handlers and equipment. As elements of docking operations become automated, such as autonomous Visual Docking Guidance System (VDGS), automated FOD detection, and autonomous connecting of the Ground Power Unit (GPU), there is an urgent need for a control framework to ensure integration and coordination among all these components [20] (Bouyakoub et al., 2017). The core issue is the absence of an integrated orchestration system to manage the interplay between human operators and multiple autonomous systems. This gap not only poses risks of inefficiency and safety incidents, but also hinders Schiphol’s ability to fully leverage automation to achieve its vision of 2050 autonomy.

1.3 Scope and Limitations

Scope: This research focuses on designing an orchestration framework for Schiphol’s airside operations, specifically to orchestrate semi-automated and autonomous processes during aircraft docking and ramp operations until the turnaround servicing phase (see figure 1 for the included steps). The study analyses the existing ecosystem, current operational workflows, and automation plans at Schiphol. It also explores the different roles of involved stakeholders (frontline) in the context of these processes. The research focuses on proposing an adaptive orchestration system that is designed to Schiphol’s specific operations, with an emphasis on system and ecosystem design and guideline development.

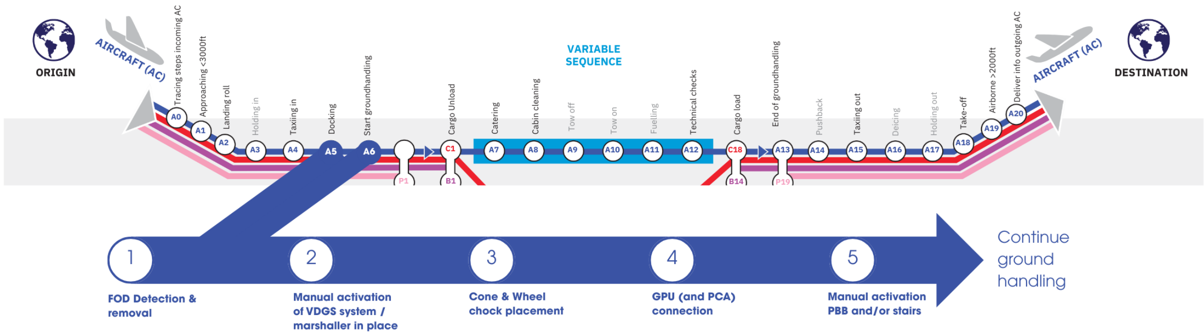

Figure 1: Included steps docking process

Limitations: Several factors limit this research. The research is future forward and includes long-term hypotheses which cannot be tested in a live operational setting. Validation is based on expert feedback rather than real-time performance data. Third, the study is limited to ramp and docking operations, excluding passenger terminal services and en-route air traffic management. Lastly, some assumptions are made about the continued development of autonomous technologies, which may require future refinement as conditions evolve.

1.4 Airport Automation Trends

Airports worldwide are undergoing rapid digital transformation, driven by increasing efficiency demands, safety concerns, and sustainability goals [43, 54]. The shift toward Airport 4.0 involves deploying IoT networks, robotics, and AI-driven decision-making to enhance operational effectiveness [20]. Many airports are investing in autonomous ground vehicles for tasks such as aircraft towing, baggage transport, and apron logistics. Examples include self-driving baggage carts, robotic cleaning systems, and automated fuelling equipment [47].

Accordingly, Schiphol's transition aligns with the automation strategies adopted by other major airports. For example, London Heathrow is implementing autonomous technologies, including driverless pods, touchless passenger processing, and digital crowd-management systems, while also shifting to an all-electric ground vehicle fleet to enhance sustainability [28]. Elsewhere in Europe, airports are integrating automation to reduce emissions and improve operational efficiency and passenger experience. Frankfurt Airport [29] has launched the AI@Fraport initiative, which leverages digital innovation to streamline processes and redesign work environments, aiming to improve outcomes for airlines, passengers, and employees [28]. In addition, Fraport is testing autonomous baggage and cargo tractors as a step toward more efficient, cost-effective ground handling [29]). Similarly, Munich Airport is pursuing digital transformation by introducing "smart" baggage trolleys equipped with interactive tablets to guide passengers and conducting feasibility studies for autonomous apron vehicles [31]. Globally, similar initiatives are underway: Changi Airport in Singapore, in collaboration with Aurrigo International plc, is deploying a fleet of autonomous electric baggage-handling vehicles [39]. These examples reflect a broader trend across the aviation industry to integrate automation for greater sustainability, efficiency, and passenger satisfaction, reinforcing Schiphol's strategic direction toward fully autonomous airside operations by 2050.

A major driver of automation is sustainability, reducing ground handler workload, and preventing staff shortages. The use of autonomous electric ground vehicles helps reduce carbon emissions and air pollution, supporting airports' environmental objectives [66]. Additionally, AI-driven predictive analytics enable proactive maintenance, reducing equipment failures and delays [1]. However, integrating these autonomous processes requires a centralized control system to manage operations holistically [65].

In summary, as airports progress toward greater autonomy, the orchestration of these technologies remains a critical challenge. While individual automation solutions, such as autonomous baggage handling, AI-driven scheduling, and robotic ground support, offer efficiency and sustainability benefits, their effectiveness depends on how well they are integrated into a cohesive operational framework. This underscores the importance of designing human-centred control systems that enable seamless collaboration between autonomous systems and human operators. To develop such a framework, it is essential to first understand existing control centre architectures, human-machine collaboration models, and digital integration strategies. The following chapter reviews the relevant literature on these topics, providing the foundation for addressing the research problem.

2 Literature Review

The transition toward autonomous airport operations presents significant opportunities and challenges in the aviation industry. As major international airports like Schiphol aim to integrate automation across airside activities, critical issues surrounding control mechanisms, human-machine collaboration, and interoperability arise. While technological advancements in artificial intelligence, the Internet of Things (IoT), and robotics have enabled greater automation, the successful orchestration of these systems within an operational framework remains a key research focus.

This literature review explores three fundamental themes essential to understanding the progression toward autonomous airside operations: (1) control areas and their evolution in overseeing autonomous processes, (2) human-machine collaboration in hybrid automation environments, and (3) interoperability challenges among autonomous systems. By analysing existing research on these topics, this review identifies knowledge gaps in the design and implementation of a structured orchestration framework that effectively integrates autonomous operations while maintaining necessary human oversight.

2.1 Control Areas for Autonomous Operations

Modern control centres play an important role in managing intricate operations in aviation and logistics [1]. They serve as centralised areas for real-time monitoring, coordination, and decision-making across various systems. For example, ProRail's railway control center (OCCR) integrates data streams and multi-stakeholder communications to rapidly resolve disruptions, allowing quick decision making during unpredictable events, a capability comparable to crisis management in airport operations [58]. Similarly, airports rely on air traffic control (ATC) towers as central control units to manage the movement of aircraft and ground vehicles on the airfield [4]. Over the years, these aviation control centres have embraced digital tools to improve efficiency. The introduction of electronic flight data systems in ATC towers, for instance, replaced manual flight-progress strips with digital displays, reducing human error and streamlining communication [77]. In parallel, the adoption of IoT-based "smart airport" technologies has enhanced control-centre capabilities by enabling real-time data exchange for passenger flow, baggage tracking, and logistics management. Notably, while IoT networks decentralise some monitoring tasks, they still require a unifying orchestration layer, a top-level control framework, to ensure overall efficiency, safety, and coordination. In other words, even highly connected systems benefit from a centralised layer to synchronise autonomous activities across the airport [21, 45].

As airports transition toward greater automation, their control centres must evolve to manage both automated processes and human-supervised processes. NASA's concept of adjustable autonomy illustrates how human operators and machines can dynamically share control: systems can switch between fully autonomous mode and human-supervised mode as needed [53]. This model is particularly relevant in transitional phases where an airport control room oversees a mix of semi-automated and fully automated systems [78, 57]. During these phases, different operations on the airside (e.g. ground handling, ramp operations, and ATC tasks) demand varying levels of human intervention depending on the situation. A major challenge is orchestrating these mixed automation levels so that autonomous tasks and human-dependent tasks integrate smoothly. One emerging solution is the use of remote control and teleoperation technologies as a safety net. For example, Tener and Lanir [74] emphasise that teleoperation can serve as a fallback mechanism, allowing a human operator to remotely take over an autonomous vehicle if it encounters a problem or unexpected scenario. By establishing remote "control areas" or assistance centres, a small team of human experts could supervise multiple automated operations from afar, intervening only when necessary. This approach balances automation with human oversight, producing efficiency gains while still leveraging human flexibility and problem-solving ability in specific or unexpected cases.

Another key consideration for autonomous airport operations is interoperability and data-sharing between systems. As more autonomous technologies are introduced, they often originate from different manufacturers and use custom communication protocols, which makes integration difficult. Incompatible interfaces and data formats can result in siloed subsystems that cannot directly communicate, undermining the potential efficiency gains of automation. Researchers have proposed solutions like blockchain-based data architectures to facilitate secure, decentralised information exchange across heterogeneous platforms [41]. For instance, implementing a blockchain framework could enable seamless data-sharing between autonomous aircraft tugs, baggage robots, and ramp sensors. Still, such decentralised networks would require an overarching oversight mechanism to coordinate actions and maintain situational awareness. Indeed, even with IoT connectivity and blockchain, a central control authority is needed to ensure all autonomous components work toward common operational goals [45]. Prior studies on connected vehicles echo this point: differing communication standards can lead to conflicts or miscommunication among autonomous units. To mitigate this, middleware solutions have been suggested to bridge the gaps between systems. For example, Agbaje et al. [6] demonstrate how interoperability gateways can act as translators, allowing diverse autonomous vehicles to exchange information and collaborate. Likewise, airport IoT platforms can integrate data from various automated processes, helping to break down data silos and improve system-wide awareness [21]. In summary, interoperability challenges must be addressed through common protocols and integration frameworks so that a future airport control area can effectively manage a network of autonomous technologies as one cohesive system [65].

Finally, as automation expands, control centres face the task of balancing automated operations with human oversight. The industry's progression toward autonomy is a gradual journey rather than an overnight change. Many studies note a gap in guidance for managing the intermediate stages where humans and automation co-exist. Traditionally, research and practice have focused on either fully manual operations or fully autonomous systems, leaving organisations less certain about how to orchestrate hybrid environments where both operate together.To conceptualise the progression toward autonomy, several maturity models have been developed across different industries. These models serve to structure the transition from human-controlled operations to full autonomy. While each framework has its own focus, they generally follow a common trajectory: starting from manual control, introducing assistance and automation, incorporating conditional autonomy, and ending in full autonomous operations.

For instance, ARC Advisory Group outlines a six-level maturity model for autonomous operations, ranging from Level 0 (No Autonomy), where all functions are human-performed, to Level 5 (Full Autonomy), where operations can run independently without human intervention, even under abnormal conditions [38]. Similarly, Yokogawa defines a framework beginning at Manual (Level 0) and progressing to Autonomous Operations (Level 5), emphasizing how orchestration and synchronization across systems become more intelligent and coordinated at higher levels [24].

In the context of mobility and robotics, the Society of Automotive Engineers (SAE) presents a six-level model of driving automation, from Level 0 (No Driving Automation) to Level 5 (Full Driving Automation). This framework is widely cited in autonomous vehicle research and highlights critical thresholds in decision-making transfer from human to machine [75].

Figure 2 (below) provides an overview of these models, illustrating the evolution from human-led to system-led operations. While designed for different domains, these models collectively inform how autonomy can be phased and orchestrated in complex environments.

Advancing through these maturity levels offers the promise of increased efficiency, safety, and operational performance. However, the transition is not purely technical. Each phase introduces distinct human and organisational challenges that require careful alignment of roles, responsibilities, training, and workflows. As seen across models from ARC Advisory Group and Yokogawa, the middle stages, where operations become semi- or conditionally autonomous, introduce coordination challenges. Here, automation takes on various tasks, yet human oversight remains essential for exceptions, adaptation, and intervention [38, 24].

Many airports are about to enter these intermediate stages, where ramp and gate-side processes are partially automated but still depend on human coordination and judgement. Literature across sectors emphasizes that such hybrid settings demand the refinement of orchestration strategies, especially within control areas since these will take on additional functions. Without clear structures for monitoring, communication, and role transition, the efficiency benefits of automation risk being undermined by miscommunication or uncer-

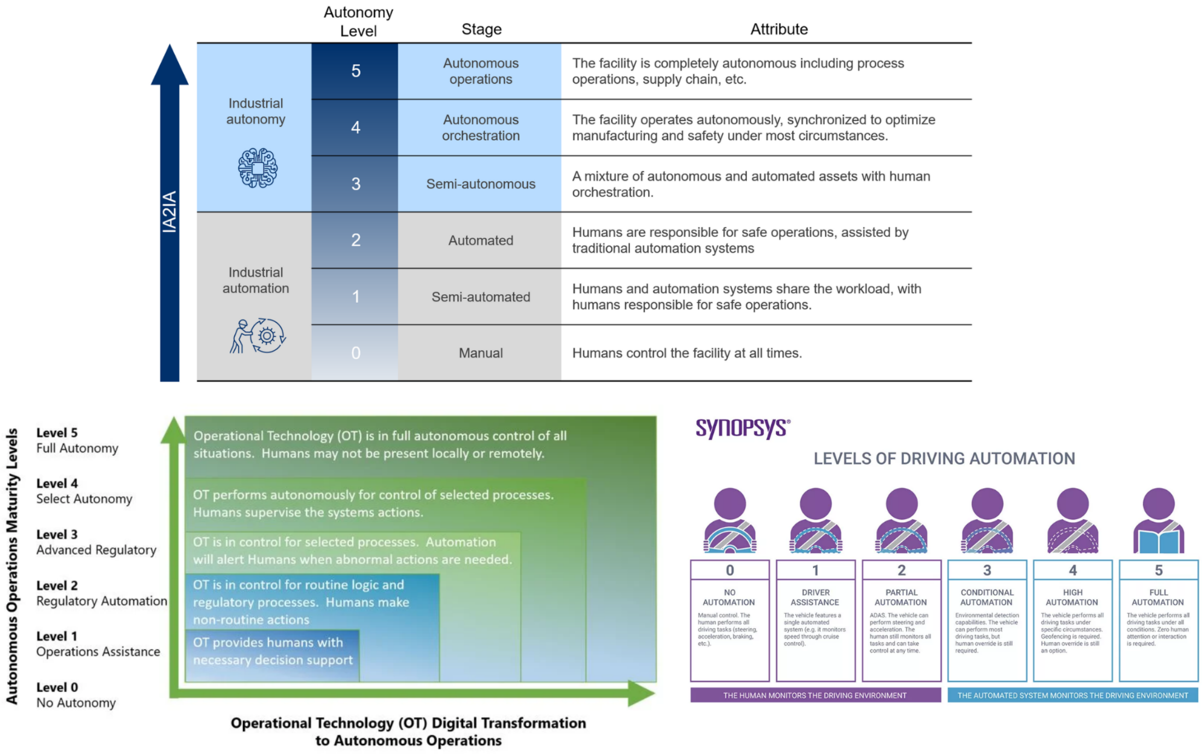

Figure 2: Overview maturity models.

tainty during human-machine handoffs [38, 24]. For airports like Schiphol, this highlights the need to proactively design control architectures and operational cultures that can manage increasing autonomy while maintaining clarity and safety.

Overall, lessons from current control systems in aviation and other industries highlight the importance of centralised coordination, common data standards, and adaptable control frameworks when integrating higher levels of automation [45, 65]. Future airport control areas must strike the right balance between distributed autonomous capabilities and a central orchestrating authority [57]. Achieving this balance ensures that diverse systems – from self-driving baggage carts to AI decision-support tools – can function cohesively while human supervisors maintain situational awareness and the ability to intervene when necessary. These principles form a foundation for guiding the design of new control area architectures capable of handling Schiphol's move toward autonomous operations.

In summary, existing control centres in aviation provide a strong centralised mechanism for oversight, but they must evolve or there must be a new solution to integrate a growing variety of autonomous systems. Key insights from the literature include the need for a top-level orchestration layer to coordinate connected autonomous technologies, and the importance of maintaining human oversight (through adjustable autonomy and teleoperation) during the transition to higher automation levels. Airports will likely require enhanced control room architectures that can manage mixed human–machine operations and ensure different autonomous subsystems work in unison. This leads to Sub-Theme 1: **Architectural Orchestration Baseline** along with Sub-Research Question 1: *"What architectural and orchestration mechanisms control centres can guide the design of a future-proof control area for managing semi-autonomous and autonomous airside operations at Schiphol?"*

2.2 Human–Machine Collaboration and Decision-Making

As aviation operations increasingly integrate automation, the nature of human work is changing. Routine tasks once done manually are now handled by machines or algorithms, pushing human operators and workers into more supervisory and high-level decision-making roles

[3]. Effective collaboration between humans and autonomous systems is therefore essential to maintain safety, efficiency, and adaptability in this new environment. Research in this area examines how to redefine roles and responsibilities, overcome adoption challenges, implement human-centric decision frameworks, and prepare personnel (such as ground handlers) for AI-assisted roles. A central question is how to ensure that growing automation enhances human performance rather than undermines it, preserving necessary human oversight and expertise [9]. In other words, how can human operators and autonomous systems collaborate so that airports gain the benefits of automation without sacrificing safety or efficiency?

2.2.1 Defining roles and responsibilities in hybrid systems

A first step toward effective human–machine collaboration is clearly defining which tasks are handled by automation and which tasks remain under human control. This clarity is critical during the transition from manual to autonomous processes [78]. The maturity models map out how human roles evolve at each stage of automation. In theory, as systems become more autonomous, human operators shift from direct, hands-on control to high-level supervision. In practice, however, organisations often struggle to define these new roles. A recent qualitative study at Amsterdam Airport Schiphol found that practitioners had differing interpretations of automation's role and how humans should remain involved, leading to ambiguity, even internally, around automation implementation [35]. Gómez-Beldarrain et al. [34] further observed that many companies underestimate the complexity of adapting to "automation," assuming humans will naturally step back as machines step forward. Instead, deliberate effort is needed to design socio-technical systems that support new forms of teamwork between humans and AI. Alix et al. [10] argue for a socio-technical approach to adaptive human–machine collaboration, emphasizing that increasing automation requires rethinking workflows, trust, and responsibility allocation in a holistic way (See Figure 3).

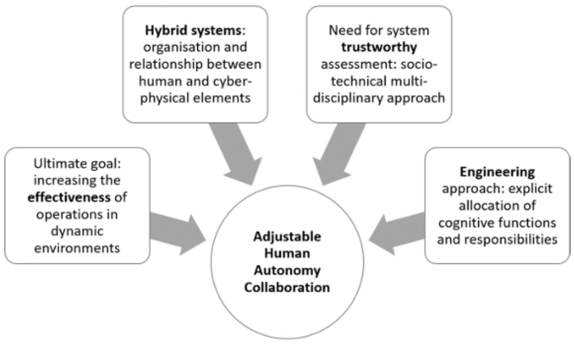

Figure 3: *Adjustable human autonomy collaboration rationale*

In such frameworks, adjustable autonomy is key: the level of automation should be flexible so that control can vary between human and machine as situations demand. This requires that each function in the operation has an explicitly assigned "owner", whether human, AI, or a shared responsibility, and that both parties understand when and how to hand off control. In hybrid control setups, AI can handle routine, data-heavy tasks, but human experts are still relied on for oversight and intervention in high-risk or novel scenarios. Ensuring this collaboration is effective also demands building trust: operators must trust the automated systems to do their part, and the systems must be predictable and transparent enough to earn that trust. For instance, an autonomous tug or gate docking system must behave consistently so that ground controllers know what to expect from it, and it should be able to explain its actions in terms that humans find credible. Overall, redefining human roles in an automated environment means shifting humans into supervisory positions, but with support so they can oversee multiple processes and step in only when necessary. This often requires new skills, such as the ability to monitor remote operations or interpret AI outputs, and a mindset that AI is there to support, not replace, human expertise [40]. Successful human–machine teams use the strengths of both parties: the speed, precision, and data-handling capacity of machines, combined with the judgment, flexibility, and contextual understanding of humans.

2.2.2 Challenges in automation adoption

Introducing advanced automation into airport operations is not just a technical upgrade; it is a socio-technical transformation that can encounter significant organisational resistance. Gómez-Beldarrain et al. [35] report that at Schiphol, efforts to automate airside processes were complicated by misaligned stakeholder priorities and governance issues, leading to delays and inefficiencies. Different departments and personnel often have differing expectations and fears about automation, making it difficult to implement changes smoothly. A common issue is the unclear definition of human roles when automation is introduced. If it's not specified how an operator's job will shift once a task is automated, the operator may become disengaged or, conversely, may interfere unnecessarily. Ideally, automation should shift human responsibilities away from tedious manual control and toward higher-level supervision and decision-making. However, without clear role boundaries and proper training, staff may either over-rely on automation (becoming complacent) or under-use it (distrusting the system and reverting to manual methods), both of which hurt overall performance. Moreover, employee attitudes toward automation vary widely. Some workers embrace new technology as a tool that makes their job easier, while others fear that automation will make their skills obsolete.

Gödöllei and Beck [33] note that fear of job displacement can lead to active resistance or even sabotage of automation initiatives. These human factors are as crucial to address as the technical ones. Effective change management strategies, including transparent communication, involvement of end users in design, and assurances about job security or upskilling opportunities are vital to encourage a positive view of automation. Another adoption challenge is ensuring the workforce has the skills to work with AI systems. The different maturity models of increasing autonomy implicitly calls for incremental upskilling: as operations progress from one level of automation to the next, employees should be gradually trained to handle their changing duties. A phased approach allows the organisation to learn and adapt at each stage, rather than attempting a disruptive, all-at-once change. Practical steps include comprehensive training programs. For example, Hajam and John [40] advocate for structured training that teaches personnel how to oversee AI systems, interpret automated alerts, and intervene appropriately. Such programs build AI literacy and confidence, helping staff to see automation as an aid rather than a threat. Building trust in automation is also critical; this can be achieved by demonstrating the reliability of new systems and involving operators in testing and feedback loops. Overall, research suggests that successful automation adoption requires a holistic approach: clear role redefinition, proactive workforce training and engagement, and strong governance to align all stakeholders [35, 33]. By tackling these organisational and human factors, airports can transition to automated operations more smoothly, ensuring that efficiency gains do not come at the cost of human disempowerment or safety.

2.2.3 AI-assisted decision-making in control rooms

One area where human–machine collaboration is especially visible is in operational decision-making. Traditional decision processes in airport control (e.g. air traffic control) have been enhanced by automation for decades, as seen with the earlier move from paper flight progress strips to computerised systems. Today, artificial intelligence is enabling a new generation of decision-support tools (DSTs) that help human controllers manage traffic, resources, and disruptions [5, 77]. For these AI-driven systems to be effective, researchers emphasise that they must behave like "good teammates" to the human operator.

According to principles outlined by the Federal Aviation Administration [5], this involves several key design attributes: **predictability, observability, and directability** of the AI's actions.

First, mutual **predictability** is important, the automation should act in consistent ways that a human controller can anticipate. In an ATC context, for example, controllers are accustomed to inferring each other's intentions by observing actions on their consoles; similarly,

an AI tool should present its suggestions or take actions in a predictable manner so the human isn't surprised by its behaviour. This consistency builds trust, as controllers learn what the AI will do in various scenarios and can plan accordingly.
Second, system **observability** (transparency) is needed, the AI should make its reasoning process and status visible to the user. Rather than being a "black box," the system might display information about how it came to a recommendation, what alternatives it considered, and how confident it is in the result. Such transparency allows the human operator to understand the AI's suggestions and to spot if something might be amiss. Mietkiewicz et al. note that a good DST filters out noise and highlights relevant data, reducing the controller's cognitive workload rather than adding to it [51].
Finally, **directability** is crucial, the human operator must always retain ultimate control and be able to redirect or override the automation when necessary. The system should be adaptable to human input; for instance, if a controller disagrees with an AI's course of action, they should be able to modify it or instruct the system to pursue a different solution. In practice, this means AI tools must submit to human judgment in unusual or high-stakes situations and support the controller's authority, rather than displacing it. Industry guidance stresses that automation in the tower should act as an assistant, not an independent decision-maker that unilaterally controls operations.
By incorporating predictability, transparency, and directability, AI decision-support systems can significantly enhance decision-making in control rooms. They can accelerate routine decisions, flag potential issues that a human might overlook, and provide data-driven recommendations, all while allowing human experts to apply their experience and intuition to the final call. When designed with a human-centred approach, these tools increase efficiency and safety, yet they keep human expertise in the loop for oversight. The evolution we are seeing and anticipating in ATC and operations control due to AI assistance is a template for how other airport roles will change. Just as tower controllers will team with AI in decision-making, ground operation managers and dispatchers are starting to work with AI-based systems that help allocate resources and predict bottlenecks [64].

2.2.4 Evolving roles of ground handlers and operators

Nowhere is the transformation of work more apparent than on the airport apron, where automation is changing how ground handling and servicing tasks are performed. Traditionally, ground handlers engage in manual labour, guiding aircrafts into gates, loading baggage, refuelling vehicles, etc. With the advent of robotics and autonomous equipment, many of these physical tasks are on the brink of being automated, and the role of ground staff is shifting from hands-on labour to system oversight and technical management. In a smart airport environment, IoT-enabled systems can track baggage in real time, coordinate vehicle fleets, and even automate baggage transportation and sorting. These innovations greatly improve efficiency and reduce delays, but they also demand that the workforce adapt to monitoring and managing automated processes rather than executing them directly. For example, an autonomous baggage cart may handle the actual transport of luggage, but a human supervisor must ensure the system is functioning properly and intervene if an issue arises. Studies of airports at the forefront of automation illustrate this changing skill set. KLM Ground Services is retraining its staff in response to the introduction of autonomous ground vehicles and AI tools. Instead of focusing purely on manual skills, ground handlers are learning digital infrastructure management, system oversight, and AI-assisted coordination capabilities [42]. A key challenge is providing sufficient training so that workers can transition from physically handling equipment to overseeing and troubleshooting the robots and systems that now handle the equipment. This requires not only technical know-how (understanding how the autonomous systems work and how to reset or repair them) but also changes in mindset and procedures (e.g. knowing how to efficiently manage multiple autonomous units simultaneously). Research by Kovynyov and Mikut [48] highlights that ground operations are becoming increasingly data-driven, for instance, big-data analytics are used to forecast passenger flows and dynamically allocate staff. Ground handling agents are thus leveraging digital tools to improve service reliability and operational efficiency, which means their roles include more decision-making based on data insights. While the nature of the work changes, there is evidence that automation does not simply eliminate jobs, but rather redistributes tasks and creates new types of work. As Acemoglu and Restrepo [3] argue, automation often shifts human labour toward tasks requiring higher technical and cognitive skills, instead of making humans redundant. In the airside context, this means a baggage handler might evolve into a fleet supervisor for robotic vehicles, or a gate agent might become a remote monitoring specialist for an autonomous boarding bridge system. The core judgment and problem-solving abilities of employees remain in demand, but now applied alongside overseeing automated systems. To ensure this transition is successful, airports need to invest in reskilling and upskilling programs. Workers should be provided with training in IT and data analytics, robotics basics, and safety management for automated operations [40]. With these skills, staff can collaborate with AI systems by focusing on exceptions, maintenance, and continuous improvement, rather than performing repetitive manual tasks. Thus, rather than replacing human workers, automation in ground handling is reshaping their roles, placing greater emphasis on cognitive and supervisory duties over physical labour.

One tangible outcome of these changes is that human operators are now able to manage multiple processes simultaneously (with the help of centralised dashboards and control centres), something not feasible in the purely manual era. This shift in workforce dynamics extends beyond ground handling. Similar changes are happening with maintenance crews (who use drones and AI for inspections) and security staff (who use AI-driven analytics for surveillance). All of these cases underscore a broader need for organised frameworks that facilitate smooth interaction between automation and human oversight across all airport operations. In every case, clearly defining how humans and machines collaborate, *who is responsible for what, and how information flows between them*, is crucial to ensuring that increasing automation leads to safer and more efficient operations, rather than confusion or new types of error.

In summary, the literature on human–machine collaboration in aviation consistently highlights that technology alone cannot deliver the desired benefits; it must be accompanied by human-centric design and management. Whether it's controllers working with AI decision-support tools or ground crews supervising robots, keeping humans "in the loop" in a well-defined way is critical. Successful automation initiatives treat operators as partners to the technology, adapting organisational structures, roles, and training so that humans remain empowered to guide, correct, and augment autonomous systems.

As airports introduce more automation, they must redefine human roles and workflows to ensure effective collaboration. Clearly assigned responsibilities, mutual trust, and transparency between humans and AI are essential for maintaining safety and efficiency. Organisations need to proactively manage cultural resistance by involving stakeholders and providing adequate training, ensuring staff embrace new tools rather than fear them. Instead of being displaced, human operators should transition into supervisory and decision-making roles, supported by AI to enhance overall performance.

These findings lead to Sub-Theme 2: **Human-Machine Role Transition and Oversight** along with Sub-Research Question 2: "*How should roles and responsibilities processes evolve to enable safe and efficient collaboration between humans and autonomous systems across automation maturity levels in airport operations?*"

2.3 Autonomous Systems and Interoperability

Airports are rapidly exploring autonomous technologies on the airside, the ramps, gates, taxiways, and service areas, in a bid to improve sustainability, health, and efficiency. Drivers for this shift include the need to reduce emissions, optimise resource use, and increase operational throughput [27]. By automating various ground processes, airports aim to minimise delays, cut fuel consumption (for example, by reducing waiting times), and lower the risk of human error. Strategic use of automation can decrease the waiting time of ground vehicles and ensure more predictable turnaround operations, contributing to sustainability goals. Recent advancements in artificial intelligence, robotics, and connectivity (IoT) have made it feasible to develop fully or semi-autonomous systems for many airside tasks. While some automation has existed for years (such as automated baggage sorters or fuelling timers), the new generation of autonomous systems goes further by leveraging AI for real-time decision-making and coordination with minimal human input. Importantly, these systems cannot operate in isolation, they must collaborate and integrate with the airport's existing infrastructure and with each other. A sturdy digital infrastructure (including reliable wireless networks, sensors, and data platforms) is required to integrate autonomous units into airport operations.

Jaradat et al. [45] emphasise that without strong integration capabilities, even the smartest autonomous robot will struggle to add value in the complex airport environment. Indeed, airports like Schiphol are finding that to gain the benefits of autonomy, they must upgrade legacy systems, establish connectivity standards, and ensure that data flows freely between old and new platforms. The potential benefits of autonomous airside operations are significant, studies suggest reductions in aircraft turnaround time, enhanced safety by removing humans from dangerous tasks, and lower labour and operational costs. However, implementing these technologies at scale comes with challenges. Key hurdles include ensuring interoperability among a diverse set of automated systems, aligning a multitude of stakeholders (airlines, ground handling companies, regulators) to embrace new processes, and restructuring existing control and oversight mechanisms to accommodate autonomous workflows. In essence, adding autonomous vehicles and equipment into the mix necessitates a new approach to coordination. As Gómez-Beldarrain et al. points out, it's not enough to deploy smart machines; airports must also adapt their operations management so that these machines can be effectively coordinated within the broader airport system [35, 26].

This section provides an overview of the emerging autonomous technologies on the airside and then examines the coordination and interoperability challenges that arise, along with best practices for integrating these technologies into existing airport operations. The guiding premise is that while individual automation projects can produce improvements, the real gains will come only when they function as part of a unified, orchestrated system [57].

2.3.1 Autonomous airside technologies in development

A variety of autonomous and semi-autonomous solutions are currently being tested or implemented on the airside, targeting different aspects of ramp and airfield operations. Notable examples include:

- *Foreign Object Debris (FOD) detection robots:* FOD on runways and taxiways (e.g. stray hardware or trash) poses serious safety risks. Autonomous FOD-detection units are in development to scan runways for FOD.

- *Autonomous passenger transfer vehicles:* Several airports (including Schiphol) have piloted driverless apron buses to transport passengers or staff on the airport. These electric autonomous buses use AI navigation systems and vehicle-to-infrastructure communication to move people between terminal gates and remote aircraft stands. By optimising routes and adjusting to real-time ground traffic, they aim to reduce transfer times and congestion, as well as lower emissions compared to traditional diesel buses [7].

- *Automatic Passenger Boarding Bridges (APBB):* Automating the docking of passenger boarding bridges to aircraft can save time during boarding and deplaning. Explorations are underway with APBB systems that use sensors and AI to align the bridge with the aircraft door without direct human operation. In practice, these systems can perform the docking procedure autonomously; however, human oversight is still required to ensure proper alignment and to intervene in case of malfunctions or unexpected situations [69]. Typically, a gate controller would observe via cameras as the bridge docks itself and take over if needed.

- *Visual Docking Guidance Systems (VDGS):* VDGS are not new, many airports use them to guide pilots to the correct stop position at the gate, but traditionally they have been semi-automated and activated by ground staff. Advances are being made to fully automate VDGS so that they detect an approaching aircraft and activate on their own.

Automating their operation (in place of a human marshaller with wands or manual activation) reduces reliance on human input and can avoid waiting times. At present, VDGS still require a person to initiate the system when an aircraft arrives [56].

- *Autonomous ground power unit (GPU) connection:* After an aircraft parks at the gate, ground staff typically connect a GPU to supply electrical power while the engines are off. Efforts are underway to automate this process using robotic arms or self-driving GPU units. The goal is for a robot to maneuver to the aircraft and plug in the power cable without human assistance, speeding up the provisioning of power and reducing workforce requirements [68].
- *Automated baggage handling:* Automated baggage handling is being explored and piloted (pilots only in baggage halls, not yet apron). The next step is deploying autonomous guided vehicles to transport baggage and cargo between the terminal, the tarmac, and aircraft. Pilot programs involve self-driving luggage carts that follow optimised routes to aircraft cargo holds. These systems aim to reduce misloaded or delayed baggage, reduce workload, and improve overall logistics efficiency [48, 67].

Together, these technologies represent a wave of innovation transforming airside operations. Each one targets a specific workflow (safety inspections, passenger transfer, aircraft servicing, etc.) and promises to make it faster, safer, or reduce workload. However, if deployed independently, they risk becoming isolated solutions: one team or vendor might operate the autonomous buses, another the baggage robots, and so on, with no mechanism to coordinate their activities. The literature warns that without a higher-level orchestration strategy, these innovations could remain improvements in their silos rather than contributing to a seamlessly efficient airport system [65]. In fact, the introduction of multiple autonomous systems has begun to reveal pain points in coordination, as discussed next.

2.3.2 Coordination challenges and system fragmentation

While individual autonomous projects have shown promise, a major challenge is getting them to work in unison. Currently, many automated airside systems operate independently and were not designed to communicate with one another. For example, an autonomous baggage vehicle might be developed by one company and an autonomous aircraft tug by another, each with its own proprietary software and data format. Without integration, each system requires separate monitoring and control, preventing the airport from a seamless, integrated operation. Jaradat et al. [45] note that the lack of standardised communication protocols and data formats among autonomous platforms leads to fragmentation and inefficiencies. In practice, this means one control room screen (or one dedicated team) might be needed to oversee the baggage robots, and another for the tugs, because the two systems don't share information. The result is that the overall performance benefit of automation is reduced – the airport doesn't get the full compounding effect of these technologies since each piece is still effectively in its own lane. Indeed, siloed operations can cause suboptimal decisions or even conflicts on the apron. For instance, an autonomous baggage cart that is unaware of an autonomous fuel truck's movements could end up in a near-collision or delay simply due to lack of coordination. A human supervisor can only do so much to manually de-conflict such scenarios; ideally, the systems themselves would exchange status data and negotiate right-of-way or timing automatically.

A related issue is that existing airport control structures were not designed with these autonomous systems in mind. Air traffic control towers and apron management teams primarily focus on aircraft movements and gate scheduling, not managing robots on the ramp. As autonomous ground vehicles proliferate, there is a gap in oversight: ATC is still controlling aircraft, but no one is centrally coordinating the ground robots in a unified way. Studies confirm that current airport control centres and procedures do not account for autonomous ground operations, resulting in ad hoc management or neglect of those systems. While automation on the airside is increasing, airports have yet to establish a dedicated "orchestration centre" to holistically manage all these new automated processes. In Schiphol's case, for example, there is currently no single control room that has a real-time picture of both manned activities and autonomous elements on the apron [44]. This organisational shortcoming means each automated system might be running on its own schedule, potentially interfering with others or at least not optimising simultaneously.

Another significant challenge is technical interoperability, ensuring that all autonomous systems can communicate and collaborate effectively. As mentioned, different parties may use different protocols; without translation layers or standards, a baggage robot doesn't "know" the position or intentions of a fueling robot, for instance. San Emeterio de la Parte et al. [65] emphasise that many IoT-based automation efforts suffer from data silos where information isn't shared between platforms. In their study, even systems intended to be part of a connected airport ecosystem could not fully cooperate due to differing governance and lack of real-time data exchange. They conclude that such silos hold back true collaboration between systems, underlining the need for standardised communication protocols across all autonomous equipment. Without addressing this, an airport might have the latest high-tech vehicles and robots, but still operate in a fragmented way similar to having multiple "mini" control centres for each automation domain. In dynamic and safety-critical environments like an airport ramp, these coordination issues are more than just inefficiencies, they can translate into incidents or operational delays. For example, imagine an autonomous pushback tug starting to move an aircraft while an autonomous apron bus is passing right behind it, because neither machine was aware of the other's presence; such a scenario must be prevented through either centralised control or direct vehicle-to-vehicle communication. Real-time coordination is essential, and it demands both technological integration and an organisational mechanism to manage it.

Given these challenges, experts and industry groups have pointed to best practices and strategies for integrating autonomous technologies into airport operations effectively. A recurring theme is the need for a unified control framework or "central orchestrator." In essence, airports may need to establish a new kind of control centre (or adapt existing ones) that oversees all autonomous activities on the ground. This would function like an expanded Operations Control centre that interfaces with each autonomous system, aggregates their data streams, and can send commands or schedule tasks to avoid conflicts. (Earlier in this review, we noted that even decentralised autonomous networks still benefit from a top layer of oversight.) Implementing a central coordination system could ensure, for example, that an autonomous baggage train waits because it knows an autonomous fuel truck is crossing its path, or that multiple robots servicing the same aircraft are sequenced properly to avoid interference.

In technical terms, standardization is a key best practice: adopting common communication protocols or middleware solutions that translate between systems can greatly enhance interoperability. Agbaje et al. [6] demonstrated a middleware gateway approach in the context of connected vehicles that could be applied here, essentially, all autonomous units would connect through a central gateway that handles data format conversion and conflict resolution. Such a solution can be introduced incrementally as new systems come online. Another promising avenue is using IoT and cloud platforms as integration layers: Bouyakoub et al. [21] showed how an IoT-based airport management system can integrate data from various processes in real time. Building on that idea, airports can require that any new autonomous equipment be compatible with an existing data platform or API to facilitate plug-and-play integration. Secure data-sharing is also crucial; Hardjono et al. [41] propose blockchain technology as one method to ensure trust and traceability in the communications between autonomous agents. By securing communications and standardizing their format, multiple parties (airport authorities, airlines, ground handlers, etc.) can confidently plug into the same network of autonomous systems without fear of data tampering or loss.

In addition to technology integration, new operational protocols need to be developed. Best practices include creating rules of the road (right-of-way rules) for autonomous vehicles on the apron (much like traffic rules for driverless cars on public roads), establishing procedures and prioritisation if a unit loses connectivity or malfunctions, and defining clear escalation paths to human controllers when needed. All of these should be validated through simulations and field tests to refine the interactions between systems. Another critical best practice is maintaining a human-in-the-loop for oversight, especially during the transition period. A human supervisor or team should have a comprehensive view of the autonomous operations and the authority to intervene. This ensures that if the autonomous network encounters a scenario it can't resolve, an unexpected obstacle, a sensor failure, a novel situational conflict, humans can quickly step in to re-coordinate or take control. Studies strongly caution that ignoring the human element in orchestration could obstruct adoption or lead to unmanageable systems [57]. Therefore, any orchestration framework must incorporate interfaces for human monitoring and control. In practice, airports might designate a kind of "ramp automation manager" in the control centre who oversees the automated systems.

In summary, to integrate autonomous airside technologies effectively, both architectural and organisational solutions are needed. A digital platform or interoperability layer can enable real-time data-sharing and coordinated control, while industry-wide standardisation can reduce fragmentation. Additionally, a dedicated coordination function and clear operational procedures will ensure these technologies work together rather than in isolation. Airports must move beyond fragmented automation toward a unified orchestration strategy. Lessons from warehouses and smart factories show that orchestration software can successfully manage diverse autonomous systems. Similarly, airports must treat automation as an integrated part of their operational framework, not as isolated tools.

Despite advancements, there is no established framework for orchestrating a multi-system autonomous airside operation under human supervision. Without a unified control strategy, the benefits of automation will be limited by inefficiencies and lack of coordination. Establishing standardised communication, centralised control, and human oversight is key to unlocking its full potential. This leads to Sub-Theme 3: **System-Level Interoperability and Scalable Orchestration** along with Sub-Research Question 3: "*What technical and organisational strategies are required to orchestrate diverse autonomous systems into one cohesive, real-time operation under unified human oversight at Schiphol?*"

2.4 Knowledge Gap

Despite extensive developments and pilot programs for individual autonomous processes on Schiphol, there remains a clear gap in both knowledge and practice when it comes to orchestrating all these systems within a single, unified control framework. Today's airport control infrastructure, essentially the ATC tower and various operations centres, is designed for managing aircraft movements and certain logistics, but not the coordination of fleets of autonomous ground vehicles and service equipment. As a result, new automated systems risk operating in silos, functioning independently without integration into a cohesive system. This siloed deployment leads to inefficiencies and even potential safety risks in a busy airport environment [65]. Essentially, airports like Schiphol currently lack a central approach to supervise and coordinate multiple autonomous activities simultaneously.

What is missing is a structured orchestration strategy or framework that enables seamless interaction among the diverse autonomous technologies on the airside. Such a framework would provide the architecture for real-time data exchange and joint decision-making between systems, ensuring interoperability as more processes become automated. It would act as the centre that ties together autonomous tugs, buses, baggage systems, and other processes, scheduling and harmonising their actions within the overall operation. Important to note, any effective orchestration model must be human-centred, it needs to include roles for human operators and supervisors in the loop. The literature suggests that ignoring the human element could obstruct the adoption of automation or result in systems that are impractical to manage in the complex, unpredictable airport context [34]. In other words, the goal is not to remove every human aspect, but to integrate automation in a way that human oversight and input are still present where needed. Furthermore, the transition from today's manual operations to a fully autonomous future should be handled in phases, as indicated by the different maturity models [38, 75, 24]. However, what is lacking in research is concrete guidance on how to design and implement an orchestration framework that embodies these principles, one that balances automated control with human oversight, and that adapts dynamically to the fast-changing conditions of an airport ramp.

This is the essence of the knowledge gap: we have numerous pieces of the automation puzzle (each individual autonomous system), but we do not yet have an established framework for integrating them under one cohesive control system. Academic and industry literature to date has explored components of this problem (communication standards, human factors, case studies of single automated systems), but an overarching solution framework remains hard to pin down. Addressing this gap is crucial for airports to move from isolated automation projects to an integrated, centrally managed autonomous airside operation.

To bridge this gap, this research explores three themes:

1. Architectural Orchestration Baseline: *What architectural and orchestration mechanisms control centres can guide the design of a future-proof control area for managing semi-autonomous and autonomous airside operations at Schiphol?*

2. Human-Machine Role Transition and Oversight: *How should roles and responsibilities processes evolve to enable safe and efficient collaboration between humans and autonomous systems across automation maturity levels in airport operations*?

3. System-Level Interoperability and Scalable Orchestration: *What technical and organisational strategies are required to orchestrate diverse autonomous systems into one cohesive, real-time operation under unified human oversight at Schiphol?*

These themes directly contribute to the main research question:

"*How can we design an adaptive orchestration model and maturity-aligned transition plan to facilitate the implementation of autonomous airside operations at Schiphol Airport, while ensuring stakeholder integration, scalable automation, and human–machine collaboration across multiple levels of autonomy?*"

Answering this question is vital to unlocking the full potential of airside automation. Without a guiding architecture, airports risk deploying high-tech autonomous systems that deliver only fragmented improvements. With a well-designed orchestration approach, however, the various autonomous innovations can be transformed into a harmonised system that significantly enhances efficiency, safety, and reliability. This research aims to contribute toward filling that gap by proposing and evaluating a model for such an orchestration framework, thereby helping Schiphol realise the promised benefits of autonomous operations in a coherent and controlled manner.

3 Context at Schiphol

3.1 Introduction to the Current Operational Landscape

Building on the literature review, which outlined the orchestration challenges surrounding automation, human-machine collaboration, and interoperability, this chapter grounds that theory in the actual operations of Schiphol Airport. Field research was conducted through walking along with stakeholders during their work sessions, live observations, and informal conversations with key stakeholders such as ground handlers, ground handler coordinators, and control room operators. These methods provided a first-hand understanding of the airplane docking/arrival process, from FOD detection to the start of turnaround services.

While the literature provides structural insights, the complexities of day-to-day coordination, communication, and task distribution, are best understood through real-world context. The research exposed how operations rely heavily on human routines, verbal exchanges, and workarounds when systems fall short. This chapter outlines the current orchestration, identifies its limitations, and lays the groundwork for framing future opportunities.

This chapter presents a comprehensive mapping of the current docking process, illustrating how coordination currently works across roles, tools, and systems. Through detailed walkthroughs, stakeholder analysis, and visual representations, this chapter demonstrates where processes rely on implicit knowledge, manual decision-making, or non-integrated technologies.

In doing so, the chapter acts as a bridge between theory and practice, showing what orchestration and communication currently is, and where the pain points lie as a foundation for innovating and iterating on this current landscape. It provides the necessary contextual understanding to inform the next phases of this research: identifying design opportunities, generating orchestration strategies, and ultimately proposing ideas that are grounded in operational reality.

3.2 Mapping the Current Docking Process

This section maps out the current airplane docking process at Schiphol in order to make the interaction between stakeholders, tasks, tools, and timing visible. By visualising the intended sequence of operations and the actors involved, this chapter provides the foundation for understanding both how coordination currently takes place. The swimlane diagram and accompanying narrative provide a step-by-step walkthrough of the docking workflow from a broad perspective, which is then followed by a closer look at the specific roles and responsibilities of key stakeholders.

3.2.1 Process walkthrough

The swimlane diagram shown in figure 4 represents the sequence of events and stakeholder interactions that make up the current airplane docking process at Schiphol. It follows the process from the initial communication of flight data to the completion of docking and handover to service teams. The figure indicates which parts of the process are within my scope by shading the other sections, the original, unshaded version can be found in appendix **??**.

The process begins when the airline communicates flight and passenger data to the 'regiecentrum', using CISS: the central information system used by various departments including ATC, bus regie, gate planning, and ground handler coordination (see figure 5 for a CISS overview). Based on this data, the gate planning team assigns a gate to the incoming flight. In doing so, gate planners aim to optimise operational efficiency by planning flights in proximity to the ground handling company responsible. This minimises the need to transport ground handling teams across the airport, which saves time and reduces logistical strain.

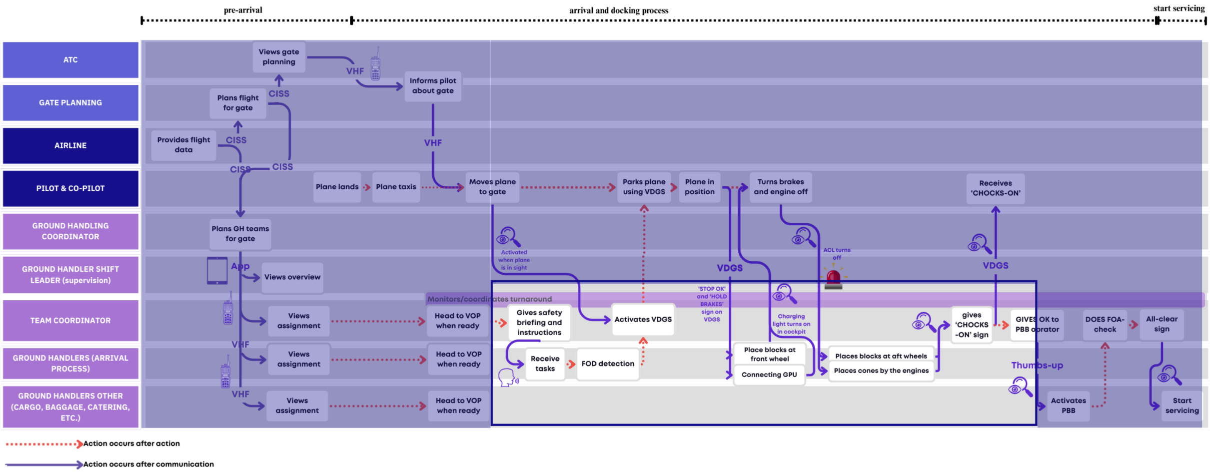

Figure 4: Airplane arrival process.

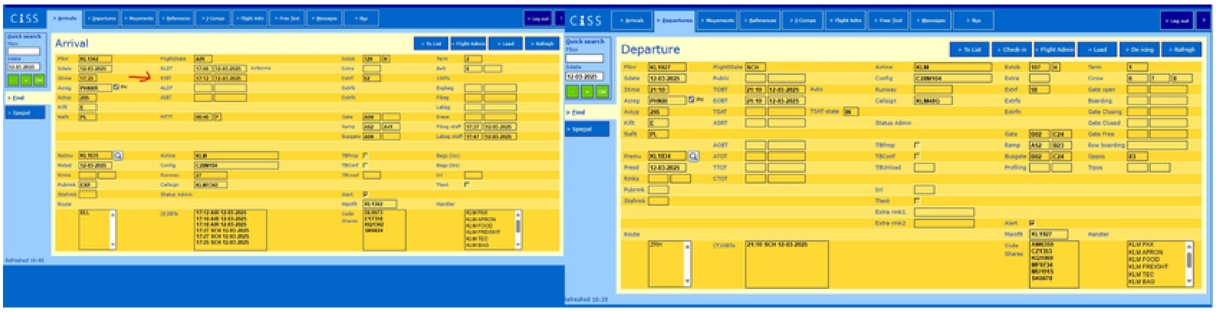

Figure 5: CISS.

Once the gate is assigned, the ground handling coordinator (regisseur) or planner begins assigning ground handling teams to each flight. This planning is performed digitally and submitted through a central platform that allows all ground handling staff to access their assignments. When on the airside, ground handlers can see their tasks using their handhelds. In addition to digital tools, some ground handling companies also post the daily planning on a whiteboard in the operations office for visibility.

The planning is then shared with the shift leaders, who receive the information via a specific application (this varies per ground handling organisation). The application provides a clear overview of flight schedules, assigned teams, and operational details for the shift. Ground handling teams access their assignments either by viewing the physical board or through updates relayed via the portophone (VHF radio). After completing a prior task, the team proceeds to their next assigned gate.

Before docking can begin, the arrival ground handling team must be present at the stand. Their first responsibility is to conduct a FOD check to ensure the stand area is clear and safe. Only once this is completed can the VDGS be activated by a trained ground handler or the team coordinator (TC). This system guides the pilot into the correct position on the apron. The VDGS is usually only turned on once the ground handler has the airplane in sight.

After the airplane has parked and the anti-collision lights are turned off, the wheel blocks are placed and the GPU is connected. Once connected, the pilot can disengage the brakes and shut off the engines and/or APU. The team then places safety cones around the aircraft and activates the PBB.

With these steps completed, the aircraft is considered docked and secure. The arrival ground handling team has fulfilled its tasks, and the airplane is handed over to other service teams such as catering, baggage handling, fueling, and cleaning. These teams begin their work only after receiving clearance from the team coordinator, who remains present to coordinate the entire turnaround.

Throughout this process, coordination is heavily reliant on human timing, verbal updates, and system workarounds. Disruptions at any point, such as last-minute gate changes, delayed equipment, or staffing issues, can quickly snowball, impacting the overall turnaround efficiency.

3.2.2 Stakeholder and roles

The aircraft docking process involves a network of stakeholders, each with different responsibilities, tools, and points of interaction. These roles are layered across strategic planning, real-time coordination, and physical execution. Understanding who is involved, how they communicate, and where their responsibilities and values intersect is crucial to identifying where breakdowns in orchestration may occur.

General Stakeholder Landscape

Figure 6 provides a stakeholder map of the broader airside orchestration landscape. It positions key actors along the axes of power and interest. While actors like Schiphol, airlines, and technology developers hold high-level influence over investment decisions and system design, the daily success of any orchestration strategy is highly dependent on those working directly on the apron.

A key insight from this mapping is the underestimated role of ground handlers. Though they are often seen as having low power in organizational hierarchies, their practical influence is crucial. The successful implementation of any new orchestration or automation technology depends on their willingness to collaborate, trust new tools, and adapt workflows. Without their cooperation, even the most advanced systems are at risk of becoming underused or bypassed in practice. This research highlights the need to reframe their role as central to the transition, especially as operations move toward semi- and fully autonomous workflows.

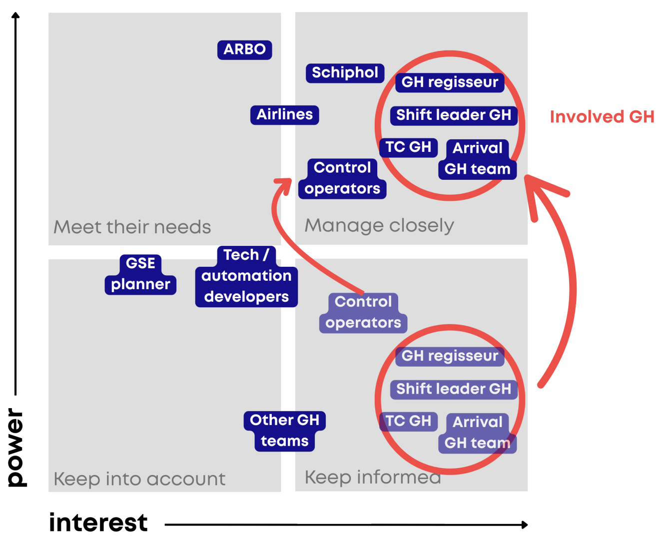

Figure 6: power-interest matrix.

Ground Handling Roles in the Docking Process

Figure 7 provides an overview of the core ground handling roles involved in the docking process.

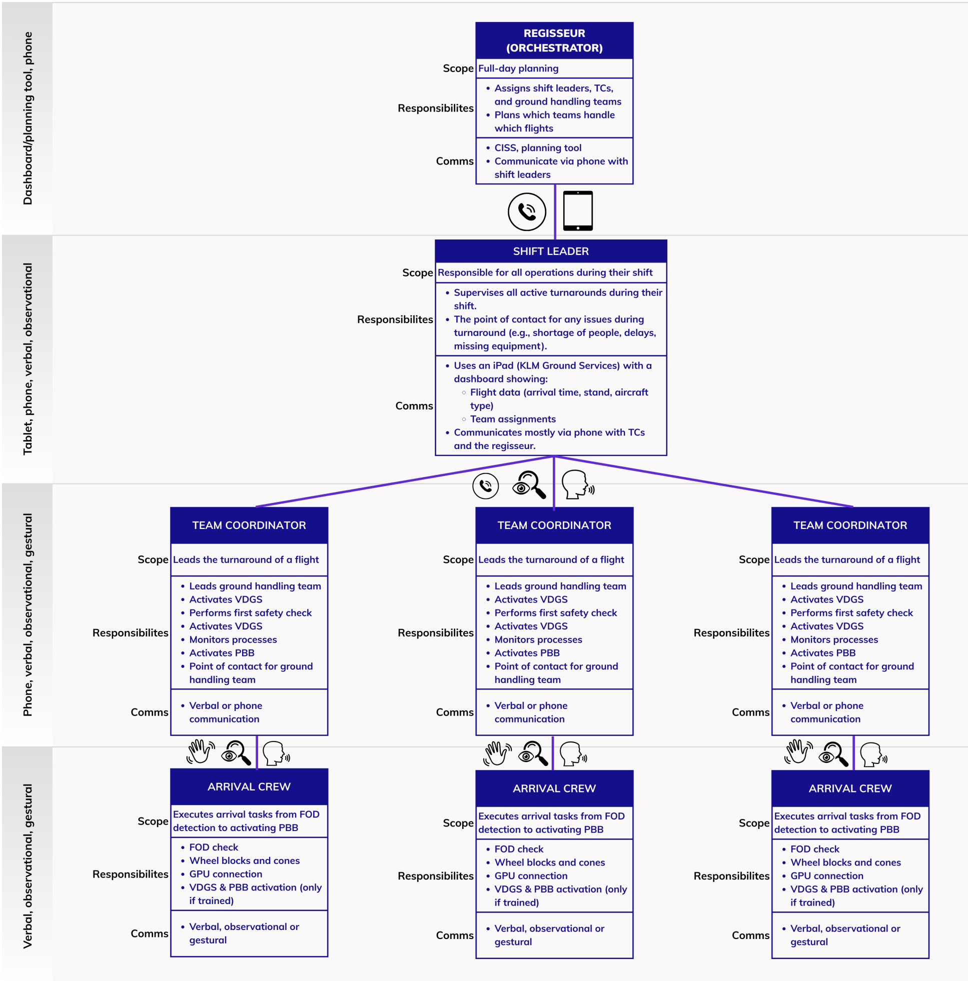

Figure 7: Ground handler dynamics.

Regisseur (Ground Handling Coordinator):
The regisseur is responsible for the strategic daily planning of all ground handling operations. They assign shift leaders, team coordinators, and ground handling teams per flight. This role is typically carried out from the back office and involves working with internal planning systems that connect to CISS. While not directly involved in the execution phase, the regisseur's decisions shape the groundwork for operational efficiency and resourcing throughout the day.

Shift Leader:
The shift leader manages all aircraft turnarounds during a particular shift, often using an application on a tablet (e.g., KLM's shift overview app) to access live flight data, team assignments, and equipment availability. When disruptions occur, such as staff shortages or missing ground service equipment, the shift leader is the central point of contact. They coordinate directly with the regisseur, ground service equipment teams (GSE), and team coordinators to resolve issues quickly, via phone.

Team Coordinator (TC):
Assigned per flight, the team coordinator is the on-site lead for the entire docking and turnaround process. They instruct the ground handling team and monitor and coordinate the process. They perform the first safety checks once the aircraft is parked, activate the VDGS, confirm the placement of wheel chocks and cones, and activate of the Passenger Boarding Bridge (PBB). Once the aircraft is secure, the TC signals service teams (catering, fueling, baggage, etc.) that they can begin. The TC also oversees timing and ensures that procedures follow sequence, often communicating through direct verbal instructions and manual gestures.

Arrival Crew:
This team is responsible for executing the physical docking tasks. Their work begins with

a FOD check on the apron, followed by placing wheel chocks, cones, and connecting the Ground Power Unit (GPU). If trained, team members may also activate the VDGS and PBB. Communication occurs mostly via visual gestures or verbal cues.

While these are the core ground handling dynamics, they form only part of a much larger web of stakeholders. GSE teams, Schiphol's gate planners, airline representatives, security personnel, and others all contribute to the broader operational ecosystem.
Together, these ground handling roles demonstrate an intrinsic balance of task execution, supervision, and informal coordination. Each role is essential to the current system, yet the tools available to them are fragmented, and communication often relies on workaround-based solutions rather than integrated orchestration systems.

3.2.3 Disruption scenarios

While the intended docking process at Schiphol follows a structured sequence, field observations reveal that operations are frequently subject to real-time disruptions, delays, planning changes, etc. These disruptions often arise from staffing limitations, equipment unavailability, last-minute flight or gate changes, or communication gaps between actors. Often times, these disruptions can set off a chain reaction, disrupting the entire day's operations. To illustrate these vulnerabilities, Figure 8 shows the same swimlane as in Section 3.2.1, but now annotated with potential disruption points and their consequences.

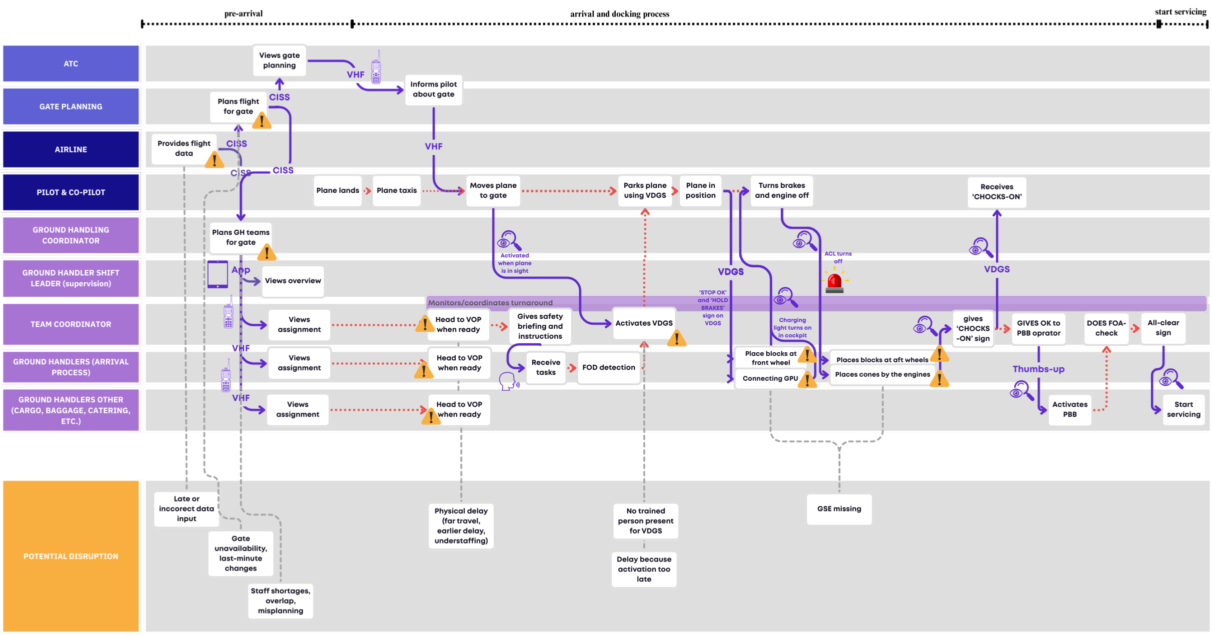

Figure 8: Airplane arrival process disruption scenarios.

Common disruption scenarios include:

Staffing shortages or delays:
If the arrival ground handling team is delayed, whether due to previous task overruns or lack of available personnel, the entire docking process is delayed. Without a completed FOD check and activated VDGS, the arriving aircraft cannot approach the gate. This not only delays parking but also delays the handover to service teams and may block other airplanes (causing them to delay), creating a ripple effect throughout the turnaround process.
GSE unavailability:
A lack of timely access to ground service equipment can delay safety steps such as engine shutdown, power transfer, and PBB activation. These delays often require the shift leader to coordinate manually with equipment providers via phone.

Gate or flight changes:
Last-minute changes to assigned gates require fast recoordination across teams. While the updates may be reflected in CISS, they are not always communicated effectively or in time, resulting in teams being dispatched to incorrect stands or needing to be redirected manually.

Communication fragmentation:
In the event of any disruption, communication between stakeholders remains highly fragmented and mostly manual. Instructions are relayed via radio, phone, or verbal messages, methods that are prone to delays and misinterpretation. With no central or live overview accessible to all parties, teams rely heavily on improvised signals and role-specific updates.

Tool or system downtime:
If a display fails, a portophone channel is overloaded, or a digital planning app is inaccessible, operators often revert to fallback strategies such as personal phone calls or walking across the apron to coordinate. These workarounds, while adaptive, underscore the fragile infrastructure currently supporting orchestration.

Though these disruptions may appear isolated, their effects cascade quickly due to their interdependence. As observed, even short delays at the apron level can affect gate availability, bus scheduling, and the broader airside flow.

3.2.4 Communication tools

Coordination in the aircraft docking process depends on a mix of digital systems, analog tools, and informal communication. While systems like CISS serve as central databases for flight and gate information, the orchestration of daily operations is still highly fragmented and manually driven. This hybrid tool environment introduces friction, as information is not always updated in real time or accessible across roles.

CISS (Central Information System Schiphol) CISS acts as a shared information system across different departments, including ATC, bus regie, gate planning, ground handling regie and apron control. Ground handling planners use CISS to receive flight and passenger data from airlines and to input planning information. However, CISS does not actively support real-time orchestration. It functions primarily as a data repository rather than a dynamic and adaptive system.

Planning Applications and Whiteboards Ground handling coordinators use digital planning platforms to assign teams and tasks. This planning is done manually and input into internal systems and, in some cases, made visible to shift leaders via mobile applications. However, not all team members have access to these digital systems. To bridge this gap, many ground handling companies continue to use physical whiteboards in the break room to display the daily planning.

VHF Radio, Phone Calls, and Verbal Communication Most real-time coordination between team members is conducted via VHF portable radios, which are used to send and receive instructions. In periods of heavy traffic, these channels become overloaded, leading to delays or confusion. In such cases, team members often resort to phone calls or verbal instructions delivered in person. The reliance on these informal communication channels reflects a lack of integrated systems.

Gestural and Observational Signals Many ground handlers use manual gestures, such as thumbs-up signals, waving, or pointing, to communicate readiness or confirm task completion with pilots or teammates. While practical in the moment, these non-verbal cues are not logged or tracked, making it difficult to reconstruct actions or assess performance retrospectively.

In summary, the current tool landscape is functional but fragmented. The lack of a shared, real-time interface across teams means that orchestration relies heavily on personal initiative, experience, and workaround-based coordination. As automation increases, the need for integrated, human-centred communication tools will become even more critical.

3.2.5 Operational pain points

Building on the disruption scenarios and communication analysis, several recurring pain points have emerged that may obstruct efficient orchestration in the current docking process. These are not exceptions but structural pain points that repeatedly surface in daily operations. While some are rooted in technological limitations, others arise from procedural gaps or dependency on human workarounds.

The following pain points are constructed from observed breakdowns and systemic challenges:

- Critical dependency on arrival ground handling teams
- Fragmented communication channels
- Lack of shared situational awareness
- Manual planning and real-time adaptation (instead of predictive or anticipatory)
- Tool/GSE limitations and overload

These operational pain points form the foundation for identifying design opportunities.

3.3 Analytical layer: Systemic interactions

To better understand the systemic tensions that underline the current orchestration challenges, a cross-impact matrix (see figure 9) was developed. This matrix combines four interdependent dimensions, automation, human roles, decision-making, and system integration, and describes how they influence one another in the current operational landscape.

	Automation	Human Roles	Decision-Making	System Integration
Automation		Shifts tasks from manual execution to supervision; unclear role transitions	Reduces need for routine decisions; raises need for AI suggestions	Requires connected systems to scale automation across functions
Human Roles	Must adapt to new forms of supervision; need clarity and training		Humans still hold responsibility, but informal processes dominate	Disconnected systems create workload and ambiguity
Decision-Making	Needs to evolve from manual to supported	Lacks structure; relies on experience and verbal updates		Decision-making is slowed by fragmented or missing data
System Integration	Currently limited; constrains orchestration	Poor access and interface design limit user control	Hinders timely, informed decisions across stakeholders	

Figure 9: Cross-impact matrix.

This analysis reveals that the difficulties faced in daily coordination are not isolated but rather part of deeper misalignments. Automation increases expectations for efficiency and reliability, yet human roles and communication structures have not evolved simultaneously. Similarly, decision-making remains reactive and siloed, partially caused by the absence of real-time system integration.

By mapping these dynamics, the matrix highlights the need for a phased orchestration framework, one that supports evolving roles, improves visibility across actors, and enables automation to function in a coordinated and human-centred manner. These insights form the analytical bridge to the design opportunity space that follows.

4 Strategic framing

This chapter builds on the previous analysis by translating the theoretical and observed insights in the current docking process. It proposes a structured design opportunity space that is both grounded in operational reality and aligned with Schiphol's goals toward increased automation. The framework introduced here sets the foundation for future design interventions.

The literature reveals the pressing need for a centralised, human-centred orchestration strategy to manage increasingly autonomous airport operations. However, observations at Schiphol show that current orchestration is upheld by human physical presence, rather than any digital system integration. Ground teams rely on informal communication, real-time improvisation, and fragmented planning tools, which obstruct coordinated action. Moreover, stakeholders lack shared situational awareness, relying instead on siloed data streams and personal updates.

These findings reinforce the idea that a successful orchestration framework must be gradual, role-aware, and designed for mixed-autonomy environments. It must empower both automation and human decision-makers by supporting visibility, handovers, and resilience through hybrid control structures.

4.1 Problem synthesis

Despite appearing seamless at first glance, the current aircraft docking process at Schiphol is built on manual coordination, verbal communication, and fragmented tools. This human-driven stability creates the illusion of efficiency, but it relies heavily on informal routines, physical presence, and system workarounds.

As automation enters the airside environment, these manual dependencies begin to show strain. Stakeholder observations and the cross-impact analysis (Figure 9) reveal how four critical elements: automation, human roles, system integration, and decision-making, interact and influence one another. These interactions expose systemic frictions that are not easily addressed through isolated technology deployments.

A strategic challenge emerges: while individual technologies or automation initiatives may perform well in isolation, their orchestration across stakeholders remains weak. There is no integrated framework for shared visibility, adaptive workflows, or consistent roles across automation phases. As Schiphol moves from manual to semi- and fully autonomous operations, the lack of such a framework risks disjointed transitions, unclear responsibilities, and fragmented communication.

4.2 Design Insight Mapping

To visualise these findings in a structured way, Figure 10 summarises the core insights into a trajectory: from current state, through identified tensions, toward clearly scoped design opportunities. This table was informed by shadowing sessions, interviews, and literature synthesis.

Design Insight Mapping Table

A small clarification is worth noting in the first row of the table: the current docking process is often perceived as seamless, but this is true only under the assumption of full human control and continuous manual coordination. As automation increases, the meaning of "seamless" must evolve to include new digital and interoperable systems.

4.3 Design opportunity space

The docking process at Schiphol will be transitioning from manual, ground-based coordination to semi-automated and eventually autonomous operations. This shift introduces the need for a future-oriented orchestration model, one that not only accommodates different phases of automation but also reinforces clarity, safety, and cooperation across stakeholders.

This orchestration model must:

Current Insight	Tension / Challenge	Design Opportunity
The current manual process is seamless	As there is partial automation, this process may not remain as seamless—especially as coordination shifts from physical to digital.	Define what 'seamless' orchestration means in a mixed human-machine environment and establish design principles to preserve that experience.
Roles and responsibilities in the docking process are currently based on physical presence and routine practice	As automation increases, these roles will shift, but there is no framework for how responsibilities change across different automation levels.	Redesign the distribution of roles and responsibilities across all involved stakeholders, ensuring clarity and adaptability in both human and machine-led phases.
Monitoring and coordination are performed physically on the apron with minimal remote monitoring or intervention	The transition to remote orchestration requires new monitoring and intervention capabilities, but it's unclear who sees what, when, and acts on it.	Design the remote monitoring and intervention functions of the orchestration system — defining accessible data, visibility thresholds, and protocols for human input.
Communication between stakeholders (handlers, coordinators, controllers) is verbal and unstructured.	Automation will demand more consistent, digital communication flows, yet current practices are dependent on informal and siloed channels.	Redefine communication flows to be transparent, coordinated, and accessible across both physical operations and digital systems.
Schiphol is undergoing a gradual shift from manual to autonomous docking processes	There is no shared vision or structure for how roles, tools, and communication should evolve during this transition.	Develop a transitional orchestration model that evolves alongside increasing automation.

Figure 10: Design opportunity framing

- Be scalable and adaptive to align with different levels of operational maturity.
- Enable flexible workflows, accommodating both human-led and machine-led coordination and the phases in-between.
- Ensure real-time visibility and accountability through improved system integration.

Within this transition, several design opportunities emerge:

- **Redefining roles, ownership and responsibilities** between all stakeholders involved in the docking process, ensuring clarity as human and machine roles evolve.
- **Envisioning monitoring and intervention functions** for remote control areas, including what information is accessible, how it's communicated, and when human input is required.
- **Reimagining communication flows** between ground handlers, control operators, and other actors to ensure coordination is clear, transparent, and actionable, across both physical and digital systems.

These opportunities are all underlined by the commitment to human-machine collaboration, not only from a tool perspective, but from an organisational and cultural one. This includes designing for trust, shared understanding, and the real conditions in which operational staff work. This design space is focused on human-machine collaboration, not only by defining potential future tools and systems, but also by aligning the operational culture and communication with different stakeholder values.

4.4 Phased Development Model

To structure the transition toward orchestrating semi-autonomous and fully autonomous airside operations, this thesis introduces an adaption to the various existing maturity models [38, 24, 75]. This tailored model provides a phased framework for assessing and designing orchestration strategies at Schiphol Airport, based on progressive integration of automation, interoperability, and human–machine coordination.

While these models offer useful theoretical insights, none are directly applicable to the highly collaborative, physically constrained, and safety-sensitive context of Schiphol operations. Accordingly, this Six-Level Maturity Model has been tailored specifically to the orchestration challenges identified. It focuses on the remote ramp and docking coordination, reflecting the unique operational layering between human teams, AI tools, robotisation and centralised decision environments.

1. Manual or semi-manual operations – largely human-controlled processes with minimal automation.
2. Controlled and optimised operations – basic automation supports humans and enhances specific tasks.
3. Intelligent operations – advanced AI systems support routine decisions or offer operational suggestions.
4. Remote operations – humans can supervise or control systems off-site through teleoperation tools.
5. Resilient operations – automated systems can handle disruptions with minimal human support.
6. Autonomous operations – full autonomy is reached, with human involvement primarily for exception handling.

By progressing through these levels, Schiphol can shift toward higher automation while maintaining safety, workforce trust, and operational clarity. This model also supports the design of orchestration strategies that align with current readiness levels of both the technology and the organisational culture, ensuring that automation is implemented as a guided transition.

4.5 Programme of Demands and Wishes (MoSCoW)

To ensure the orchestration framework aligns with Schiphol's strategic and operational values, the following demands and wishes are defined using the MoSCoW method:

Must Have:

- The system must enable a seamless orchestration experience, both for humans and autonomous systems.
- The design must facilitate the human-machine collaboration to reduce ground handler workload.
- The solution must contribute to emission reduction goals by minimising delays and unnecessary vehicle use.

Should Have:

- Compatibility with existing systems like CISS and future smart airport platforms.
- Adaptive workflows that changes across different automation levels.
- Real-time data visibility for all stakeholders.

Could Have:

- Learning mechanisms that improve orchestration through feedback loops.
- Predictive planning tools for staffing and planning.

Won't Have (for now):

- Full removal of human oversight in the orchestration process.
- One-size-fits-all automation levels across all apron zones.

4.6 Selection of design focus: Envisioning Remote Monitoring and Intervention Functions

From the design opportunity space presented in section 4.3, three possible design directions were proposed. Based on the MoSCoW prioritisation, field observations, and systemic analysis, this project will move forward with the opportunity:
"**Envisioning monitoring and intervention functions for remote control areas**, including what information is accessible, how it's communicated, and when human input is required." This direction was selected for various strategic reasons.

First, remote monitoring stands as the basis for the other two opportunity areas: role redefinition and communication restructuring. Without clarity on how coordination and monitoring happens remotely, it would be premature to redefine roles and responsibilities or redefine communication flows. Remote orchestration determines who remains physically present, what needs to be communicated, what tasks must be done manually, and how responsibility is distributed in this human-machine collaboration.

Second, focusing on this direction allows for designing with assumptions: roles and communication flow can be substantiated by assumptions shaped by the functions of remote orchestration whilst the other way around does not work.

Third, this direction is directly aligned with Schiphol's long-term vision. In order for the airside to operate autonomously, there must be a way of remote and centralised orchestration to facilitate these operations.

Last, the current airside operations present a clear gap: there is no central or shared overview, or structured way for a remote operator to intervene or assist. Coordination relies heavily on verbal updates, system workarounds, and physical presence. Envisioning a remote orchestration system addresses these gaps whilst also providing a foundation for the broader transformation.

In essence, orchestration is not just a tool or a dashboard, it is a socio-technical nervous system that enables anticipation, synchronisation, and adaptive response across interconnected components. In Schiphol's future operations, this system must evolve beyond reactive oversight into a proactive, intelligent layer that governs how autonomy and human judgment are interconnected in real time.

This focus area also acts as a convergence point for the three strategic themes that emerged from literature. It responds to architectural orchestration needs by defining the monitoring and control functions required in a future-proof remote control area. It addresses human–machine role transition by clarifying when and how operators intervene in automated processes. Finally, it supports system-level interoperability and scalable orchestration by outlining how diverse autonomous systems can be supervised through a unified interface and coordination logic.

This direction offers clear value to Schiphol by supporting adaptive workflows, real-time situational awareness, and structured human-machine interaction, presenting it as a facilitator for the transition toward safely and effectively coordinated and monitored autonomous airside operations. Importantly, remote orchestration is not only a temporary bridge to autonomy but also a long-term architectural foundation. It enables scalable oversight, structured human fallback, and orchestration logic that matures alongside automation, serving both transitional and end-goal operations.

In conclusion, the central theme that will guide the methodology is:

Design an adaptive orchestration model and maturity-aligned transition plan to facilitate the implementation of autonomous airside operations at Schiphol Airport's docking operations, while ensuring stakeholder integration, scalable automation, and human–machine collaboration across multiple levels of autonomy. the concept of remote orchestration and

the envisioning and designing of the transition towards autonomous airside operations. In the given context of airside operations, more specifically: the docking operations on the apron, remote orchestration refers to a centrally coordinated system in which human operators supervise and intervene in autonomous process from a remote control area. Focusing on this direction ensures the human elements, in terms of oversight, collaboration, and decision-making, remains crucial even as automation increases. At the moment, Schiphol's airside operations do not include a centralised oversight mechanism for coordination and intervention. Introducing a remote (top) orchestration layer addresses this gap by providing real-time situational awareness and structured human-machine interactions. This layer is envisioned as a facilitator for the transition towards toward safely coordinated and monitored autonomous operations. In other words, remote orchestration serves as the basis: it enables human operators to manage multiple autonomous systems remotely, building confidence and reliability on the way to full autonomy. This strategic direction has shaped the use of the different methods, ensuring that technological foresight is grounded within the operational possibilities and maintaining human-in-the-loop control.

5 Methodology

5.1 Introduction

This research mainly uses a foresight-driven approach in planning the transition towards autonomous airside operations at Schiphol. Rather than relying only on current trends, a combination of future-oriented methods has been applied to methodically explore and design the long-term transition [18]. A main reason for this approach is that the future of remote monitoring and intervention is a vision that serves as an extension to the existing future vision of the autonomous airport operations. Accordingly, the research uses this extension of the existing vision as a strategic base point, enabling the transition to be constructed backwards through backcasting [18] while also incorporating a forward-looking approach through forecasting (drawing on three horizons) [25] to ensure the inclusion of human factors, technological readiness, and organisational adaptability throughout each phase. Other used methods include the enablers and blockers analysis, the futures wheel [17], and a cross-impact matrix [37]. Together, these methods provide a planned trajectory based on the future vision and an exploratory analysis of how these changes may develop. This diverse approach is suitable for an intricate, long-term transformation, as it considers the uncertainty involved.

A key structural decision in the methodology was to align the backcasting and forecasting tools with the six-layer maturity model. The six layers mentioned earlier were integrated with the three horizons model [25] and were later further expanded:

- Current situation includes layer 1 (manual operations)
- **Horizon 1** includes layer 2 (controlled operations)
- **Horizon 2** includes layer 3 (intelligent operations)
- **Horizon 3** includes layer 4 and 5 (remote and resilient operations)
- **Future vision** includes layer 6 (autonomous operations)

This alignment provides a structured development road, ensuring each horizon builds logically on the previous one.

5.2 Backcasting Approach

The first method used was a backcasting exercise centred on the long-term vision of fully autonomous airside operations, which includes a functioning and automated remote orchestration system. Backcasting is a planning approach that begins with a defined future goal and works backwards to identify the steps required to reach it [18]. It is a well-suited approach for intricate, long-term transitions due to its goal-oriented and problem-solving nature. In this study, the end-goal (future vision) is defined as "Autonomous Operations", a state of fully autonomous, intelligent, and robust airside operations at Schiphol by 20250. This vision aligns with Schiphol's 2050 vision: "day-to-day airside operations... will be intelligent and autonomous," with an interconnected fleet of self-driving, zero-emission vehicles and automated processes [66]. Using an extension of this vision that includes the orchestration and monitoring of this interconnected fleet and automated processes as the starting point, the backcasting mapped out the necessary in-between states and conditions (H3, H2, and H1) that must be achieved in order to enable the desired future.

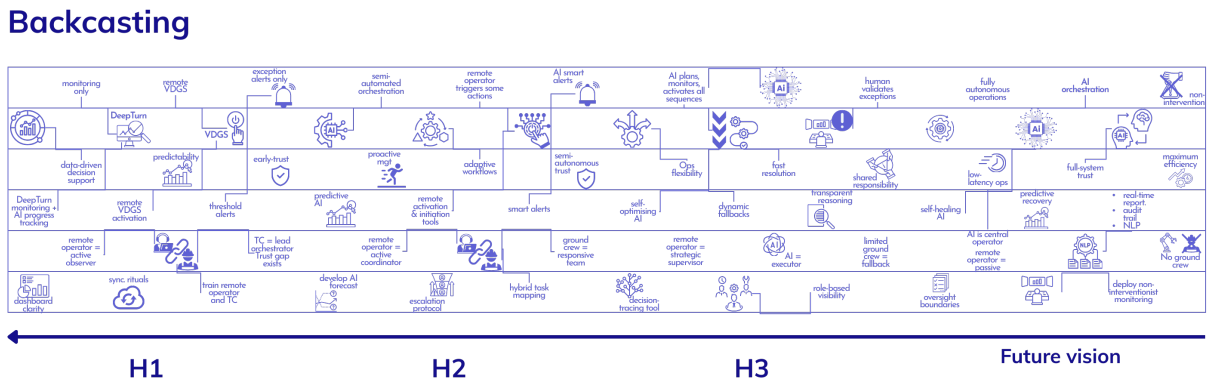

Figure 11: Backcasting Method

Through backcasting, these horizons were defined and mapped in reverse. Starting from the future vision, the analysis identified what must be accomplished prior to each horizon. See figure 11 for an overview of the backcasting analysis. This backward logic ensured that short-term actions are aligned with enabling the long-term goal. For the descriptive backcasting methodology, see Appendix **??**.

5.3 Three Horizons Forecasting Framework

To complement backcasting, a forward-looking approach was also developed. This tool was used to ensure no factors (especially human) would be overlooked in the case of solely relying on a backcasting tool. Forecasting helps structure how a system can evolve in the short, medium, and long term. The model helps map how innovation disrupts existing systems and becomes the new norm through overlapping phases of development. It acknowledges that features of the future can already emerge in the present and that legacy system may persist well into the transition period [25]. By applying this approach, the forecast focused on the social, procedural, and trust-building aspects necessary to, in the end, realise autonomous operations, ensuring not just technical feasibility but also organisational alignment, workforce adaptation, and stakeholder confidence. The descriptive version of figure 12 method can be seen in Appendix C.

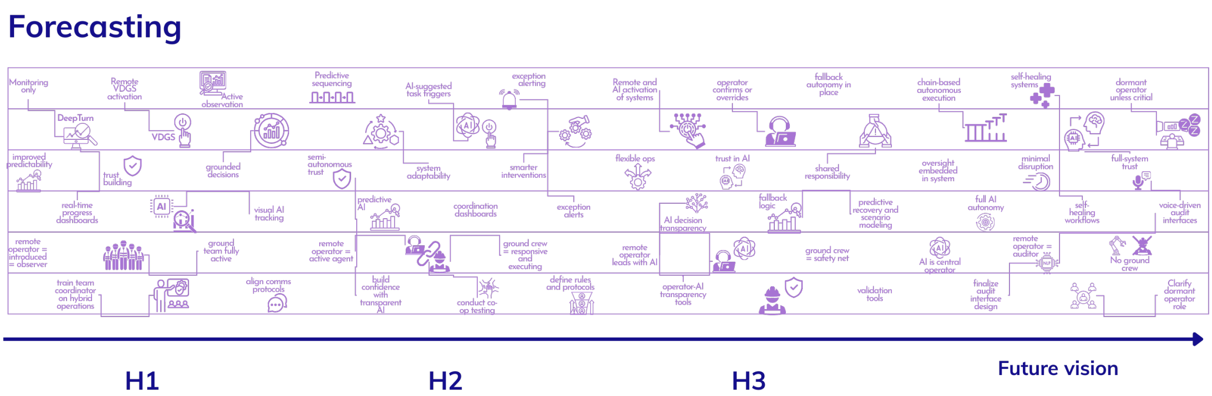

Figure 12: Forecasting Method

Applying the forecasting tool ensures the transition plan is not only technologically sound but socially and organisationally grounded. Each horizon builds upon the last, trust and familiarity in H1, collaboration and transparency in H2, and resilience and oversight in H3. See figure 12 for an overview of the forecasting analysis. This narratives emphasises that autonomy is not just a technical implementation, but rather an all-rounded transformation, through iterative validation, structured communication and the inclusion and empowerment of both human and machine actors. Finally, the forecast shows that preparing humans to work with and oversee AI is just as important as the systems themselves in achieving effective change.

5.4 Overlay Backcasting and Forecasting Methods

To ensure a holistic and adaptive roadmap for the transition towards autonomous airside operations at Schiphol, this study applied both backcasting and forecasting methods. Then an overlay of figure 11 and figure 12 was generated in Appendix D. Based on this overlay, a merged version emerged and is shown in figure 13. Additionally, a comparative table was created to clearly provide the difference in outcomes (Table 1)

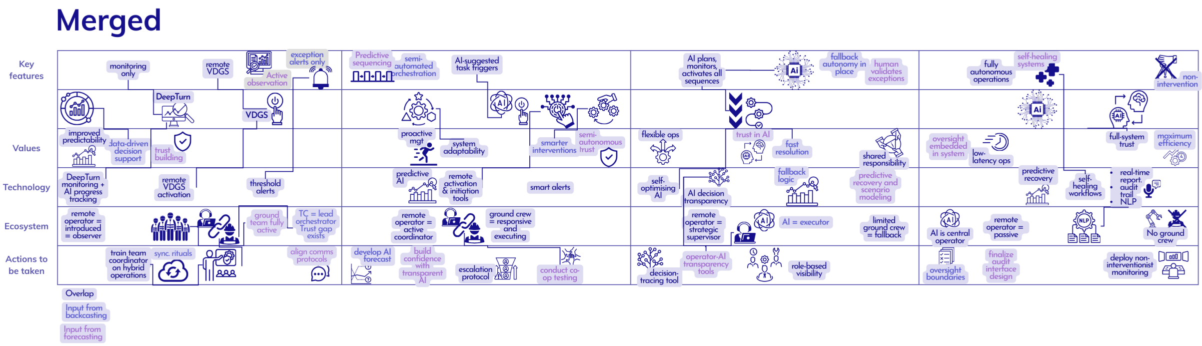

Figure 13: Merged Backcasting and Forecasting Tools

Horizon	Backcasting Key Features	Forecasting Key Features
Future Vision	Fully autonomous, AI-executed operations; human role dormant	Chain-based autonomous execution, self-healing systems; audit via voice
H3: Resilient Ops	AI orchestrates all tasks; operator validates exceptions	AI activates most systems; remote operator confirms/overrides; fallback autonomy
H2: Intelligent Ops	AI offers smart alerts; operator triggers tasks; hybrid logic	Predictive sequencing; AI-suggested workflows; human-AI collaboration
H1: Controlled Ops	Monitoring only; DeepTurn, remote VDGS, exception alerting	Passive monitoring; remote operator as observer; trust-building focus

Table 1: Comparison of Backcasting and Forecasting Key Features by Horizon

The **backcasting approach** begins with the final state, fully autonomous operations, and works backward to define the necessary in-between horizons. This tool provides strategic clarity, focusing on what must be achieved across system capabilities, technology, roles, and governance to realise the 2050 vision. It is particularly useful for defining future-dependent requirement like auditability, fallback autonomy, and legal frameworks. Backcasting assumes a progressing maturity of AI system and technology, with humans gradually transitioning out of execution roles and into oversight or dormant auditing functions.

Contrarily, the **forecasting approach**, structured through the Three Horizons model, starts from the present and considers how innovation and adoption unfold through concurrent phases of change. the forecasting focuses heavily on human and organisational readiness, including cultural adaptation, trust building, and training. It acknowledges that features of the future system may emerge earlier, but also that legacy systems may persist longer into later phases. the forecasting tool is especially valuable in defining how roles (e.g. remote operator, ground handlers) evolve, and how collaborative human-machine workflows are introduced.

Together, the two tools provide a holistic framework:

- **Backcasting** ensures all steps are aligned with the long-term autonomous goal.
- **Forecasting** ensures these steps are realistically adopted by people, processes, and organisations.

5.5 Enablers and Blockers Analysis

This section positions the enablers and blockers as dynamic forces that shape the development of trends and their impact. This section views enablers and blockers as dynamic forces influencing the evolution of trends. Specifically, it examines the shift from manual airside operations to remote autonomy, highlighting how this trend is shaped by opposing forces that accelerate, delay, or redirect its development over time. This perspective acknowledges the complexity of change, avoids linear assumptions, and demonstrates turning points that are strategically significant. The trend underpinning this research is the shift from decentralised, manual airside operations to a model based on remote orchestration and eventually full autonomy. This transformation is not inevitable, it requires mitigation of blockers and leveraging enablers. This is not just a technological shift, but a deep structural transformation involving operational roles, infrastructures, regulation, trust, and organisation culture. To understand how this trend might develop, we examine:

- The current influence of the trend on Schiphol operations.
- Underlying enablers that promote the trend.
- Blockers that obstruct or divert the trend.
- Potential turning points that could mark sharp accelerations or slowdowns.

5.5.1 Enablers: Drivers of Remote Orchestration

Enablers are dynamic forces that may increase the momentum of the trend. The following enablers are currently shaping or are expected to shape the shift toward remote airside orchestration:

- **Operational inefficiencies and labour shortages:** Repeated bottlenecks, increasing personnel turnover, and recurring capacity shortfalls (notably during peak hours) provide strong incentives to reduce dependency on manual labour and reactive handling)[13].
- **Leadership vision and policy alignment:** Schiphol's strategic ambition for autonomous operations by 2050 legitimises and formalises the trend. Organisational support acts as a continuous enabler by securing investment, setting priorities, and creating company narratives [66].
- **Technological availability:** Improvements in perception system (e.g. DeepTurn), predictive analytics, and semi-autonomous vehicles provide the capability base necessary to support remote orchestration pilots [52].
- **Rapid growth of AI agents:** The continuous evolution of AI agents (generative planning systems, anomaly detection algorithms, self learning orchestration models, etc.) accelerates the capability curve. These AI agents enable the coordination of elaborate, dynamic environments, providing more opportunities within remote orchestration [19].
- **Public and organisational familiarity with automation:** Societal exposure to automated/autonomous systems in logistics, urban transport, and service environments contributes to increased social readiness and decreases potential cultural resistance to autonomous orchestration [46].

5.5.2 Blockers: Friction in the System

Blockers are the forces that deflect, obstruct, or delay the progression of a trend. They often stem from legacy systems, organisational stagnation, or socio-cultural resistance:

- **Sociocultural resistance to autonomy:** Operational staff may fear job loss, loss of control, or reduced purpose, especially in roles tied closely to physical presence. This resistance is cultural, not so much technical [60].
- **Regulatory uncertainty:** Aviation's safety-critical context demands proven safety standards. Existing rules often lag behind the realities of autonomous capabilities, particularly concerning remote control [79].
- **System fragmentation:** Ground support at Schiphol involves multiple operators and groups with different systems and procedures. this decentralised system slows integration and makes coordinated orchestration harder [22].
- **Infrastructure dependencies:** Orchestration relies on reliable wireless connectivity and shared task visibility. These require upfront investment and coordination, and gaps here can delay the trend [2].
- **High visibility of failure:** The risk that even a minor incident involving an autonomous or remotely operated process could undermine broader trust across the public and regulatory unions [70].
- **Risk of depreciation:** Fast innovation in AI and autonomous systems increases the risk that early investments may become outdated before they deliver the full expected value. Organisations may hesitate to invest in technology that might need replacing within a few years, especially if long authorisation and/or certification processes delay its use [80].

While these blockers present friction points, many can be mitigated through early co-creation efforts, flexible regulatory engagement, and adaptive rollout strategies.

5.5.3 Turning Points: When Change Accelerates

While it is important to understand the different enablers and blockers separately, their true impact emerges when their interaction is examined, strengthening or weakening the force of a trend. In this project, the different forces converge and clash across different horizons, creating points of inflection in which the development of remote orchestration may accelerate, stagnate, or risk being diverted. The trend from manual airside operations to autonomous remote orchestration does not progress linearly. Rather, it has moments of tension in which the enablers outweigh the blockers, or vice versa. Mapping these interactions allow for the anticipation of crucial moments of influence, suggesting where and when to intervene, preparing stakeholders, and safeguarding the progress (Figure 14).

In some cases, a trend can function as both an enabler and a blocker, depending on how it evolves, making it multifaceted. For instance, as shown in Figure 14 using the red dotted lines, sociocultural resistance is currently positioned as a blocker. However, this same cultural dimension also holds the potential to act as an enabler. Once familiarity with new technologies and systems is established among workers, what was initially a source of resistance can transform into an growing source of support. This shift, from resistance to acceptance and familiarity, marks a critical point of inflection, after which the adoption curve may accelerate as trust and reliance on the new system grow.

Another example is system fragmentation. While fragmentation is indeed a source of complexity, inefficiency, and integration challenges, it also offers a silver lining. It allows the broader problem space to be broken down into smaller, more manageable units. This modularity can become an enabler of change, as it enables focused interventions and experimentation without the need to restructure the entire system at once. In this way, some challenges may also contain parts of their own resolution.

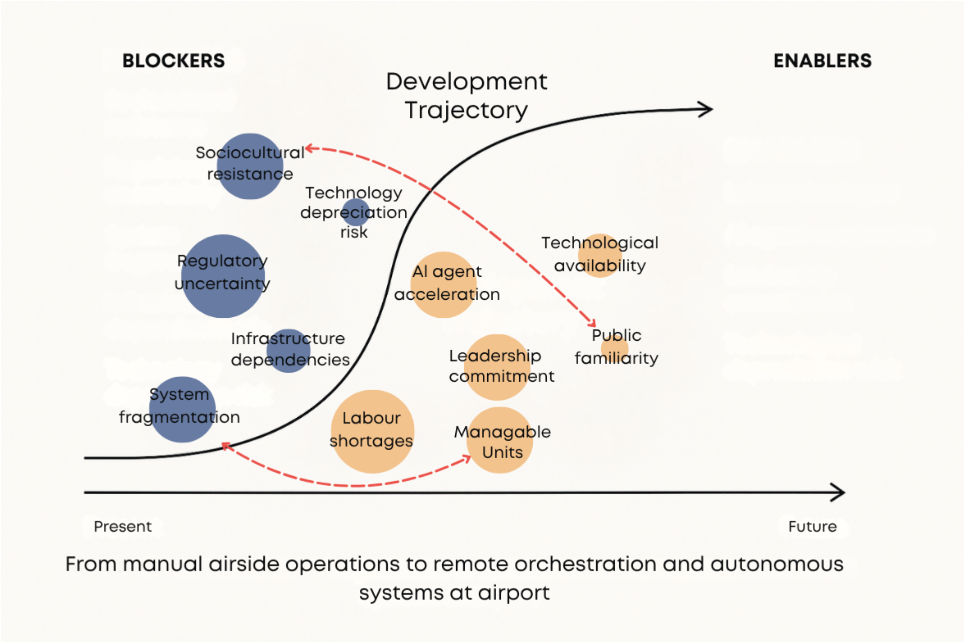

Figure 14: Enablers and Blockers

- **Horizon 1 - Controlled Operations:** At this stage, blockers exert a stronger resistance. Regulatory ambiguity, fragmented systems, and sociocultural concerns lead the transition. However, the enablers like operational pressure from labour shortages and early leadership visions begin to also exert forward momentum. These forces are further reinforced by the rising advancements and availability of AI-based monitoring tools. Effective mitigation efforts, such as stakeholder co-creation, visible pilot success, and early rule-setting, are essential turning points, which depend on successfully managing risk perceptions and initiating the cultural shift through safety pilots and visible performance gains

- **Horizon 2 - Intelligent and Remote Operations:** At this point, the momentum increases. Enablers begin to accumulate: AI agents now assist in planning and intervention, pilot results build legitimacy, and more coordinator ecosystems enable possible scaling up. However, new blockers emerge too: interoperability between systems, partial regulatory alignment, and the retraining of the workforce may lag behind. The turning points here revolve around the successful integration of the different autonomous components underneath the shared remote orchestration layer, which would significantly lower the resistance. Clear communication protocols and co-development of orchestration roles between remote operators and ground staff will be essential to avoid operational delays.

- **Horizon 3 - Resilient Operations:** By this point: the balance leans more towards the enablers. The AI systems are mature, remote orchestration is established, and cultural and regulatory standards have adapted mostly. However, blockers still persist in the shape of systemic fragility: exception case handling, cybersecurity risks, and the cost of maintaining the evolving infrastructure (risk of depreciation). The final major turning point involves ensuring that autonomy is not obstructed by its own difficulties, which requires mitigation strategies such as the deep integration of fallback systems, explainability, transparency of AI logic, auditability, and the standardisation of fallback escalation protocols.

Through these horizons, certain blockers (like sociocultural resistance or system fragmentation) need to be weakened at the start, while others (like technological depreciation or cybersecurity risks) require consistent mitigation. On the other hand, some enablers (like AI maturity and regulatory support) must be actively leveraged over time, through investment, leadership, and proven success.

5.6 Futures Wheel

To explore the 'ripple' effects of achieving autonomous airside operations, the Futures Wheel method was applied. The method places the central scenario: "full autonomous airside operations coordinated via remote orchestration", as its centre. It then systematically maps out its first-order, second-order, and third-order impacts, see figure 15. The tool is intended to identify not only the most immediate impacts, but also the long-term implications across operational, economic, regulatory, social, and ethical dimensions [17]. Notably, it complements the above back- and forecasting analyses by expanding the view beyond direct consequences to include the unintended or conflicting dynamics.

While the Futures Wheel traditionally emphasises outward-moving impacts, this analysis also considers how certain impacts can influence each other in feedback loops. These bidirectional relationships are crucial in complex systems like airside operations, where one outcome can reinforce or constrain another.

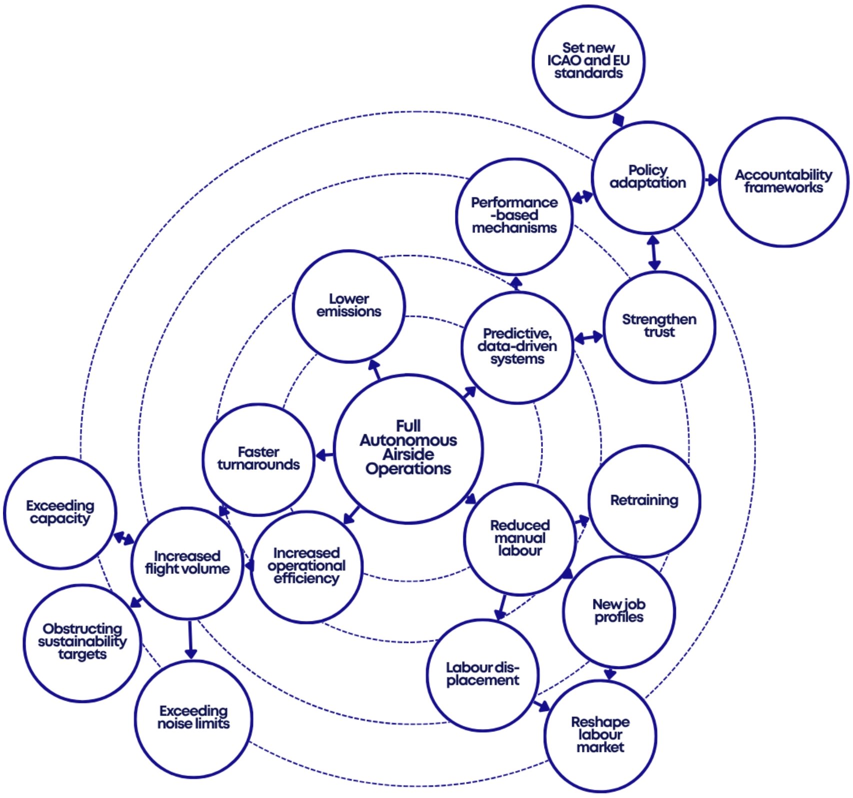

Figure 15: Futures Wheel

In this analysis, the first-order impacts are expected and align with the end-state conditions described in Horizon 3 and the future vision. These include significantly increased operational efficiency, due to predictive sequencing and real-time AI orchestration, along with reduced dependence on manual labour, lower emissions, and improved turnaround time. Human operators take on supervisory or dormant roles, while the autonomous system executes tasks through a loop of prediction, coordination, and correction. These shifts reflect the technological and structural maturity explored in earlier chapters, particularly the development of resilience, auditability, and fallback autonomy in Horizon 3.

The second-order impacts arise from these improvements and insinuate more complex systemic changes. For instance, with less manual work required, ground handling roles are transformed or phased out, necessitating retraining, new job profiles, or even labour displacement in some cases. While operational gains improve efficiency, they may also encourage airlines to increase flight volume, putting pressure on available capacity, sustainability targets, and noise limits. Here, some impacts do not merely follow a linear trajectory. For example, increased flight volume and capacity exceedance form a reinforcing loop: more flights lead to congestion, which can in turn force schedule adjustments that either constrain or redistribute growth. Similarly, predictive, data-driven systems build trust among regulators, which then supports policy adaptation, this in turn feeds back into greater adoption of data-driven practices, reinforcing the cycle.

Simultaneously, the shift to predictive, data-driven systems could strengthen the trust among regulators and insurance companies, leading to policy adaptation and performance-based certification mechanisms. At an organisational level, the airports begin transitioning from an environment of physical and manual operations to one of digital governance, requiring new ways of oversight, accountability, and data transparency.

Moving even further outward, the third-order consequences reveal deeper and sometimes less intuitive implications. The changing workforce dynamic could reshape the labour market, reducing entry-level roles while increasing demand for tech supervisory and maintenance positions. This, in turn, raises broader social responsibility questions for Schiphol and its partners: how should the airport support those workers replaced by automation, and how can inclusivity in new jobs be ensured? From a global perspective, Schiphol's success in employing autonomous orchestration may position it as a standard example, impacting how regulators, airports, and aviation organisations approach autonomy. This may accelerate standard-setting at the ICAO or EU level, but could also raise geopolitical or market challenges around exclusive systems and technology exportability.

Ethical concerns also arise in the third-order layer. As decision-making is increasingly delegated to AI systems, questions surrounding liability, transparency, and societal trust grow more pressing. For example, in the event of an accident involving an autonomous vehicle, it must be clear who is accountable: the manufacturer, the airport operator, or the AI system itself? In addition, as AI becomes the new operational standard, public perceptions of safety, fairness, and human oversight become crucial in maintaining trust. This insinuates that governance, interface design, and audit protocols are not merely technical considerations but rather societal ones, necessitating alignment with public expectations and ethical norms.

By mapping these interconnected layers of impacts, the Futures Wheel reinforces the need for a multi-dimensional transition strategy. While the roadmap may prioritise capabilities such as AI reliability, system fallback logic, and orchestration UI design, it must also proactively address organisational readiness, regulatory changes, and cultural adaptation. In conclusion, the wheel highlights that successful automation does not only depend on technical feasibility, but on the alignment of values, expectations, and accountabilities across all stakeholders.

5.7 Cross-Impact Matrix

To conclude the foresight-driven methodology, a cross-impact matrix was employed to identify how the major influences of the autonomous transition at Schiphol interact with and impact each other. This tool is a multi-factor analysis that is intended to reveal hidden dependencies and feedback loops that influence the evolutions of the overall system [37]. The six identified systemic drivers included in the analysis are:

1. AI and technology maturity
2. Stakeholder trust
3. Workforce adaptation
4. Regulatory evolution
5. Infrastructure readiness
6. Data and connectivity

Each driver was positioned along both the rows and columns of table 2, and the cells contain insights into how the evolution of each impact, influences the other. This structure provides a relational and dynamic understanding of different impacts within the transition. For instance, AI maturity influences regulatory evolution by providing evidence of safe performance, which can stimulate rule updates or trial initiatives. Simultaneously, regulatory flexibility can accelerate the integration of AI by reducing procedural time lag. Similarly,

	AI and Technology Maturity	Stakeholder Trust	Workforce Adaptation	Regulatory Evolution	Infrastructure Readiness	Data and Connectivity
AI and Technology Maturity		Strong performance boosts trust in automation	Enables upskilling and role redesign	Reliable systems push regulators to adapt rules	Depends on robust physical systems to scale	Needs high-quality data streams to perform reliably
Stakeholder Trust	Trust in AI grows with visible safety and performance		Trust improves if workers feel included and supported	Regulators may act faster if public sentiment is supportive	Trust rises when infrastructure is reliable and visible	Transparent data handling and explainable AI increase confidence
Workforce Adaptation	Workforce understanding accelerates AI integration	Trust is built through staff inclusion and communication		Worker readiness supports policy reform and regulator assurance	Training programs depend on infrastructure availability	Workers must engage with real-time systems and dashboards
Regulatory Evolution	Clear rules enable safe AI deployment	Public support influences regulatory pace	Labour law and training standards shape workforce transitions		Regulation dictates infrastructure standards and approvals	Drives requirements for data access, traceability, and auditability
Infrastructure Readiness	AI relies on connected and automated hardware	Trust grows if systems function consistently and visibly	Smart infrastructure reshapes job requirements	Must comply with evolving regulatory mandates		Enables system coordination, uptime, and shared awareness
Data and Connectivity	Real-time data is essential for AI orchestration	Data transparency drives trust	Workers rely on connected tools and intuitive interfaces	Supports evidence-based policy and certification	Critical for infrastructure monitoring and performance	

Table 2: Cross-Impact Matrix: Interdependencies Between Key Transition Drivers

workforce adaptation and stakeholder trust are closely tied together: when workers are well-trained and informed, public trust in the system also tends to increase.

Some key insights from the matrix include:

- **Mutual reinforcement:** AI maturity, data flows, and infrastructure tend to strengthen each other in positive feedback loops. As AI systems improve, they require more data inputs and dependable infrastructure, which in turn become easier to justify investing in [8, 11].
- **Trust as a multiplying factor:** Trust not only depends on technical performance, rather, it shapes everything from regulatory openness to workforce attitude. Accordingly, human-centred design and communication strategies are crucial [10, 5].
- **Sequential hurdles:** Infrastructure readiness and regulatory delays can create system-wide obstructions and delays. If either lags behind, this can slow down progress significantly in other areas like AI or workforce training [36, 72].

By demonstrating how progress in one area may hinder or enable others, the cross-impact matrix emphasises the need for coordinated development across all domains. It aligns with prior foresight tools by encouraging a systemic approach, ensuring that no essential enabler is left behind, and that interconnected dependencies are anticipated in the transition roadmap.

5.8 Synthesis and Roadmap Integration

This final stage of the methodology combines the various methods into a unified, strategic plan toward autonomous airside operations at Schiphol. Synthesising insights from the backcasting and forecasting methods, alongside the enablers and blockers analysis, futures wheel, and cross-impact matrix, provides a well-rounded design that balances long-term vision with short-term feasibility.

The synthesis begins by merging the backcasting, starting from the 2050 vision of autonomous, self-orchestrated operations, with the forecasting model. Where backcasting describes what must happen to reach the final state, the forecasting tool offers insights into how this transition will unfold in real organisational, cultural, and procedural terms. This double logic allows the roadmap to align technical milestones with social readiness, ensuring each step is both necessary and actionable. The trends identified through the forecasting set the context and urgency for many roadmap initiatives, ensuring that planned actions align with anticipated developments in technology, policy, and demand.

The enablers and blockers analysis contributes to this combined trajectory by providing potential accelerators or reasons of delay. If regulatory evolution arises as a crucial enabler, the roadmap incorporates proactive engagement with aviation authorities and policymakers. On the other hand, if workforce resistance of system fragmentation appears as a blocker, targeted mitigation strategies (role redefinition strategies or interoperability standards) may be introduced early. These frictions and turning points help shape the timing and sequence of interventions as well as the contents of the roadmap.

Identified enablers were deliberately leveraged in formulating the roadmap's strategies, providing supportive conditions to accelerate implementation, such as strong political commitment or industry collaboration. At the same time, known blockers influenced the roadmap's path. For each major obstacle, corresponding actions or measures were built into the plan.

The futures wheel reinforces this by drawing attention to ripple effects and long-term consequences that may otherwise be overlooked. For example, improved operational efficiency may lead to increased flight volumes, requiring balancing measures to address sustainability and community impacts. The mapped impacts from the Futures Wheel guided the roadmap to include measures that strengthen positive outcomes and mitigate negative ones.

Finally, the cross-impact matrix ensures dependencies are harboured and taken into account. It emphasises how progress in areas like AI maturity, infrastructure readiness, and stakeholder trust must occur simultaneously and suggests the timing of actions in the roadmap.

In summary, all the foresight insights were synthesised to shape a roadmap that is both ambitious and realistic. Ambitious in pursuing the long-term vision, yet grounded in observed trends and prepared for foreseeable obstacles. By combining the trends, impacts, enablers, and blockers, the resulting strategy reflects a dynamic, adaptive process.

In conclusion, the methodology does more than define analytical tools, it integrates them into a system-wide design logic. This foresight-driven synthesis transforms the transition into a structure and adaptable process. It ensures that each milestone is feasible, that risks and consequences are pre-emptively addressed, and that the final roadmap reflects a balanced strategy. This provides a solid foundation for the transition strategy described in the next chapter.

6 Strategic Roadmap

	H1: Controlled Operations 2025-2030	H2: Intelligent Operations 2030-2035	H3: Remote Operations 2035-2040	*Future Vision 2050*
Key Features Core enabling features	**Observational Monitoring** • Task recognition • Remote Activation • Predictive Forecasting	**Assistive Monitoring** • Condition-Based Alerts • Contextual Recommendations • Predictive Analytics	**Supervisory Monitoring** • Autonomous Activation • Exception Handling • AI-driven Fallbacks	**Autonomous Monitoring** • Self-Healing Systems • Multi-Tiered Fallbacks • Predictive Self-Management
Values Core benefits and principles	• Trust building • Better predictability • Data gathering	• Disruption mitigation • human-machine collaboration • Confidence and trust	• Reduced workforce reliance • Fast issue resolution • Operational clarity and trust	• High system resilience • Maximum operational efficiency • Minimal physical labour reliance
Technology Enabling technological systems	 Smart Progress Monitoring System	 Predictive AI and smart alerts	 Full remote orchestration, monitoring, and activation	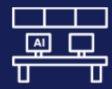 Self-healing AI systems with predictive fallbacks
Ecosystem Stakeholder relationships and environment	 Remote monitoring supports on-site control	 Shared orchestration between humans and systems	 AI-led orchestration with human validation	 Self-orchestrating systems with exception oversight
Actions to be taken Key steps for implementation	• Co-develop collaborative workflows • Train for hybrid workflows • Establish clear protocols for remote activation • Define communication protocols	• Train for human-AI collaboration • Involve direct stakeholders • Establish data governance and feedback loops • Build escalation protocols	• Clarify exception protocols • Co-creation of SOPs • Audit AI decisions	• Establish legal and ethical frameworks • Secure oversight for dormant operators • Formalise emergency interaction systems

Involvement
Estimated relative involvement of human, AI, and autonomous processes

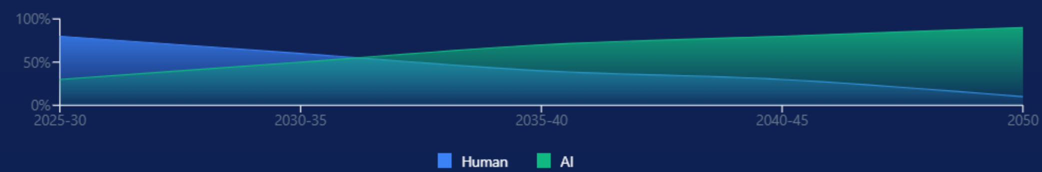

7 Strategic Roadmap Narrative

This chapter presents a roadmap for transforming Schiphol's airside operations from the current manual practices to a future of autonomous orchestration. It builds on the foresight methods and six-layer maturity model introduced earlier, translating them into a clear transition across multiple horizons. Starting with the current state of the art and then horizons 1 through 3, and finally the future vision, each phase is described in detail. The focus is on how monitoring and intervention capabilities change over time through a remote orchestration model. For each phase, the key features, technologies, ecosystem configuration, underlying values, and the recommended actions needed are outlined. This narrative explains not just what changes in each phase, but also why. The rationale behind each element and its anticipated impact. The interactive roadmap (as visualised in the digital tool) provides an overview, while the sections below dive into the specifics. For clarity, the current state of the art with relation to monitoring and intervention functions is described first as a comprehensive starting point. The QR code and link to the interactive roadmap can be found in Appendix E.

7.1 Current phase: Manual Operations

7.1.1 Key features

The current airside operations at Schiphol are predominantly manual and fragmented. Ground handling crews, ground handling coordinators, and team coordinators rely on direct human control and coordination for each task. There is no centralised oversight system for apron activities, coordination happens through radio calls, face-to-face communication, and physical presence. Monitoring of events (e.g. the approach of an aircraft, docking events, etc.) is done by physically present crew. Interventions are reactive and localised: if an issue arises (like a delay or equipment failure), the response is handled in the moment by the team on site or it is verbally passed on. This means that situational awareness is fragmented, no single actor has a real-time complete picture of the turnarounds, which can lead to inefficiencies or missed obstructions. Coordination often depends on implicit knowledge and routine, with workers adapting using informal workarounds when systems or tasks fall short. Overall, the current state of the art is characterised by reliable but labour intensive processes, with limited support from automation or preventive handling and a heavy reliance on on-site workers for problem-solving.

7.1.2 Technologies

In this manual phase, technology exists but in isolated, assistive manners rather than integrated orchestration. For instance, a VDGS at gates provides pilots with docking information (replacing a marshaller), but it operates independently and is activated by on-site ground crew. Communication systems (VHF, CISS, etc.), contain flight and gate information but are not unified into one control platform. To communicate from one platform to another, all information needs to be manually input. Some sensors and cameras are used, but not yet with the intention of centralised monitoring and coordination. In conclusion, the current technology landscape is fragmented: each department has its own tools, and there is no central, real-time shared display of all ongoing docking/turnaround activities. Importantly, remote monitoring or intervention tools are absent, an off-site operator cannot easily observe or assist the process. The enabling technoplogies at this stage are thus minimal for orchestration goals.

7.1.3 Ecosystem

The current ecosystem involves many actors with clear roles operating in a sequential coordinated workflow. Key actors include, aircraft arrival crews, team coordinators, ground handler coordinators, and gate planners. Each actor focuses on specific tasks and communicates with others as needed. The process follows a fixed routine, for instance: ATC guides the aircraft near the gate, VDGS takes over for final parking, then ground crews chocks the wheels, connect ground power, etc., each step triggered by the previous one. Monitoring is distributed: ATC monitors aircraft movement towards the gate, then ground crew oversees and completes the docking tasks. Each team intervenes in its own domain. There is no top-level orchestrator in the current ecosystem, instead, coordination arises from standard operating procedures and face-to-face communication. This ecosystem, while functional, easily leaves gaps in the information flow and can strain efficiency especially in the case of anomalies or delays since all problem-solving is done reactively.

7.1.4 Values

The current state has strong values of safety and established trust. Safety is maintained through multi-actor safety checks. There is a deeply ingrained trust in the experience and judgment of experienced crew who know the operations through and through. These crew members are usually appointed team coordinator and certain tasks may only be completed by them. These human centric values mean new technology is often doubted, trust must be earned before automation takes over. Efficiency and optimal capacity-use are important, but when trade-offs occur, safety and reliability win over efficiency. additionally, values like clarity in communication are emphasised, everyone sticks to established procedures to avoid mistakes. There is also a clear value of accountability with humans in charge of each task, responsibility is clear and hierarchical. This value framework, while ensuring a well-understood operation, can make automation or remote control challenging, as it requires shifting trust from known human processes to new systems. Accordingly, any change must prove it can uphold these values: uncompromised safety, trust-building through transparency, and gradually improving efficiency.

7.1.5 Actions to be Taken

The transition from today's fragmented, manual operations toward the controlled, monitored environment of Horizon 1 requires both systemic and cultural transformation. This shift is not primarily about deploying new technologies, it is about redefining workflows, establishing new relationships between human and machine actors, and introducing orchestration logic into what is currently a sequential, communication-driven ecosystem.

The most immediate action must be the structured co-creation of Horizon 1 capabilities with the relevant workforce, especially ground handlers, shift leaders, and team coordinators. Without active involvement from these groups, the remote orchestration layer risks being perceived as an external and unrelatable control mechanism rather than a collaborative and valuable enabler. Ground staff must not only be informed, but also invited to shape how the remote operator's role is introduced, how system instructions are delivered, and how collaborative decision-making should occur. Previous organisational changes at Schiphol have shown that lack of clarity and involvement may lead to scepticism and resistance. Horizon 1 must learn from these experiences.

To support co-creation, Schiphol should facilitate workshops, pilots, and iterative feedback loops during the development of smart monitoring tools. Ground handlers should test these systems in controlled environments, co-designing elements such as alert timing, activation permissions, and escalation flows. Training programs should focus not only on technical capabilities but also on communicating the "why" behind the orchestration model, emphasising the values and benefits like reduced workload, improved predictability, and operational resilience.

Simultaneously, organisational actions are needed to formalise the remote operator function. Role descriptions must be established, including qualifications, escalation responsibilities, and limits of authority. Operational procedures, such as when and how to activate a VDGS remotely, or how anomalies are to be indicated, must be standardised and validated through pilots and early field deployments.

Additionally, integration of technological systems must begin with small, controlled use cases. The deployment of AI video analytics for smart apron monitoring and VDGS tracking should initially focus on a subset of gates, allowing teams to observe performance, gather data, and build procedural familiarity. The predictive forecasting models can be trained at the same time, gradually introducing real-time decision support into the dashboards used by remote operators.

Importantly, this transition requires cultural alignment. Operators and supervisors must be introduced to the concept of orchestration not as a displacement but as a shift in where and how decisions are made. Communication protocols must reflect this, they should frame the remote operator as an assistant and source of extended awareness, not as a replacement or external controller. On-site crew should maintain the ability to validate or override remote suggestions in early phases.

Lastly, it is essential to integrate continuous evaluation mechanisms. Early deployment phases should include structured observation, performance tracking, and user experience surveys. These will be crucial in adjusting system logic, retraining AI models, and refining coordination before remote capabilities are scaled further in Horizon 2.

Overall, the actions leading into Horizon 1 must balance careful technical deployment with deep human integration. The orchestration platform must not only 'technically work', it must earn its place in the existing ecosystem. That means proving its value, aligning with involved stakeholders, and building trust.

7.2 Horizon 1: Controlled Operations - *Foundation*

Horizon 1 covers the present and near-future. In this phase, Schiphol begins to implement human control with basic automation and remote oversight. The operations remain largely human-driven, but the key change is the introduction of a remote orchestration layer in a monitoring and (limited) control capacity. This is an initial stage in which building trust and familiarity with new systems while humans remain in charge stands at a central point. Data gathering also plays a crucial role in this horizon, as this data will function as the base for features to come. As such, this first horizon stands as the *foundation* of all future horizons.

7.2.1 Key Features

Horizon 1 is the start of the shift from manual processes to a hybrid human-machine operation. While ground handling tasks are still performed by humans, a new remote orchestration layer is introduced in an observational and advisory capacity. Remote operators are situated in a remote control room monitoring live camera feeds and sensor data from the apron. This enables the operators to be virtually present at each stand, with real-time situational awareness of incoming flights and turnaround status. They rarely intervene, instead they monitor and gather data. This data is then used for forecasting and so anticipating anomalies. These predictions are then shared with on-site operators pre-emptively, and, collaboratively, the remote operator and on-site ground handlers plan solutions or mitigation strategies. Additionally, as Horizon 1 progresses, the remote operator gains increasing remove activation capabilities. For instance, remotely activating the FOD detection or the VDGS, once all pre-docking' conditions are confirmed. Using existing technologies, remote operators can immediately know when an aircraft has parked, improving safety alerts and predictability [11, 12]. Overall, Horizon 1 establishes the first "human-in-the-loop" feedback loop: machines enable remote monitoring and recommending, but humans still approve and execute all actions. By the end of Horizon 1, Schiphol has a functioning centralised docking monitoring system and has demonstrated a few remote-controlled actions without any loss of safety or efficiency. Importantly, this phase builds confidence, both the workforce and the technology learn to collaborate in a low-risk setting, collecting early data to inform subsequent horizons. Additionally, enabling remote FOD detection activation and VDGS activation, is the first step towards reducing the reliance on the physical presence of ground operators. Most tasks will still be completed manually, while their autonomous alternatives are in early stages of development; however, these tasks will now be tracked and monitor remotely. Ground crews will carry out these tasks as before, while remote supervisors verify (data gathering) that all steps are completed and whether there are anomalies. In short, Horizon 1 is about supplying human operations with greater visibility and control, without yet automating the entire physical execution of the docking process.

7.2.2 Technologies

The technological focus in Horizon 1 is on enhancing visibility and communication rather than automating tasks outright. A core system deployed is an AI-driven video analytics platform that monitors turnaround activities in real time [12]. It processes live apron camera feeds to automatically recognise events and their completion such as aircraft arrival on stand, passenger bridge connection, cone placement, and more. This provides a live "dashboard" of apron operations, effectively creating a remote window into the turnaround. Through this smart progress tracking system, remote operators can instantly see the progress of the docking process (e.g. aircraft on blocks, bridge connected, GPUs in use) and get alerts for anomalies or delays. Another tool can track the timing of VDGS activations and their outcomes. This tool logs when the VDGS was turned on relative to aircraft arrival and whether any delays occurred, giving insight into how remote VDGS activation affects performance [11]. This tool stands at the forefront of enabling remote VDGS activation. Together, these technologies allow for smarter monitoring and controlled intervention.

The data generated by these monitoring systems will be used to train predictive models capable of anticipating operational delays or obstructions, enabling more proactive coordination. Functioning as a central data point, the dashboard can use predictive forecasting in combination with external factors to output likely future anomalies. For instance, GSE maintenance, poor weather conditions, staff shortages, are all taken into account in combination with the forecast to create a realistic prediction of operations. This prediction can then be used to anticipate problems pre-emptively. It is crucial in this phase, that the communication infrastructure is also strengthened. Control centre operators can use their dashboards to communicate with ground handling teams for effective and efficient communication.

Essentially, all Horizon 1 systems are designed to be human-supervised and failsafe. AI-generated insights or system activations are reviewed and confirmed by human operators. In the event of anomalies or technical issues, remote operators and on-site workers collaborate to mitigate these obstructions. In this way, the technologies introduced in Horizon 1 are not intended to replace human expertise but to enhance situational awareness and support the transition towards centralised, remote orchestration. See figure 16 for a full overview of the dashboard. For additional features, view Appendix F.

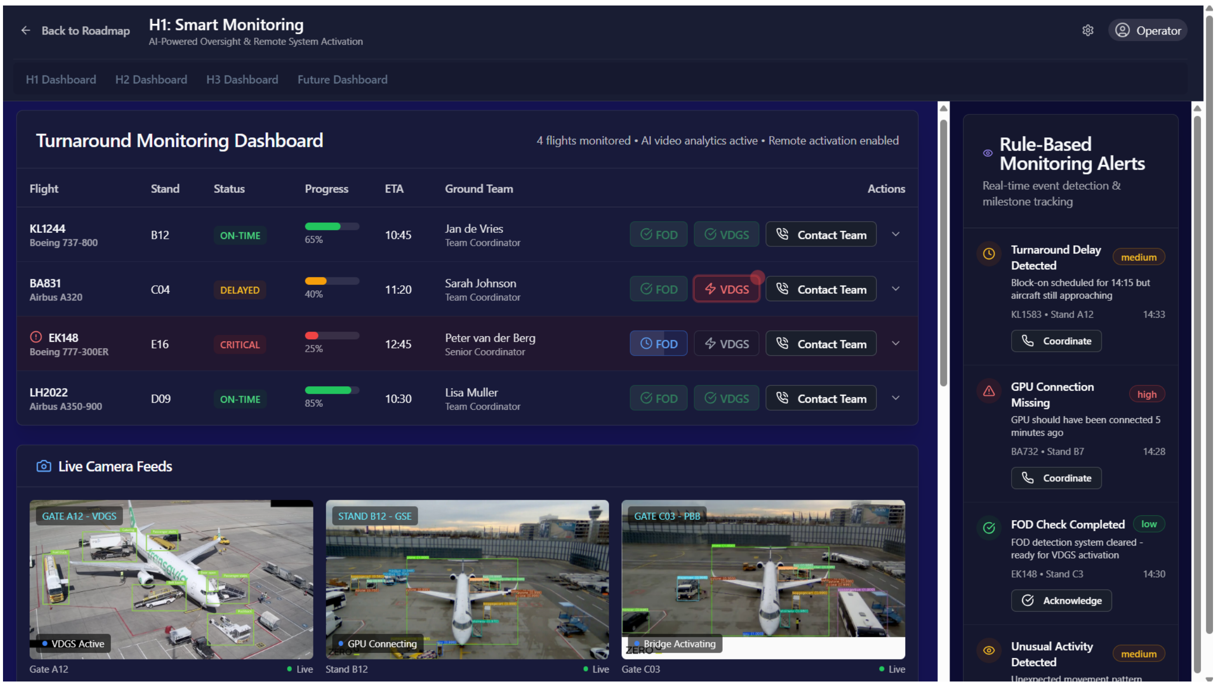

Figure 16: Dashboard Horizon 1

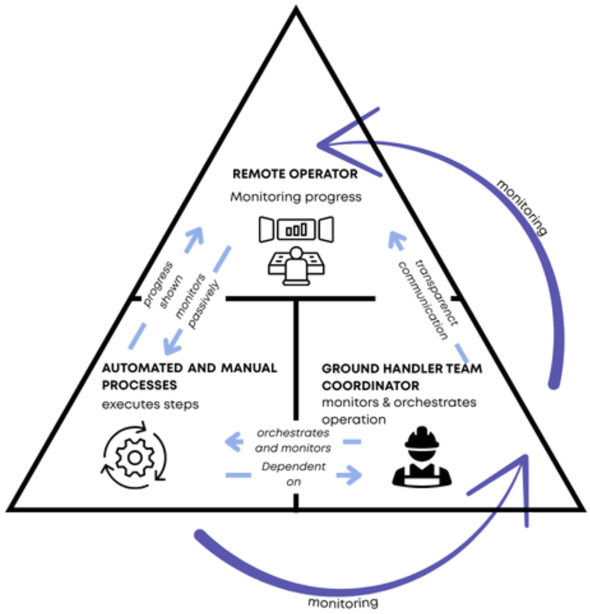

Figure 17: Ecosystem of stakeholders in H1

7.2.3 Ecosystem

The ecosystem in Horizon 1 still involves the same actors as current operations, but with new roles and interactions coming into play. The most notable addition is the remote controller, a trained operator (or a small team of operators) situated in the newly established control centre. These remote operators oversees multiple stands or processes simultaneously via the dashboard. Their job is to monitor progress and to provide guidance or intervention when necessary. For example, when the system is further along and notices an anticipated problem, the operator may inform the relevant ground handling crew pre-emptively. Ground handling crews now collaborate with the control centre. The remote operator may collaborate with the ground handling coordinator with regards to team planning and assignment, the shift leader with regards potential anomalies and so points of attention, and the team coordinator with regards to in-the-moment events or concerns on the apron. The process flow thus becomes more centralised, key milestones of the turnaround are tracked by the remote operator, and exceptions flagged. However, on-site staff remain the ones physically executing tasks and the physically present supervisor (team coordinator) remains present for supervision. Also providing the ground handler coordinator with the new 'shared' dashboard enables further transparency and trust. This empowers the ground handlers to see the value and feel involved in the process.

Also introduced in this ecosystem is an early AI support role, while not a person, the system itself takes on a role of a silent supervisor and coordinator. All actors undergo training to adjust to this configuration, but mostly the team coordinator, ground handling coordinator, and the remote operator since they are especially affected as well as the remote operator being an entirely new function. The culture begins shifting toward one of shared awareness, everyone knows there is an overarching system that oversees the operations. The Horizon 1 ecosystem is thus a combination of old and new, maintaining the familiar human-driven process, but with an added coordination and monitoring layer that connects actors closely together and starts to bridge the gaps between silos, see Figure 17 for the interactions.

7.2.4 KPIs

In Horizon 1, success is not only measured through traditional efficiency and safety standards, but also through the quality of interaction between remote and on-site crew and the effectiveness of new remote capabilities. A central KPI in this stage is the consistency and clarity of communication between remote operators and apron ground teams. The introduction of a centralised orchestration layer relies on seamless collaboration, and so the ability to maintain high-quality, timely exchanges of information, becomes a measurable indicator of progress. This can be assessed through operational logs, feedback loops, and incident response reviews, verifying whether communication breakdowns occurred and if information was shared pre-emptively to avoid issues.

Another critical KPI is the number of prevented disruptions as a result of early forecasting and remote oversight. The orchestration platform's AI components begin flagging delays or conflicts before they unfold, such as potential FOD risks, equipment misalignment, or delays due to staff shortages. These predictions, when acted upon by remote operators or on-site teams, lead to pre-emptive solutions that maintain process flow. The effectiveness of this predictive capability can be measured by tracking incidents that were identified in advance and resolved before escalating into operational disruptions.

Remote operational capabilities, such as VDGS and FOD system activations, also introduce a new performance indicator. Their success can be tracked by the frequency and reliability of remote system activations, and by the extent to which these actions enable more efficient workflows. When remote operators are able to prepare a gate for arrival independently, activating the VDGS only when conditions are validated and safe, or confirming stand clearance via autonomous FOD checks, ground handlers are granted slightly more flexibility in their timeliness, since the aircrafts will not be stuck waiting in front of the apron in case of ground crew delay. A KPI in this area is the reduction in manual dependencies, for example, fewer calls or physical interventions needed to initiate aircraft docking or gate readiness procedures.

Additionally, situational awareness among both remote and on-site personnel remains a key benchmark. With remote visibility tools and AI alerts, the expectation is that ramp supervisors and operators have a clearer, real-time understanding of the turnaround status, equipment locations, and ground crew movement. This can be evaluated through periodic assessments of user feedback and observational studies, checking whether staff feel better equipped to make decisions and coordinate efforts.

Lastly, safety performance under a remote monitoring (and activation) system must be maintained. Even if steps are activated remotely without on-site supervisors (during VDGS or FOD detection), the rate of safety incidents should not increase. Metrics like the number of turnarounds completed without safety interventions, or the rate of successfully completed remote validations (e.g. chocks confirmed in place, stand clear of FOD), provide a tangible way to track safety continuity.

In summary, the KPIs for Horizon 1 are designed to assess whether remote support capabilities are not just functioning, but improving the predictability, resilience, and collaborative efficiency of docking operations. The focus is less on manual task speed and more on a smooth operation, whether communication, foresight, and remote capabilities are enabling a smarter, less human-reliant operation while preserving the core standards of safety and precision.

7.2.5 Purpose and Objectives

The overarching purpose of Horizon 1 is to establish a solid foundation for more autonomous operations while delivering immediate operational improvements. This is intentionally a conservative, learning-focused stage. By keeping humans in charge but introducing selective automation in a controlled manner, Schiphol can validate new technologies in real conditions and refine procedures before scaling them up. A primary aim is to prove that the concept of remote orchestration is feasible and beneficial: that a central team can effectively monitor multiple gates and even execute simple control tasks (like turning on a VDGS or activating the autonomous FOD detection) without compromising safety or efficiency. Demonstrating this is CRUCIAL for gaining stakeholder trust for the more advanced automation planned in later horizons. Horizon 1 is also about allowing the workforce to get used to new ways of working. Operators learn to trust sensor data and software recommendations, while on-site crew learn to collaborate with remote colleagues and act pre-emptively rather than only reactively. The airport adopts an iterative approach: test a new capability in a pilot, validate it, then implement it more widely. In essence, Horizon 1's purpose is to show that automation can coexist with and even enhance human operations, thereby building confidence and excitement for the bigger transformations to come.

Specifically, by the conclusion of Horizon 1, Schiphol intends to achieve several objectives. One objective is to integrate central monitoring across the airport: have a fully operational central dashboard covering all active gates, with live feeds (video, sensor, turnaround status data) accessible to remote operators and ground handler coordinators. This includes leveraging the smart monitoring system on all aprons so that every flight's turnaround events are being tracked. A second objective is to demonstrate remote activation capabilities under safe conditions. For example, Schiphol aims to leverage remote VDGS activations and possibly remote initiation of the autonomous passenger bridge at a subset of gates, and document that these actions can occur without incident over various trials. Successfully doing so validates the remote orchestration concept. Another important objective is to improve efficiency metrics such as the prevention of docking delays and OTP. Meeting set targets would confirm that even in this early stage, the new system provides value. Finally, a crucial objective is to earn trust and support from the people involved, from ramp workers to managers to regulators. By the end of Horizon 1, the goal is that these stakeholders have seen the technology in action and are optimistic rather than fearful about its expansion. They must be involved in the development of the communication protocols and have a say in the overall creation of this new way of working. In summary, Horizon 1's purpose is *foundational*: to get all pieces (technology, people, process) in place and prove that remote assisted and monitored docking can work. The objectives tied to this purpose ensure that the airport not only prepares for the future but also realises short-term gains in safety and efficiency.

7.2.6 Rationale and Values

The primary reasoning for launching Horizon 1 comes from the persistent and structural shortages in ground handling workers at Schiphol. Reports indicate that staff shortages in ground handling and security services have been significant contributors to operational disruptions at Schiphol, leading to flight delays and passenger dissatisfaction. These shortages are not temporary but reflect a broader industry trend where rising traffic levels are not matched by corresponding increases in available ground staff, but rather the opposite [13].

Horizon 1 addresses this challenge by initiating a shift in how key monitoring and intervention tasks are conducted during the docking phase. Instead of placing additional demands on an already strained physical workforce, it begins slowly transferring parts of the responsibility to a remote orchestration layer. This approach does not eliminate the human factor but reallocates it to more scalable and sustainable positions, such as remote operations roles.

Automated progress tracking and alerts act as an extra safeguard. In short, Horizon 1 provides quick wins in efficiency by leveraging technologies that are already quite established or at least low-risk in the industry, such as video analytics and remote monitoring. Importantly, these gains are enabled while keeping human operators in the loop.

A major rationale for Horizon 1 is to build trust in automation, both among employees and regulators, by demonstrating that humans and machines can safely collaborate. The controlled nature of this phase where no autonomous action proceeds without supervision, is designed to show that technology can assist without overriding human judgment. This aligns with Schiphol's innovation philosophy, often referring to the first steps as the "earth" approach. Rather than suddenly taking a disruptive leap, the airport opts for a gradual transition, ensuring each innovation directly supports day-to-day operations and builds organisational readiness for more elaborate future changes.

This phase is as much about change management as it is about technological deployment. It provides Schiphol with the space to learn, refine procedures, and identify potential weak points. For example, if remote operators are overwhelmed when monitoring too many stands at once, adjustments can be made to staffing levels or system features before further automation is introduced.

Horizon 1 is fundamentally a trust-building phase, not just between operators and new technologies, but between the remote orchestration layer and the existing ground handling teams. As remote monitoring and limited remote interventions are introduced, mutual trust, transparency, and collaboration emerge as key values crucial to this early phase in which the foundation is built.

Trust is created through operational transparency and human-in-the-loop control. Every automated suggestion is visible, explainable, and ultimately reviewed by human operators.

Also, the ground handlers have access to the dashboard (monitoring functions only), so that they grasp the value and reasoning behind it. The remote operator is not framed as a supervisor or enforcer, but as a collaborative assistant, supporting the on-site crew through better visibility and coordination. This collaborative positioning is essential in overcoming scepticism and distrust and integrating a shared sense of ownership throughout the transition.

A second core value is learning through data. Horizon 1 marks the beginning of large-scale operational data collection, which is essential for training AI models. These datasets will progressively improve the system's ability to anticipate delays, identify anomalies, and suggest optimal responses. In this way, Horizon 1 is not just about deploying new technology, but rather preparing the ecosystem for anticipatory, data-driven decision-making in future horizons.

Another value is the progression of the seamless inbound flow through remote VDGS activation. Currently, aircrafts may experience unnecessary idle times due to late VDGS activation. Horizon 1 mitigates this by allowing remote operators to safely activate the VDGS, helping reduce idle engine time, lower emissions, and streamline inbound movement. This reinforces a subtle but important value: early-stage efficiency that does not come at the cost of trust or control. repository.

Finally, communication and co-ownership form the foundation of Horizon 1. Ground crew are equipped to both receive and contribute information, enabling feedback loops, encouraging operational reporting, and allowing local knowledge to inform orchestration decisions. In this way, the organisation develops not only safer and more efficient operations but also a culture in which technological progress is co-owned by its users.

Through these values: trust, transparency, data-informed learning, seamless coordination, and human-centric communication, Horizon 1 establishes the ethical and operational groundwork essential for eventually scaling intelligent, remote systems in subsequent phases.

7.2.7 Actions to be taken

To successfully navigate Horizon 1 and build the necessary capabilities for Horizon 2, Schiphol must prioritise actions that reinforce trust, accelerate technological maturity, and enable ecosystem adaptability. These actions should unfold throughout Horizon 1, establishing the foundations for a seamless transition into the next phase of intelligent orchestration.

First, Schiphol must focus on scaling the use of smart monitoring and intervention tools. Now that several of these systems are operational, efforts should be directed toward broadening their deployment. This includes expanding video analytics coverage and integrating new detection features. Smart visual tools should be introduced to provide real-time timelines of each flight's turnaround progress, automatically updating when key steps are completed. These tools make it easier for remote operators to quickly detect delays or out-of-sequence events.

At the same time, the orchestration interface must evolve. What began as a remote monitoring dashboard should become a more interactive coordination assistant. Usability improvements must be informed by operator feedback and should include enhancements in contextual awareness, anomaly tracking, and system responsiveness. Anticipated features might involve predictive alerts, confidence scores tied to AI forecasts, and interactive tools that allow operators to simulate reassignments or interventions. These additions are critical to preparing the orchestration layer for the more dynamic demands of Horizon 2.

To support effective remote orchestration, clear intervention protocols must be formalised. This includes defining the precise conditions under which remote operators are authorised to intervene, distinguishing between observational monitoring and actionable orchestration. Standard Operating Procedures (SOPs) must be developed to reflect validation thresholds, safety requirements, and escalation paths. For instance, SOPs should stipulate that remote VDGS activation can only proceed once FOD detection has been completed and validated on-site.

Training programmes should also transition from basic system familiarisation to advanced human-AI collaboration. Remote operators need to develop fluency in interpreting AI recommendations, assessing reliability scores, and coordinating simultaneous turnaround processes. At the same time, ground handling teams must learn how to respond to AI-informed schedules and remote activation protocols through scenario-based training modules. This dual development supports a future in which humans and machines act as complementary operational agents.

To accelerate the shift, Schiphol should run shadow and hybrid orchestration pilots. These pilots would allow human operators to review and approve autonomous decisions in real time, gradually increasing AI autonomy in controlled, low-risk contexts. For example, the VDGS system could be programmed to activate automatically following validated GSE and FOD clearance, provided that flight details have been confirmed. Over time, these practices build familiarity with AI logic and confidence in its ability to manage routine orchestration tasks.

Equally important is the establishment of strong data governance and continuous feedback loops. Schiphol should define clear roles and responsibilities for collecting, interpreting, and learning from operational data. Internally developed dashboards can help visualise key metrics such as turnaround duration, delay frequency, and the effectiveness of interventions. When anomalies emerge these insights can guide refinements to sensors, AI models, or procedural logic.

As Horizon 1 advances, proactive engagement with external stakeholders becomes increasingly important. Schiphol should communicate measurable results, such as reductions in idle engine time, improvements in turnaround consistency, and faster response times to apron events, to partners including regulators, unions, and airlines. This transparency will help cultivate the trust necessary to expand remote orchestration and automation further in Horizon 2.

Finally, and critically, the success of all these actions depends on the direct involvement of frontline stakeholders. Ground handlers, apron controllers, gate planners, and other affected personnel must be involved in every meaningful step of this transition. Through co-creation sessions, feedback workshops, and regular information exchanges, their insights and values can shape the trajectory of the orchestration system. Simultaneously, this engagement process reinforces ownership and alignment with the broader strategic vision.

To support this transition systematically, figure 18 provides a concise MoSCoW table linking the prioritised actions to each horizon. This table aligns urgency levels with their corresponding interventions and helps visualise how efforts compound over time. For a complete, multi-horizon version of this table, see Appendix J, which includes the full MoSCoW overview. Additionally, Appendix K details the strategic rationale and implementation logic behind each action area.

Horizon	Urgency	Priority	Action
H1	**Must have**	Smart monitoring with AI-enabled video analytics and real-time dashboards,	Deploy AI video systems at select high-traffic gates and integrate dashboards for situational awareness
		Human-machine coordination through remote activation (VDGS, FOD)	Pilot remote activation tools with operator-in-the-loop workflows
		Ecosystem engagement and training alignment	Establish stakeholder workforce for co-development
	Should have	Structured anomaly detection and alerting logic	Define escalation paths and alert thresholds
		Integrated dashboard covering multiple stands	Design and deploy a live dashboard prototype with basic anomaly visualisation
	Could have	Long-term data readiness for AI refinement	Outline data labelling and storage strategy for future model training
		Preliminary automation planning	Document constraints and opportunities for gradual airside automation

Figure 18: H1 MoSCoW

By executing these steps, Schiphol will shift from a phase of cautious digitisation to one of intelligent orchestration. These interventions will ensure that Horizon 2 begins on a strong foundation of operational trust, system maturity, data integrity, and stakeholder alignment, making the transition not a disruptive one, but a confident and structured evolution.

7.3 Horizon 2: Intelligent Operations - *Collaboration*

Horizon 2 envisions the shift from basic remote monitoring to intelligent orchestration. This phase builds directly on the trust, tools, and operational routines established in Horizon 1. While complete autonomy is not yet the norm, Horizon 2 introduces more advanced remote capabilities and predictive systems, allowing Schiphol to move towards a hybrid human-AI coordination model. Autonomous processes begin to scale in defined domains, and the orchestration centre transitions into an active operations centre, where remote operators collaborate closely with intelligent systems and ground handlers to manage the dynamic docking process.

7.3.1 Key features

In this phase, the remote operators make a switch from basic remote supervision to collaborative orchestration. This phase builds directly on the foundation laid in Horizon 1, where data gathering, trust building, and basic remote support proved that remote monitoring could be introduced without compromising safety or staff confidence. Horizon 2 introduces a more active role for the remote operator, who now not only observes but increasingly interacts with and guides apron operations in real time. While various tasks on the apron remain manual, they are now carried out in close coordination with a remote operator who has real-time situational awareness and system support. This is the core dynamic of Horizon 2: collaborative human–machine operations, with shared control between the apron and remote control area.

A major advancement in this phase is the introduction of smart alerting. The orchestration system can now forecast disruptions well based on years of data gathering and training and generate actionable insights. These alerts are not generic status notifications, but targeted actionable messages connected to specific process anomalies, emerging risks and simply providing the operator with activation reminders. For example, if a stand's FOD detection camera identifies debris and confirms clearance, the system may automatically trigger a recommendation to activate the VDGS, pending operator confirmation. In a more basic scenario, the smart alerts will provide the operator with notifications like: "Activate VDGS at stand C05 in 5 minutes" with a button below it through which the operator can do so. This allows operators to guide events with confidence and without needing to wait for manual inputs from the apron team. Importantly, alerts are paired with system-suggested actions. Rather than simply raising an alert, the system proposes a solution and allows the operator to act or delegate (See figure19 . These suggestions are shared immediately with relevant stakeholders (e.g. ground crew or gate planner) through digital channels, creating a shared situational picture.

Predictive analytics, trained on Horizon 1 data, now play a central role. The system uses these insights to prevent problems before they escalate. If the system detects that at a certain gate there are consistently delayed GPU connection, it may prompt the remote operator to notify a technician in advance. In more advanced cases, it may even suggest rescheduling tasks or altering the ground sequence based on congestion forecasts or equipment availability. In Horizon 2, operators oversee multiple stands and are supported by dashboards that filter their attention to the turnarounds most likely to need intervention.

This phase also introduces semi-automated activation of apron tasks. Operators can now remotely initiate more functions, while still verifying their status via camera and sensor confirmation. These tasks, while still supervised, no longer require manual execution on-site. However, some systems still require joint participation. For instance, a remote operator may need to coordinate with an on-site ground handler to ensure a sensor has verified clearance before remote activation can proceed. This kind of human-machine teamwork is a defining trait of Horizon 2.

Exception handling becomes more intelligent. Alerts are ranked by urgency, and the system explains why it has flagged a case. A deviation in turnaround time will trigger not only a warning but also a proposed corrective action. For instance, the system may recommend deploying an additional handler or reassigning a vehicle to another stand, and the operator can approve or adapt that suggestion. These decisions are logged, creating a basis for learning and process refinement.

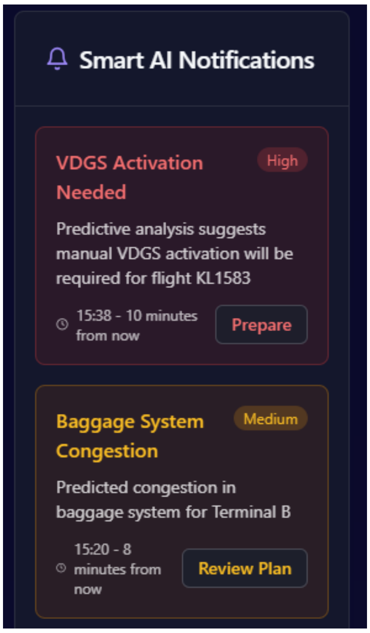

Figure 19: Smart AI notifications

In Horizon 2, remote operators and apron crews operate as a coordinated unit. Decisions are no longer made separately but informed by a system that combines prediction, live status, and stakeholder inputs. This phase transforms remote operations from passive support into real-time collaboration. While automation grows, human judgment still plays a central role. This prepares the ecosystem, both technically and culturally, for the further developments in automated orchestration and full remote execution that will characterise Horizon 3.

7.3.2 Technology

The technology in Horizon 2 reflects a step forward in integration, responsiveness, and operational intelligence. Building on the data foundations laid in Horizon 1, the orchestration platform now incorporates predictive analytics that are actively used in daily decision-making. These analytics process the accumulated turnaround data to anticipate likely anomalies or obstructions in the docking sequence and provide early warnings. Instead of responding to issues after they occur, remote operators now receive suggestions in advance, allowing for quicker decisions and, in some cases, automated mitigation.

Remote activation capabilities also expand. In Horizon 1, remote process activation required coordination and remained small-scaled. By Horizon 2, it becomes a routine part of operations. Following verified FOD clearance, operators can activate VDGS remotely without delay, ensuring inbound aircraft are guided immediately and have minimal idle times. Other processes are also being automated (like GPU units, robotic chocks, etc.). These may be activated on-site or remotely, but each system reports back to the orchestration dashboard, confirming activation status, completion time, and any anomalies.

Operators are provided with a real-time graphical interface showing live apron conditions, equipment positions, aircraft statuses, and task progress. These digital visualisations do not replace team coordinator oversight yet but enhance situational awareness, allowing for faster interpretation of complex events across multiple gates. For processes remotely activated (FOD detection and VDGS), the remote operator acts as the supervisor instead of the team coordinator. For manually exectuted tasks, the team coordinator still monitors on-site.

As these technologies become integrated in daily practice, human oversight remains central. All automation is still subject to manual override, and any task initiated by the system requires confirmation where safety might be affected. Smart alerts generated by the AI come with reasoning, allowing the operator to understand the recommendation before taking action. Each interaction is logged, creating a history that can be reviewed for accountability, training, or performance analysis.

Overall, the technological environment of Horizon 2 creates a reliable framework for distributed, semi-autonomous operations. It allows remote orchestration to move from a supporting role to a more active, directive position while still respecting the importance of human experience and input. This hybrid structure is essential for enabling trust and confidence among all stakeholders as Schiphol moves toward full autonomy in later horizons.

Interesting to note in Horizon 2, is that cognitive orchestration system will have reached a maturity that could support higher degrees of autonomy. However, this capability is restrained on purpose. The reason is strategic. While the orchestration system might be ready, organisational, on-site process, and the ecosystems are not. Therefore, the roadmap intentionally deploys cognitive autonomy in alignment with social and procedural maturity, not its own technical potential alone. See figure 20 for a full overview of the dashboard. For additional features, view Appendix G.

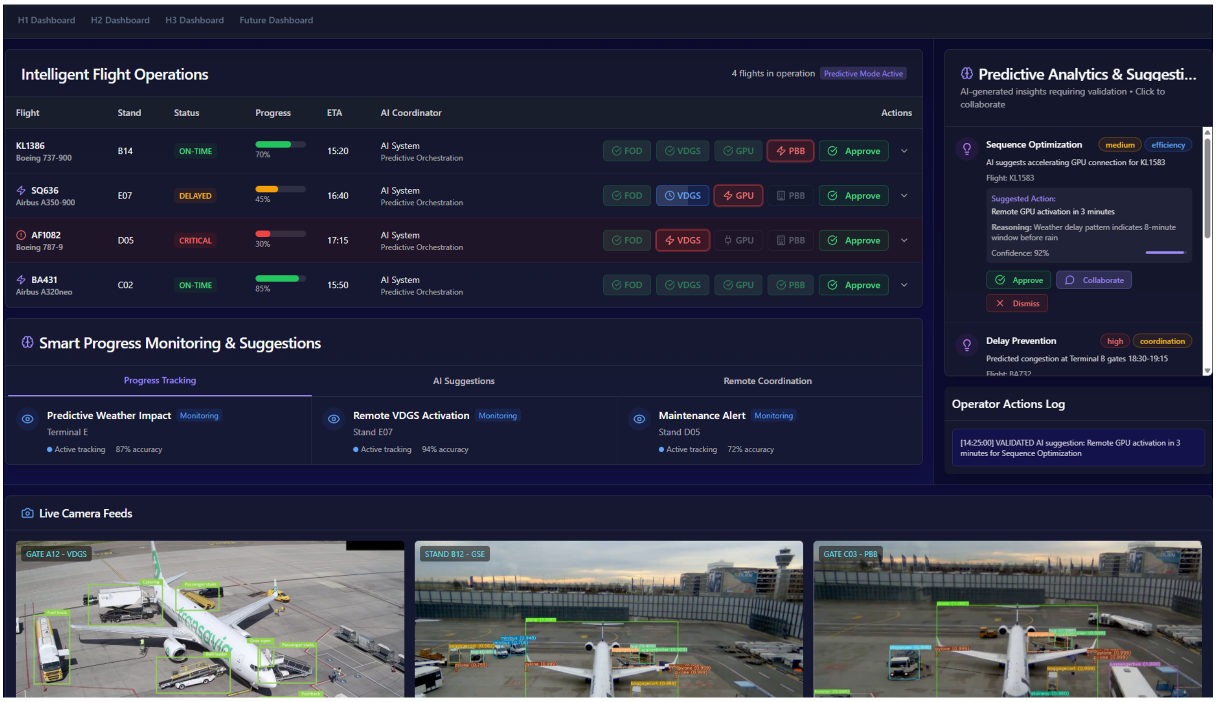

Figure 20: Dashboard Horizon 2

7.3.3 Ecosystem

In Horizon 2, the operational ecosystem begins to take on a more coordinated structure between digital orchestration and physical execution, while still remaining firmly human-centric. The foundation laid in Horizon 1, built on trust, transparency, and improved monitoring, now supports more active collaboration between remote operators and on-site crew. However, this collaboration is not yet built on shared digital interfaces or fully integrated systems. Instead, it is driven through communication, procedural alignment, and operational feedback loops.

The team coordinator continues to play an essential role as the local supervisor of turnaround execution. They remain physically present at the stand and are ultimately responsible for the safe and timely progression of all on-site activities during the docking phase. Their position at the centre of the on-ground operation allows them to make real-time decisions and have direct context perception.

Rather than replacing this role, the remote operator serves as an additional layer of situational awareness and orchestration support. The remote operator oversees multiple stands simultaneously through the orchestration dashboard. This operator monitors live updates, receives predictive alerts, and initiates certain orchestrated actions. Communication be-

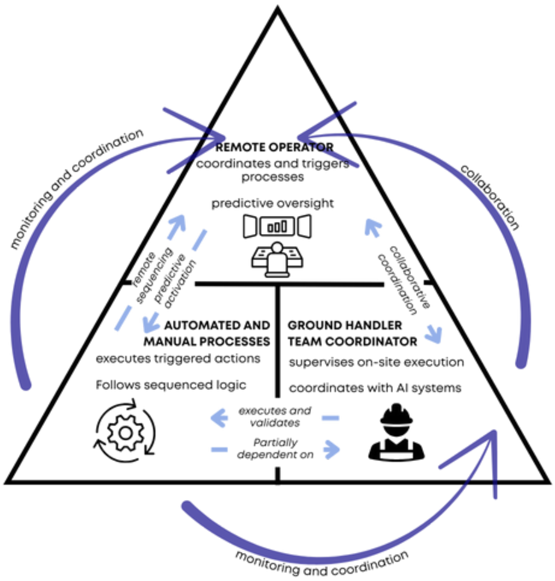

Figure 21: Ecosystem of stakeholders in H2

tween the remote operator and the team coordinator becomes increasingly central to this horizon. For example, the remote operator might notify the team coordinator when a delay is predicted in chock placement due to ground crew being delayed, allowing the coordinator to adjust personnel accordingly. Combining the overview on-sit and remote, the remote operator and team coordinator have a full-rounded overview of the operation. In collaboration, they can foresee and prevent anomalies, or at least mitigate the consequences.

This ecosystem relies on collaboration rather than automation. While AI-generated alerts and forecasts guide the remote operator, decisions about how and when to act are still confirmed with human input. There is no direct control of the on-site team by the orchestration layer, only certain tasks are now executed autonomously. However, team coordinators and ground crews continue to execute various tasks manually, while gradually incorporating autonomous executed processes into their decision-making process. For instance, a remote operator may signal that an arriving aircraft is predicted to dock early and that FOD clearance has already been confirmed. The team coordinator can then respond by readying the GPU or chocks more quickly, closing the timing gap and reducing idle time.

Training and procedural clarity are key in this phase. Team coordinators must understand not only the timing and content of remote prompts but also the logic behind orchestration decisions. Meanwhile, remote operators need to remain sensitive to the actual context of on-site execution and must avoid overloading the team with unnecessary signals. This two-way understanding requires active alignment between remote operator and the ground handling teams.

As AI begins to suggest actions like manual VDGS activation or flag potential stand anomalies, this system must be seen as a supporting actor, not a directive one. The remote operator interprets these insights within the context of real-world constraints. If suggested actions are not the right fit in the perception of the operator, he and the team coordinator always have the control to override or choose otherwise. This mutual understanding of roles and responsibilities becomes the backbone of Horizon 2's collaborative model.

Ultimately, this horizon is about forging functional partnerships between remote and local stakeholders. It prepares Schiphol for the more autonomous operations of Horizon 3 by first establishing reliable patterns of communication and a working division of responsibility. As trust increases and orchestration accuracy improves, so does the willingness of all actors to rely more on remote coordination in the future.

This growing interplay sets the stage for greater autonomy in the future, while respecting the expertise and authority of the current workforce. See Figure 21 for an overview of these stakeholder interactions in Horizon 2.

7.3.4 KPIs

To evaluate the effectiveness of Horizon 2, Schiphol will monitor a set of KPIs focused on orchestration quality, operator readiness, and system responsiveness. These indicators serve not only as standards of progress but also as feedback mechanisms for adjusting workflows and training as the system develops.

A primary KPI is the success rate of AI-suggested task sequences, which reflects how often predictive orchestration recommendations are executed without modification. A high rate indicates alignment between system-generated plans and real-world conditions, reinforcing trust in the orchestration model. In contrast, frequent overrides or dismissals may signal either limitations in the AI logic or gaps in operator-system calibration.

Response time to exception alerts is another critical metric. As predictive systems flag potential issues, the speed with which operators and ground teams act determines whether those insights translate into operational gains. By tracking average response time across event types, Schiphol can assess the practical value of smart alerts and identify where additional automation or support might be needed.

Operator confidence in the system's decision support capabilities is equally important. This can be measured through surveys and observational studies, asking operators how often they trust, rely on, or challenge AI-generated recommendations. As confidence rises, it indicates that operators see the orchestration platform as a useful partner rather than a monitoring burden.

Training completion rates provide insight into workforce readiness. As the scope of orchestration expands, operators and team coordinators must be upskilled accordingly. A strong correlation between training completion and successful orchestration performance would validate the effectiveness of the human-machine collaboration model.

Together, these KPIs support an overarching goal: to shift from reactive execution toward proactive coordination. When successful, this transition leads to tangible outcomes such as shorter turnaround times, smoother handoffs between on-site and remote actors, and fewer disruptions linked to miscommunication or delay. Horizon 2's metrics reinforce its strategic goals: smarter orchestration that enhances efficiency while progressively easing the manual workload.

7.3.5 Purpose and Objectives

The purpose of Horizon 2 is to expand Schiphol's orchestration capabilities through collaborative workflows, predictive AI, and remote activation of critical turnaround processes. This phase builds directly on the trust established during Horizon 1, evolving from basic oversight into coordinated orchestration where intelligent systems and human operators together manage the docking sequence.

The primary objective is to transition from passive monitoring to a more active, intelligent orchestration model. Remote operators are no longer only observers but begin to steer ground operations based on real-time and predictive insights. AI tools are integrated to forecast potential disruptions and propose responses, while operators retain control by validating, adjusting, or overriding suggestions when necessary. This co-decision model ensures that human expertise continues to guide the operation while leveraging the strengths of machine-driven foresight.

Dashboards evolve from status overviews into operational control centres, presenting smart alerts that flag emerging issues and recommend actions. These alerts are not isolated, they are automatically shared with the relevant stakeholders to enable faster coordination and accountability.

The rationale behind Horizon 2 lies in the operational benefits of proactive intervention. Rather than reacting to delays or misalignments after they happen, Schiphol moves toward a model where disruptions are anticipated and mitigated earlier. The system becomes not just smarter, but more valuable to its users, precisely because its suggestions are traceable, interpretable, and consistently aligned with on-site context. In doing so, the orchestration platform earns the confidence of both operators and ground handlers.

To support this evolution, several enabling projects are leveraged during this phase. These include the development of predictive alerting to detect and surface emerging delays or safety risks, and the secure remote activation of various docking processes once verification protocols are met. A key enabler is the continued development of a orchestration interface, that integrates real-time monitoring, actionable alerts, and task initiation tools into a single control environment. Finally, specific operator training ensures that remote operators can confidently manage multiple simultaneous workflows and interpret AI logic in real time.

From the Futures Wheel analysis, it can be deduced that social resistance likely will reach an inflection point during horizon 2. As operators repeatedly engage with the AI orchestration system and observe the accuracy of its predictions and the value of its recommendations, initial resistance may start to die down. Conversely, familiarity and trust will take its place. This shift from caution to confidence marks an important cultural enablers. Something Horizon 2 takes advantage of in establishing deep collaboration.

In short, Horizon 2 is about establishing human-AI collaboration and building operational routine in remote orchestration. The systems introduced in Horizon 1 now evolve into active partners, and the people using them begin to orchestrate, not just observe, the future of apron operations.

7.3.6 Rationale and Values

The strategic rationale for Horizon 2 lies in advancing from a phase of observation and trust-building toward one of operational collaboration and guided autonomy. While Horizon 1 proved that remote oversight could coexist with manual processes, Horizon 2 demonstrates that intelligent orchestration can measurably enhance those processes. The shift is subtle but impactful: instead of reacting to events, the system starts anticipating them, offering operators targeted actions at the right moment.

This horizon is marked by the operationalisation of smart interventions. Predictive analytics are no longer passive back-end systems but active participants in the workflow. They surface early warnings and suggest immediate, relevant actions. These interventions are not only automated but also situationally aware, taking into account the current context of each turnaround. When approved by an operator, alerts are shared directly with the involved teams, transforming a simple insight into a coordinated response.

Adaptability becomes a defining value. The orchestration system is designed to accommodate variation and uncertainty, adjusting plans dynamically as inputs evolve. At the same time, the people interacting with these systems, remote operators, team coordinators, and ground handlers, are supported in developing the flexibility required to shift from fixed routines to collaborative decision-making. The operator learns to manage by exception, engaging only when the system shows a non-routine situation or when manual input is required. This division of attention helps scale orchestration across multiple stands without compromising control.

Underneath this collaboration is a strengthened trust between humans and the system. Each AI-generated alert is accompanied by transparent reasoning and a confidence score, making it easier for operators to judge the validity of the suggestion. These design features are intentional values integrated in the system to maintain operator confidence and accountability. The fact that remote operators retain final authority in every decision reinforces this trust, creating a working relationship where automation is valued as a partner, not a replacement.

Horizon 2 brings Schiphol closer to a future where remote coordination becomes reliable. It does this by supporting smart, timely, and context-aware interventions, encouraging flexible and responsive teamwork, and making transparency a core part of every AI-human interaction. The values developed in this horizon are not technical enhancements, they are organisational commitments that set the tone for the fully autonomous capabilities envisioned in Horizon 3.

7.3.7 Actions to be taken

To prepare for Horizon 3, Schiphol needs to build on Horizon 2's gains by integrating the systems, practices, and collaborations that support semi-autonomous orchestration. This requires focused actions across operations, governance, training, and system integration.

First, the orchestration platform must be upgraded to support more intrinsic, predictive logic. Dashboards should allow operators to simulate suggested actions, display system confidence levels, autonomous activation of processes, and manage multiple workflows efficiently. As orchestration logic advances, autonomous asset deployment should become routine. To enable this safely, Schiphol must continue standardising communication protocols and fallback mechanisms across these systems.

Resilience planning is essential. Clear exception protocols must be defined and stress-tested, ensuring the system and workforce can respond effectively to disruptions. Supporting this, a simulation environment should be used to validate orchestration logic and trial future features under controlled conditions.

Workforce development remains a priority. Remote operators should be trained to manage larger process volumes and validate AI-driven decisions. At the same time, ground crew and team coordinators, whose responsibilities shift in this phase, must be equipped to engage with orchestration tools and adapt to collaborative workflows.

AI governance must also develop. Schiphol should establish internal mechanisms, like a governance board or assurance group, to check orchestration decisions, review exceptions, and ensure system logic remains fair, explainable, and aligned with safety standards.

Stakeholder alignment is critical throughout. Horizon 2 introduces more advanced interactions between remote operators and ground crews, which must be supported by trust and clarity. Co-creation is key: team coordinators and gate planners should be actively involved in shaping SOPs for shared orchestration. Regular workshops and feedback loops can ensure on-site knowledge informs orchestration rules, while also facilitating involvement. Without this alignment, even technically sturdy interventions may face resistance, limiting their impact.

To support this transition in a structured manner, figure 22 presents a concise MoSCoW table mapping the prioritised actions to their respective horizons. It highlights the urgency and impact of each intervention and illustrates how these efforts build progressively over time. For a complete cross-horizon breakdown, refer to Appendix J, which provides the full MoSCoW framework. Further strategic rationale and implementation logic can be found in Appendix K.

Horizon	Urgency	Priority	Action
H2	Must have	Predictive analytics and contextual alerts	Train forecasting models using Horizon 1 data and integrate predictions into dashboards
		Operator intervention capabilities with AI recommendations	Deploy interfaces that allow operators to accept/reject AI-generated action suggestions
		Expanded remote task execution	Scale up remote activations to include multiple turnaround functions
	Should have	System interoperability with airport and third-party systems	Define data and API standards for integration with CISS and equipment control
		Operator override protocols for hybrid decision-making	Develop rules for manual overrides and automated handover during exceptions
	Could have	Fallback logic for partial autonomy	Introduce early-stage fallback scenarios and test operational impact
		Trend analysis for persistent anomalies	Begin long-term logging and classification of recurring event patterns

Figure 22: H2 MoSCoW

By executing these steps, Schiphol can transition from a semi-automated environment to a resilient orchestration model, where smart systems, empowered operators, and engaged stakeholders collaborate in a high-performance, future-ready operation.

7.4 Horizon 3: Remote and Resilient Operations - *Autonomy*

Horizon 3 represents a near-realisation of the autonomous vision. In this phase, operations are remote and resilient, the system handles most situations autonomously, with humans now in a high-level supervisory layer. The orchestration is remote in that it can be managed from a central location (off-site), and resilient in that the autonomous ecosystem can adapt to disruptions or unexpected events with minimal human input. This is effectively the result of the previous steps in the roadmap, an environment in which autonomy is predominant but still under human supervision for safety.

7.4.1 Key Features

By Horizon 3, autonomous airside orchestration at Schiphol operates with high maturity and confidence, built on a foundation of self-optimising systems, layered resilience, and transparent human oversight. The orchestration platform is no longer static or based on 'rules'; instead, it is adaptive and self-improving. Through continuous learning from operational performance, disruption patterns, and contextual data inputs such as weather, aircraft arrival deviations, and ground traffic fluctuations, the AI finetunes its sequencing, task allocation, and resource deployment strategies. This ensures that the system becomes more efficient, reliable, and contextually aware over time, becoming a self-optimising orchestration layer that evolves simultaneously with the operation it orchestrates.

This operational maturity is supported by robust fallback logic. The entire system is designed around resilience. At every level, there are multi-tiered fallback mechanisms integrated to guarantee continuity. In case of anomalies, the AI immediately executes alternate action plans. These might involve dispatching an alternative vehicle, shifting the timeline, or isolating the flawed process without affecting the rest of the operation. Communication and decentralized coordination logic provide backup mechanisms so that even if part of the system degrades, the orchestration continues without collapse.

Importantly, human oversight in Horizon 3 is preserved in a strategic supervisory layer. Operators do not execute or directly coordinate tasks but instead validate high-risk decisions, inspect the logic behind AI actions, and intervene only when the system flags uncertainty or exceptions. This is enabled through advanced operator dashboards that offer real-time access to AI decision paths, confidence scores, and historical actions (Figure 23). The ability to trace every automated decision back to its origin, whether a predictive model, a sensor input, or a decision rule, reinforces both trust and accountability in a highly autonomous context.

Figure 23: AI Decision Logic

Finally, the ecosystem as a whole is interconnected and transparent. Data from autonomous vehicles, infrastructure sensors, and orchestration subsystems flows into a centralised situational awareness layer, visualised in the orchestration interface and accessible to all relevant stakeholders. This integrated monitoring environment provides comprehensive visibility across turnaround progress, equipment status, and predictive risk indicators. Such transparency ensures that while human involvement is minimal, organisational oversight and confidence remain high. Together, these features define Horizon 3 as a phase of resilient autonomy, where automation is not only functional, but trustworthy, adaptable, and robust.

7.4.2 Technologies

The technological architecture of Horizon 3 is defined by the merging of adaptive intelligence and transparent orchestration. At the centre of this ecosystem lies a self-optimising AI orchestration system that continuously processes live operational data to dynamically reprioritise and resequence airside tasks. This adaptive logic enables the orchestration system to respond in real time to fluctuating operational conditions without the need for human recalibration. These self-adjusting capabilities form the core of system flexibility and ensure resilience in a complex and dynamic environment.

To maintain accountability and reinforce trust, all AI-led decisions are demonstrated through explainable reasoning dashboards. These interfaces provide human operators with a clear trace of the conditions, thresholds, and fallback logic used in every orchestration outcome. In the event of an escalation or anomaly, the dashboard enables users to inspect system logic in detail, often supported by natural language explanations and decision trees (Figure 24). This level of transparency allows human oversight to remain meaningful despite the automation of most routine processes.

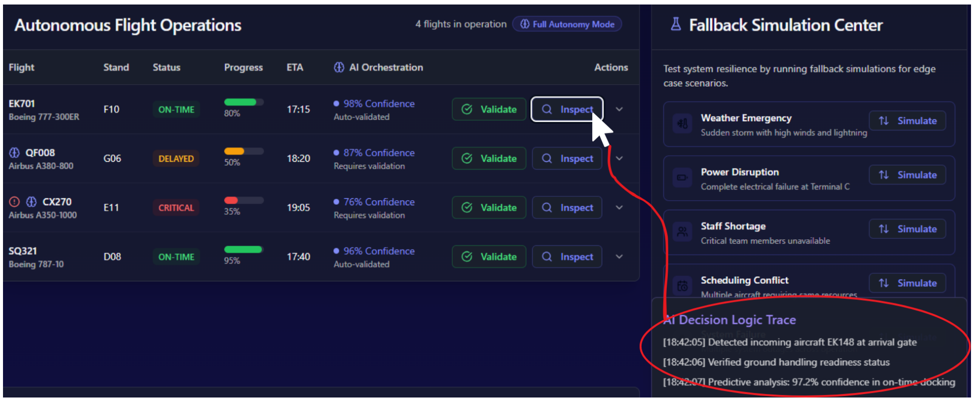

Figure 24: Inspecting functionality

A core enabler of reliability in this horizon is the implementation of dynamic fallback systems. These systems operate across multiple tiers of redundancy, ensuring that if an automated task fails the AI can trigger or suggest alternate procedures without external input. In critical cases where fallback logic is exhausted or the system's confidence drops below a defined threshold, escalation protocols alert the remote operator.This allows a human to take control if needed, ensuring safety and that operations can continue smoothly even when the AI can't handle a situation alone.

To support both strategic validation and system resilience, Schiphol deploys a live digital twin of the airside operation. This continuously updated simulation environment mirrors the real-world system using data streams from all autonomous assets, infrastructure, and external variables. The digital twin is an important component of simulation and testing protocols, allowing both AI and human operators to simulate different circumstances, run virtual trials of edge cases, and test new updates before deployment into live operations (See dashboard mock-up example in figure 25. This functionality ensures that improvements, rule updates, and operational learning can occur without disrupting actual performance.

All of these systems are supported by an integrated and secure digital infrastructure. High-speed, low-latency vehicle-to-everything (V2X) communication ensures seamless orchestration among autonomous agents, gates, and stands, while cybersecurity protocols protect the integrity of both control signals and audit trails [32]. Together, these technological components establish Horizon 3 as a resilient and intelligently adaptive orchestration environment, where trust is maintained not through control, but through visibility, traceability, and systemic robustness. See figure 26 for a full overview of the dashboard. For additional features, view Appendix H.

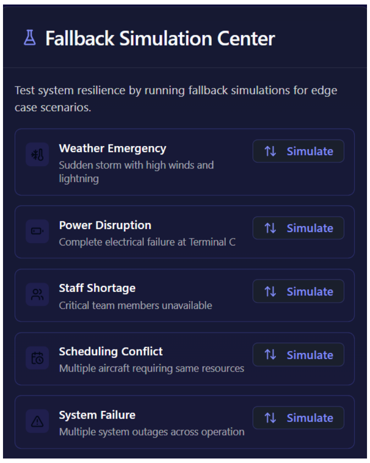

Figure 25: Simulation Tool

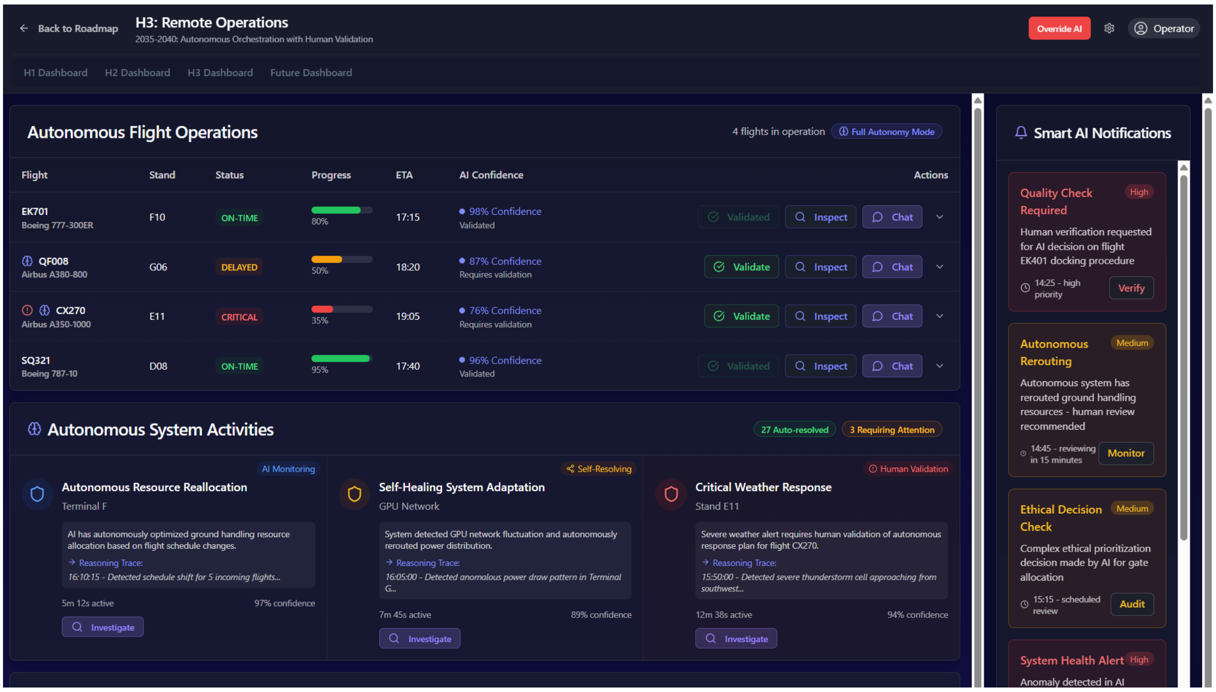

Figure 26: Dashboard Horizon 3

7.4.3 Ecosystem

By Horizon 3, the involved ecosystem at Schiphol has reached its most advanced and autonomous form. At the centre of this evolved structure is a highly integrated orchestration system, powered by AI, which assumes the primary role in planning and executing nearly all operational tasks. Human involvement, while still present, is completely redefined, shifting from day-to-day management to strategic supervision. The resulting configuration creates a highly efficient and resilient system where automation is the default, and humans serve as final fallback and oversight roles.

The ecosystem is now composed of three interdependent layers: AI-driven and autonomous processes, a reduced but strategically placed ground handling team, and a remote operator overseeing orchestration (Figure 27). The AI system assumes full control of routine operations, coordinating ground service vehicles, adjusting turnaround sequences, and responding autonomously to real-time inputs. This orchestration is not rigid; instead, it is dynamic and adaptive, with integrated fallback protocols that allow the system to self-correct or reroute when encountering disruptions. Vehicles and service units operate in sync through continuous communication with the orchestration core and with one another, ensuring seamless execution of tasks even under dynamic operational conditions.

Meanwhile, the role of the ground handling team has been significantly scaled down. Their operational presence on the apron is minimal and limited to functions that cannot yet be automated. They serve primarily as a local safety fallback layer, stepping in only when the autonomous system escalates a task that requires physical interaction or on-site validation. Importantly, they no longer initiate or manage operations themselves but operate in alignment with AI instructions.

Above these two layers sits the remote operator, who now serves as a strategic supervisor rather than a tactical controller. This individual oversees the orchestration logic through a transparent and interactive dashboard environment. When an anomaly is flagged or an automated decision requires validation, particularly in edge cases, the operator steps in to investigate or assess the reasoning provided by the AI, and either approve, override, or request further clarification. These actions are not reactive management tools but rather high-level governance interventions. The operator retains ultimate responsibility for overseeing fallback escalations and ensuring that system behaviour aligns with regulatory, operational, and ethical expectations.

The interactions within this ecosystem are designed to be both efficient and resilient. Most communication flows from the orchestration system outward, with autonomous units adjusting to AI commands and conditions. The operator is kept informed through transparent, explainable AI logs that display the rationale behind key decisions, confidence levels, and fallback logic in use. In turn, operators and ground staff can investigate or inspect system behaviour at any time but are only prompted to intervene when human input is necessary to resolve ambiguity or risk.

Figure 27 visually summarises this ecosystem. AI-driven processes operate autonomously as the system default, while both remote operators and the ground handling team stand in as fallback actors, engaged in the event of a system failure or rare exception. Arrows indicate the flow of validation, escalation, and support, revealing a robust and decentralised structure in which each layer supports the others while maintaining clear functional boundaries.

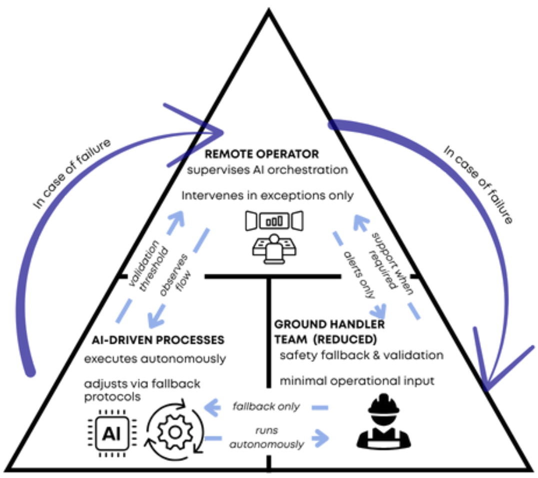

Figure 27: Ecosystem of stakeholders in H3

Through this structure, Horizon 3 demonstrates that autonomy does not imply isolation from human oversight but rather a redefinition of human roles toward strategic supervision, safety assurance, and system trust maintenance. This layered approach ensures that even in a highly automated environment, human accountability, transparency, and operational clarity are preserved.

7.4.4 KPIs

To assess the maturity and effectiveness of Horizon 3 implementation, a defined set of KPIs is established. These indicators not only function as an operational performance evaluation, but also to validate the reliability, safety, and stakeholder acceptance of the remote and resilient orchestration model.

A primary metric is the *percentage of operations handled autonomously*, which tracks the proportion of ground handling tasks fully managed by AI systems without human initiation or intervention. This indicator reflects the system's ability to execute sequences independently and serves as a standard for the degree of operational autonomy achieved.

Closely linked to system resilience is the *frequency of fallback protocol activation*. This KPI captures how often the orchestration system needs to engage backup processes, either through alternative logic paths or escalation to human operators. A low activation frequency would suggest system robustness and reliable performance, while a higher frequency may indicate ongoing edge case learning or areas needing further system development.

Trust and acceptance remain critical in highly autonomous environments. Accordingly, *public and regulatory trust indicators* are included to capture sentiment and confidence levels among key stakeholders. These may be obtained through periodic surveys or structured feedback from ground handlers, operators, and regulatory bodies. Metrics might include perceived system transparency, safety confidence, or readiness to expand autonomous functions.

Another essential metric is the *average time to anomaly resolution*. This measures the duration between the identification of an operational exception and its resolution, either autonomously or through human intervention. Shorter resolution times indicate the system's ability to manage disruptions efficiently, a core requirement in high-throughput airport environments.

Together, these KPIs support the continuous improvement of the Horizon 3 ecosystem. As automation scales, the expected impacts include a significant reduction in operational costs due to decreased manual labour, enhanced safety outcomes enabled by predictive anomaly detection, and improved resilience through integrated fallback logic. Furthermore, by enabling real-time optimisation of resources, the orchestration system promotes both economic and environmental sustainability.

Monitoring these KPIs over time will allow Schiphol to quantify progress toward its long-term vision of autonomous airside operations while maintaining transparency, trust, and regulatory alignment.

7.4.5 Purpose and Objectives

The primary purpose of Horizon 3 is to establish a resilient orchestration environment in which autonomous systems execute the majority of routine airside tasks, while human roles shift toward high-level oversight and exception management. This phase represents the establishment of technological maturity and organisational trust, where automation is no longer experimental but a stable, integrated layer of operations.

The overarching objective is to enable full AI-led sequencing and coordination of turnaround activities across the airside. These operations are to be handled with minimal delays and high reliability, driven by real-time data and predictive algorithms. To maintain human governance and transparency, operators must have access to intuitive dashboards that allow them to inspect, simulate, and override AI decisions when necessary. This ensures that human judgement remains integrated in the system, particularly for edge cases, regulatory compliance, or rare operational conditions.

A simultaneous objective is to protect resilience through robust fallback protocols. These mechanisms must allow operations to continue seamlessly during anomalies. Human operators must be able to intervene quickly and confidently when such cases arise, without compromising the flow of operations.

Together, these objectives define Horizon 3 as a phase of strategic automation, one in which system autonomy is matched with traceability, explainability, and trusted human oversight.

7.4.6 Rationale and Values

By the time Horizon 3 is reached, autonomous technologies at Schiphol have developed to a level of maturity where they are capable of independently managing the majority of routine airside operations. The orchestration system does more than automate, it learns, adapts, and optimises based on live operational feedback. This phase marks a crucial shift in the human-machine relationship: AI assumes the role of an intelligent executor, while humans transition into supervisory roles focused on validation, exception resolution, and ethical governance.

The rationale for this phase lies not only in achieving higher output and efficiency but in ensuring operational continuity and trust during disruption whilst further reducing physical labour reliance. Autonomous systems must demonstrate their ability to recover from failure, make decisions in ambiguous situations, and provide clear, explainable reasoning for their actions. This is essential for maintaining human confidence in the system, both internally among operators and externally among regulators, partners, and passengers.

A key value underlying this horizon is transparency. The orchestration logic must remain interpretable, with all decisions logged, traceable, and auditable. Explainable AI dashboards allow human supervisors to understand the decision pathways taken by the system, enabling real-time validation and post-event reviews.

Resilience is another defining value. The system's ability to detect anomalies, activate fallback protocols, and continue operating without interruption is central to its reliability. This reduces the need for reactive human intervention and ensures stable performance across varying conditions.

Efficiency is achieved not only through automation but through intelligent coordination, proactive maintenance, and dynamic scheduling, all driven by real-time data. At the same time, human value and purpose are maintained through strategic roles that focus on validation, learning, and system management.

Ultimately, Horizon 3 balances automation with accountability, offering a model of AI-led operations where performance, safety, and trust are mutually reinforced.

7.4.7 Actions to be taken

To ensure a successful implementation of Horizon 3, Schiphol must take goal-oriented steps to reinforce the safe, transparent, and resilient integration of autonomous orchestration. These actions must strengthen both the technical foundations and the human oversight frameworks that support this operational phase.

First, exception protocols must be fully clarified. While the AI system handles the majority of routine operations, it is essential to define precise conditions under which responsibility shifts back to human operators. These fallback mechanisms should include clear escalation thresholds, communication pathways, and procedural steps for validating, overriding, or pausing automated actions in real-time.

Secondly, the co-creation of standard operating procedures (SOPs) is critical. These SOPs should be developed collaboratively with all relevant stakeholders, including ground handling providers, safety regulators, and remote operators, to ensure alignment across technical, operational, and regulatory expectations. This participatory approach not only builds trust in the system but also helps to anticipate potential points of friction or misalignment during live operations.

Third, a comprehensive auditing framework must be established. All automated decisions, including fallback activations, routing changes, and task reassignments, should be logged and made accessible for review through explainable AI dashboards. These audit trails will be essential for continuous system improvement, regulatory compliance, and accountability in the event of incidents or anomalies.

To support this transition systematically, figure 28 provides a concise MoSCoW table linking the prioritised actions to each horizon. This table aligns urgency levels with their corresponding interventions and helps visualise how efforts compound over time. For a complete, multi-horizon version of this table, see Appendix J, which includes the full MoSCoW overview. Additionally, Appendix K details the strategic rationale and implementation logic behind each action area.

Horizon	Urgency	Priority	Action
H3	Must have	Delegated AI decision-making with human override thresholds	Define and enforce system autonomy boundaries across orchestration tasks
		Automated fallback and error recovery	Implement tiered fallback strategies that escalate only when needed
		Operator role transition to orchestration governance	Redefine KPIs and responsibilities to support supervisory orchestration
	Should have	Shadow-mode simulations to test AI reliability	Run full AI orchestration in parallel with human oversight and compare performance
		Transparent diagnostics for trust and accountability	Develop audit logs and explanation tools for AI decision-making
	Could have	Advanced metrics for orchestration effectiveness	Create analytic models to measure efficiency, stability, and error handling

Figure 28: H3 MoSCoW

Taken together, these actions support a controlled yet future-forward transition into a highly autonomous airside environment, where human expertise continues to play an important role in oversight, safety assurance, and ethical governance.

7.5 Future Vision

The Future Vision outlines Schiphol's ultimate operational goal: a fully autonomous, self-orchestrating airside environment governed by intelligent systems with minimal human intervention. Building on the foundational trust, collaborative orchestration, and resilient autonomy established in the previous horizons, this phase marks the completion of the transition toward machine-led operations. By 2050, all standard processes are autonomously executed and adaptively optimised in real time. Human operators no longer manage day-to-day activities but serve as strategic supervisors, regulators, and ethical stewards of the system. This vision reflects not only technological maturity but also institutional and societal readiness to trust critical airport operations to AI-driven systems.

By 2050, Schiphol will operate as a self-orchestrating, resilient airside ecosystem where intelligent systems execute all standard operations autonomously, while humans provide strategic oversight, regulatory governance, and ethical assurance, ensuring a future of safe, efficient, and sustainable airport operations.

7.5.1 Key Features

By the time Schiphol reaches its 2050 vision, the airside orchestration system will be defined by a set of fully autonomous capabilities that together enable safe, efficient, and adaptive operations with minimal human involvement. At the core of this transformation is the rise of *self-healing systems*, capable of detecting, diagnosing, and resolving their own anomalies. Whether triggered by a failing component, misalignment in scheduling, or environmental changes, the system is designed to automatically reconfigure itself, reroute assets, and restore optimal operation without human intervention.

To ensure robustness under all conditions, these systems are equipped with *multi-tiered fallback logic*. This logic is structured to address escalating degrees of disruption through layers of redundancy. These fallback strategies are integrated across all orchestration levels, ensuring continuity even when multiple failures or unforeseen scenarios occur.

Another distinct feature of the future vision is *predictive self-management*. Taking decades of historical and real-time operational data, the AI continuously anticipates potential risks and proactively adjusts scheduling, resource allocation, or activation sequences to mitigate them. This capability positions the system not as reactive, but as forward-looking and preventive in its logic, greatly reducing the occurrence of delays or failures.

Monitoring is fully *autonomous* at this stage. Instead of relying on live camera feeds or operator observation, systems continuously collect, process, and analyse data streams from all assets and environmental inputs. This allows the orchestration system to maintain complete situational awareness at all times, without the need for human oversight during standard operations.

Finally, the role of the human operator becomes largely dormant, referred to as the *Dormant Operator Mode*. Operators no longer engage with routine workflows but remain on standby, alerted only under exceptional or regulatory conditions. Their function shifts toward long-term oversight, governance, and system evaluation, ensuring human values remain integrated within autonomous processes. This model represents the peak of the human-AI transition, in which intelligent systems execute autonomously, yet remain accountable through human-designed supervisory frameworks.

7.5.2 Technology

At the core of Schiphol's autonomous future lies a self-healing, fully orchestrated intelligence framework. By 2050, all airside operations are autonomously sequenced, executed, monitored, and recovered through a highly integrated network of autonomous systems. The technological environment functions as a closed loop, constantly predicting conditions, executing plans, detecting deviations, and adapting in real-time (Figure 29).

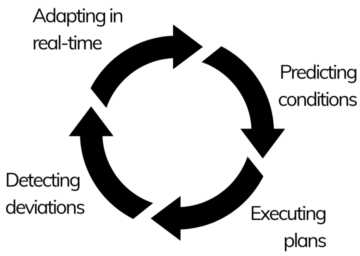

Figure 29: Closed Loop

Central to this environment are *self-healing AI systems*. These systems are capable of identifying and isolating disruptions the moment they arise, and independently deploying corrective measures. The AI adapts the operational plan accordingly without the need for manual intervention. This not only further speeds up resolution times but also significantly reduces error.

At the same time, *predictive recovery and multi-tier fallback systems* ensure that resilience is integrated throughout the orchestration logic. Rather than relying on a single recovery method, the AI employs layered response strategies, escalating through increasingly robust fallback scenarios as needed. The result is seamless operational continuity, even in the face of complex or a chain-reaction of anomalies.

All autonomous activity is traceable through a network of *real-time reporting and audit trails*. Every decision made by the orchestration system is logged, time-stamped, and stored with associated reasoning and data context. These logs are accessible via advanced dashboards that allow human supervisors, auditors, or regulators to inspect any decision on demand. This capability supports not only compliance with safety and transparency standards, but also provides a foundation for continuous performance analysis and refinement.

Finally, interaction with the system is made intuitive through a *natural language oversight interface*. Operators can communicate with the AI using voice or text prompts, like "Why was the GPU delayed at stand B07?" or "Show anomaly history for this shift." This interface reduces the cognitive load of oversight, allowing rare operator interventions to be conducted with maximum efficiency and minimal training. It ensures that even as the system operates independently, human understanding and engagement remain possible on demand.

Together, these technological advancements create an orchestration system that is not only autonomous but also adaptive, explainable, and accountable, delivering on the promise of a safe, resilient, and trustworthy airside operation. See figure 30 for a full overview of the dashboard. For additional features, view Appendix I.

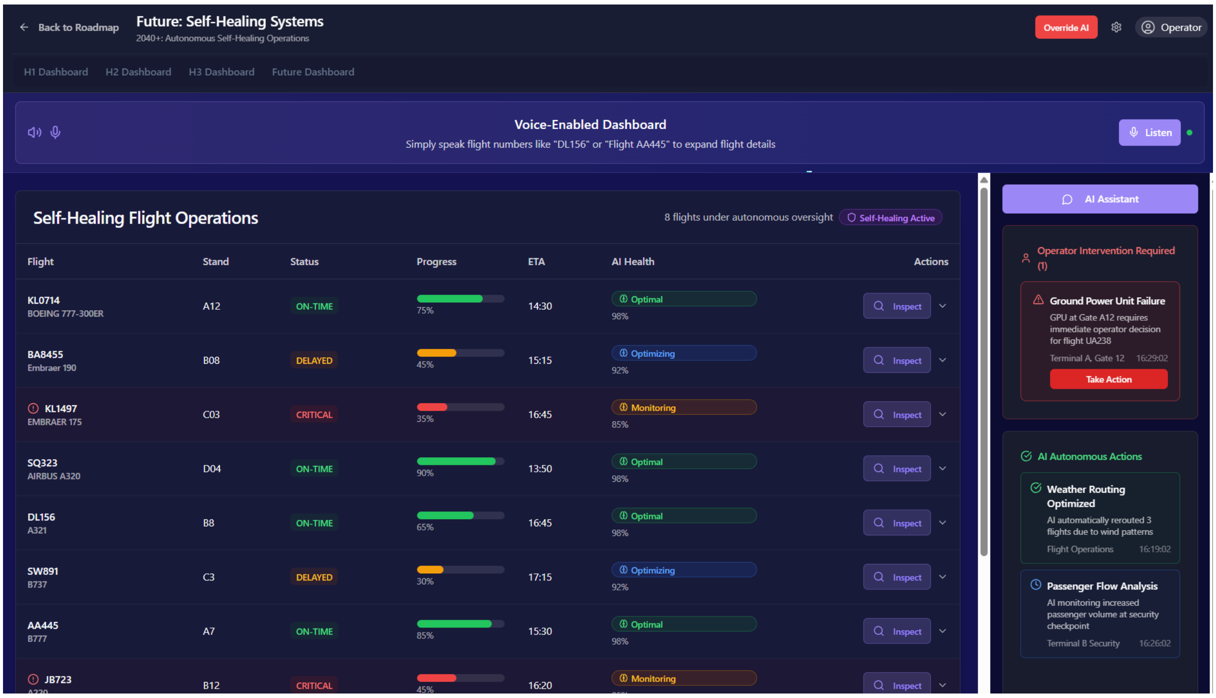

Figure 30: Dashboard Future vision

7.5.3 Ecosystem

By 2050, Schiphol's airside ecosystem functions as a fully autonomous orchestration environment, integrated within a thin supervisory layer of human oversight. The system is designed for almost-total independence, with intelligent tools carrying out all operational tasks, planning, execution, adjustment, and monitoring, while human actors are engaged only in exceptional cases that exceed system capabilities. This marks the completion of the transition from human-led coordination to AI-governed orchestration with human intervention as a last resort.

At the centre of this ecosystem is the *autonomous orchestration system*, which acts as the default operational 'operator'. It is responsible for managing all airside sequences in a self-regulating, data-driven manner. The AI platform operates independently, requiring no human input for routine tasks. Its logic is adaptive and context-aware, responding to live environmental data, performance history, and predicted disruptions. Each process is linked to the next through conditional logic, forming a self-activating chain of operations, such as the automatic triggering of FOD scanning, GPU connection, and baggage handling once aircraft docking is detected.

Oversight is still present, but reframed. The *remote operator* assumes a dormant role, providing passive supervision and intervening only when the AI system escalates a scenario beyond its confidence threshold. In such cases, the operator is presented with a summary of the incident, the reasoning behind the AI's assessment, and a menu of possible responses. Importantly, this interaction is exception-based and facilitated through intuitive dashboards and natural language interfaces. Routine visibility into the system's performance remains available but is rarely acted upon unless anomalies arise.

The *ground handling team* has been reduced to a small, highly trained unit that is only engaged under exceptional physical circumstances. These tasks are rare but critical, serving as a last-resort layer to ensure safety and regulatory compliance.

The structure of the ecosystem is defined by two principles: autonomous continuity and integrated resilience. All operations are linked in a logic-based orchestration chain, dynamically adjusting to traffic flow, environmental inputs, and real-time system states. Meanwhile, resilience is maintained through multi-tiered fallback logic, predictive analytics and handling, and redundant communication channels. The system is designed to self-correct before disruptions escalate, but retains the ability to obey human input when uncertainty or ambiguity arises.

Interactions between human and AI are thus minimal, but deeply structured. Communication is not ongoing, but triggered by need. When interaction occurs, it is transparent and bidirectional: operators may inspect the system using natural language commands, and the AI provides structured, explainable responses. This configuration maintains high system efficiency while preserving the integrity of human oversight. Simply put, when the autonomous processes fail, the multi-tiered fallback systems will attempt to solve. If the fallback systems and self-healing AI fails, the remote operator intervenes. The remote operator can activate the ground handling team as a last resort, as shown in Figure 31.

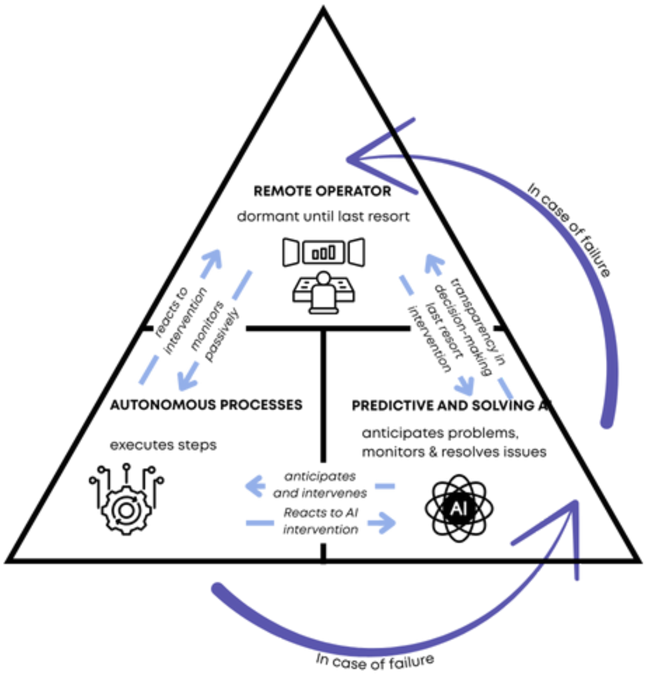

Figure 31: Ecosystem of stakeholders in the future vision phase

Ultimately, the future ecosystem at Schiphol is defined by its ability to operate autonomously under all standard conditions while preserving strategic human control in the background. This configuration enables scalability, system trust, and operational continuity, delivering on the vision of intelligent, safe, and sustainable airside operations.

7.5.4 KPIs

To assess the successful implementation of Schiphol's autonomous future vision, a series of KPIs must be monitored. These metrics must capture system resilience, trustworthiness, and strategic impact.

A primary indicator is the Autonomous Operations Rate, which measures the percentage of airside operations executed end-to-end without human intervention. This metric directly reflects the system's maturity and the extent to which the orchestration logic can reliably and independently manage elaborate workflows under varied conditions.

The System Trust Index assesses stakeholder confidence in autonomous processes. This index may be gathered through periodic surveys, feedback loops, and qualitative assessments, this indicator captures perceived transparency, reliability, and safety of the AI-led environment.

A key resilience metric is the Exception Rate, the frequency with which the AI system must escalate a scenario for human intervention. A low exception rate suggests high confidence in system decisions and robust performance across routine and non-routine events.

Complementing this is the Self-Healing Success Rate, which measures the proportion of system anomalies resolved autonomously through internal logic or fallback mechanisms. This KPI reflects the orchestration system's ability to adaptively manage disruptions without delaying operations or requiring external input.

Together, these KPIs provide a comprehensive picture of Future Vision effectiveness. When successfully realised, the expected impacts include maximum operational efficiency with minimal human oversight, full transparency through real-time audit trails, close-to-zero disruptions through predictive and preventive logic, and a substantial reduction in human worker reliance. Monitoring these indicators ensures that autonomy is not only achieved, but remains trusted, accountable, and future-proof.

7.5.5 Purpose and Objectives

The overarching purpose of the Future Vision is to realise a fully autonomous airside operation in which all ground handling processes are self-orchestrated by intelligent systems. In this state, AI governs the complete operational chain, from planning and execution to monitoring and real-time adaptation, while human involvement is limited to rare exception scenarios and high-level strategic oversight. The system is designed to be not only autonomous, but also self-healing, adaptive, and inherently transparent.

To achieve this, several key objectives must be met. First, a fully integrated orchestration framework must be deployed, enabling AI to autonomously coordinate all airside operations with minimal delay and maximal efficiency. This includes seamless task sequencing, vehicle activation, service coordination, and schedule optimisation under dynamic real-world conditions.

Second, the system must establish regulatory certification as a self-healing platform. This means demonstrating their reliable fallback mechanisms, predictive recovery protocols, and the capacity to maintain safe operations under stress without human intervention. Such validation is critical for public trust and regulatory operational approval.

Finally, it is essential to maintain public and organisational trust through ongoing transparency and accountability. This requires the implementation of comprehensive audit trails, explainable AI logic pathways, and compliance to high ethical standards. The system must be designed not only to perform independently, but to do so in a way that is fully traceable, fair, and aligned with stakeholder expectations.

Together, these objectives define a new operational setting, one in which the airport functions as a self-regulating, intelligent ecosystem, governed by autonomous systems yet guided by human values.

7.5.6 Rationale and Values

The Future Vision represents the final outcome of Schiphol's long-term roadmap, a fully autonomous orchestration system capable of managing the complexity and scale of airside operations with precision, reliability, and adaptability. At this stage, intelligent automation has replaced human execution across all standard processes, enabling operations that are not only efficient and scalable, but also environmentally sustainable and inherently resilient. Human roles have transitioned from active coordination to passive management, overseeing the system through exception-based logic, ethical governance, and strategic validation.

The rationale for this phase lies in the systemic benefits of end-to-end autonomy. Autonomous orchestration reduces delay chain-reactions by removing latency in decision-making and execution. It lowers emissions through precise task sequencing, idle time elimination, and proactive maintenance scheduling. Labour reliance is significantly decreased, allowing (limited available) human resources to shift toward high-value oversight and compliance functions. At the same time, system safety is maintained through predictive diagnostics, layered fallback protocols, and traceable decision logic.

Together, these values and interventions affirm the strategic outcome of the roadmap: a resilient, scalable, and intelligent airport system that aligns automation with transparency, efficiency with ethics, and autonomy with accountability.

7.5.7 Actions to be taken

The success of this future model depends not only on technical autonomy capabilities, but on persistent and verifiable system visibility. Trust in AI-led airport operations must be continuously upheld through transparent logic structures, comprehensive audit trails, and intuitive oversight interfaces. Even in a dormant role, human operators must retain the ability to inspect, simulate, or override any system decision. In this context, visibility overrides direct authority as the primary mechanism of trust. It ensures that at all times, the operator can understand, verify, and, if needed, intervene in the orchestration process.

To realise this vision, several interdependent technological and organisational developments are required. A fully autonomous orchestration platform must be implemented, capable of handling predictive planning, adaptive recovery, and complete end-to-end process management without external prompts. This must be supported by advanced audit interfaces that allow dormant human oversight through natural language or gesture-based inputs. These interfaces ensure that the rare moments of operator intervention are intuitive, efficient, and strategically guided. Furthermore, the infrastructure must be inherently self-healing with multi-tier fallback protocols that can solve disruptions, reroute processes, and prevent chain reaction failures in real time. Alongside this, a robust regulatory and ethical review framework must be established. This framework will govern the evolution of AI-led operations and ensure their ongoing legitimacy, safety, and public accountability.

Beyond these initial deployments, the Future Vision phase also demands long-term resilience and continuous system refinement. Regulatory compliance must be actively maintained, with autonomous systems continually aligned to evolving aviation standards and subjected to auditing and certification procedures. Machine learning and AI-training must be an ongoing process, using operational data to refine decision-making accuracy, improve prediction models, and adapt orchestration strategies over time. Finally, exception handling must be systematically enhanced through simulation environments and real-world feedback loops, improving the system's ability to detect and resolve edge cases with minimal disruption.

To support this transition in a structured manner, Figure 32 presents a concise MoSCoW overview, mapping the prioritised actions to their respective horizons. It highlights the urgency and impact of each intervention and illustrates how these efforts build progressively over time. For a complete cross-horizon breakdown, refer to Appendix J, which provides the full MoSCoW framework. Further strategic rationale and implementation logic can be found in Appendix K.

Horizon	Urgency	Priority	Action
Future Vision	**Must have**	Self-healing orchestration with minimal human input	Develop fully autonomous orchestration logic with fallback AI and anomaly handling
	Should have	Operator shift to 'governer'	Define training, audit and oversight roles for exception-based intervention
	Could have	Adaptive orchestration ethics framework	Establish ethical guidelines for decision transparency and accountability

Figure 32: Future Vision MoSCoW

These actions are continuous commitments. They form the foundation for Schiphol's ambition to operate an ethical, resilient, and high-performing autonomous airport, where operational excellence is no longer reliant on manual intervention, but on intelligence, adaptability, and embedded transparency.

7.6 Conclusion

The strategic roadmap outlined in this chapter presents a structured evolution from observational, human-led operations to a future of fully autonomous orchestration. Each horizon: *Foundation*, *Collaboration*, and *Autonomy*, builds upon the capabilities and trust established in the previous phase, allowing for a phased integration of technology, roles, and resilience. Rather than a purely technological progression, this transformation is tied to cultural, procedural, and ethical development. As illustrated in Figure 33, key initiatives across each horizon align with overarching strategic objectives, ensuring the roadmap remains coherent, actionable, and stakeholder-driven. The resulting vision is one of transparent intelligence, automation without worker alienation, and resilience grounded in transparency, positioning Schiphol as a frontrunner in ethical, AI-enabled airport operations.

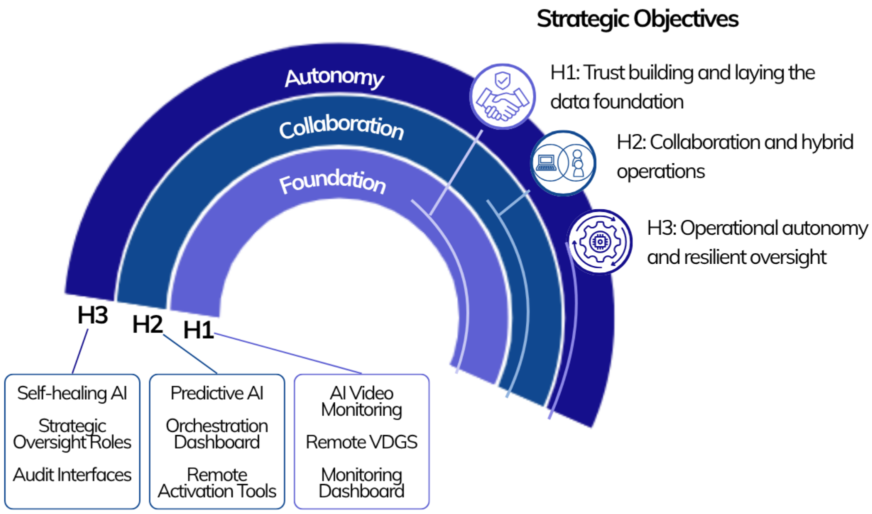

Figure 33: Progressive Arc of Strategic Horizons and Objectives

8 Discussion

8.1 Interpretation of Results

8.1.1 Generalisation and Maturity Framework

A key insight is that progress toward autonomy must align technological, organisational, and ecosystem maturity, a principle demonstrated by the three-level maturity framework of Thomson et al. (2022), seen in figure 34. This framework defines three stages of autonomous solution maturity: Level 1 (Assistance), Level 2 (Partial Automation), and Level 3 (Full Autonomy), and, describes how each level requires corresponding advances in technology, business/organization, and ecosystem integration [75]. For example, at Level 1 (Operator Assistance), firms deploy basic assistive technologies while keeping traditional business models and minimal ecosystem change. At Level 2 (Partial Automation), systems begin to handle specific tasks independently under human supervision, requiring more advanced integration of digital tools, evolving business processes, and early forms of inter-organisational coordination. By Level 3 (Fully Autonomous Operation), technology is highly advanced and self-governing, business models shift to outcome-based or service contracts, and the ecosystem becomes a collaborative network with standard interfaces and shared data and risk-reward structures. In short, Thomson's study emphasises that moving beyond isolated "islands of autonomy" requires progressive evolution of technical capabilities, internal processes, and external partnerships. An autonomy initiative will delay or struggle if any one of these dimensions delay, a point that Schiphol reinforces [76].

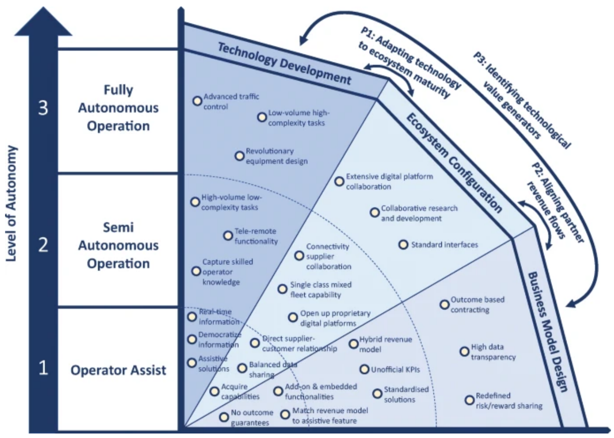

Figure 34: Maturity Framework

The phased progression at Schiphol aligns with this maturity framework. Early roadmap phases focused on operator support correspond to Level 1: introducing technology that assists workers without major workflow changes. Next phases introduce increasingly more semi-autonomous vehicles and remote operations, reflecting Level 2: technology takes over routine tasks under human supervision, prompting more significant process changes and new skill requirements, and requiring increased integration standards. The final envisioned phase of full autonomy aligns to Level 3 maturity: vehicles and ground handling systems operate with minimal human intervention, which demands not only the technical capability but also new organisational roles and ecosystem-wide arrangements (regulations, liability sharing, new business models with airlines and handlers). Importantly, the Schiphol roadmap was designed with holistic transitions in mind, it pairs each technology step with training, change management, and stakeholder agreements so that all elements of the system progress in their maturity together. This approach makes the roadmap adaptable beyond the specific scope of docking operations at Schiphol.

Accordingly, it is crucial to acknowledge that technological feasibility alone is not enough to progress to the next stage of autonomy. The organisational and ecosystem context must be ready as well. Studies of past innovations show that neglecting these aspects leads to pilot projects that never scale, advanced 'high-tech' tools end up as fragmented unused solutions because the company wasn't prepared to actually integrate them [75]. Broad industry evidence shows that transformations fail if people and processes are not ready. A significant amount of organisations fail in digital transformation because they hadn't sufficiently addressed organisational change and behaviour, despite deploying new technology [63]. Accordingly, successful autonomy scaling and implementation demands a socio-technical approach, aligning technology upgrades with workforce training, organisational updates, and ecosystem governance.

Furthermore, transitions to autonomy must proactively manage resistance to change. Even if a roadmap is technically and logistically sound, human factors can pose a bottleneck. Literature on technology adoption identifies several sources of resistance that planners should account for. Joshi's equity theory offers one perspective: users evaluate changes based on fairness and personal impact, and they will resist an innovation if they perceive an inequitable outcome for themselves [14]. This means ground crew and other stakeholders may ask, "Do the new autonomous systems benefit me or just others? Am I losing out, while someone else gains?" If the introduction of autonomous processes or AI is seen as unfair (threat of job loss, disregard to current practice), morale and cooperation will suffer. Markus's theory further argues that resistance often arises not from technology alone, but from how the technology changes power and roles in an organization [50]. A system that redistributes decision-making or undermines established expertise can provoke pushback from those who feel their status or control is being disregarded or diminished. In the case of Schiphol, these theories imply that each autonomy milestone must be introduced with attention to perceived fairness and power dynamics. Change management strategies are essential: e.g. involving employees in design (to give them voice and ownership), clearly communicating how roles will evolve (to avoid status ambiguity), and ensuring the benefits of automation (safety, reduced physical labour, upskilled roles) are shared and understood. The stakeholder co-development approach was one way to mitigate resistance, by inviting input from ground handlers, operations managers, and regulators in shaping the new automated circumstances, the process acknowledged their concerns and ideas, making the eventual changes more easy to accept. This aligns with best practices suggested by the literature, when users feel a change is for them rather than done to them, they are less likely to feel it is inequitable or disempowering [14].

While Thomson's framework offers a structured approach to evaluate maturity, its greatest asset in this case is not for generalisation across industries, but rather for generalising within the larger Schiphol ecosystem itself. The maturity-based logic developed for docking operations represents a small aspect of a broader transformation.

Although the roadmap in this study is scoped around docking operations, this function is integrated within the wider turnaround process, which itself is integrated with numerous other airside operations. Each of these domains follows its own automation trajectory, but all are increasingly shaped by shared dependencies: AI-based orchestration, remote activation systems, escalation protocols, workforce adaptation, and data infrastructure.

Accordingly, the maturity-based approach piloted here can be generalised across Schiphol's wider operational environment. It offers a systemic template that encourages readiness-based progression, phased role evolution, and cross-functional integration. This approach aligns with Schiphol's ambition of developing a cohesive, orchestrated airside ecosystem, where each domain matures not in isolation, but in strategic synchronisation.

This concept is illustrated in Figure 35. Each operational domain maintains its own progression toward autonomy. However, they converge through shared enablers: AI monitoring systems, operator training, and escalation protocols.

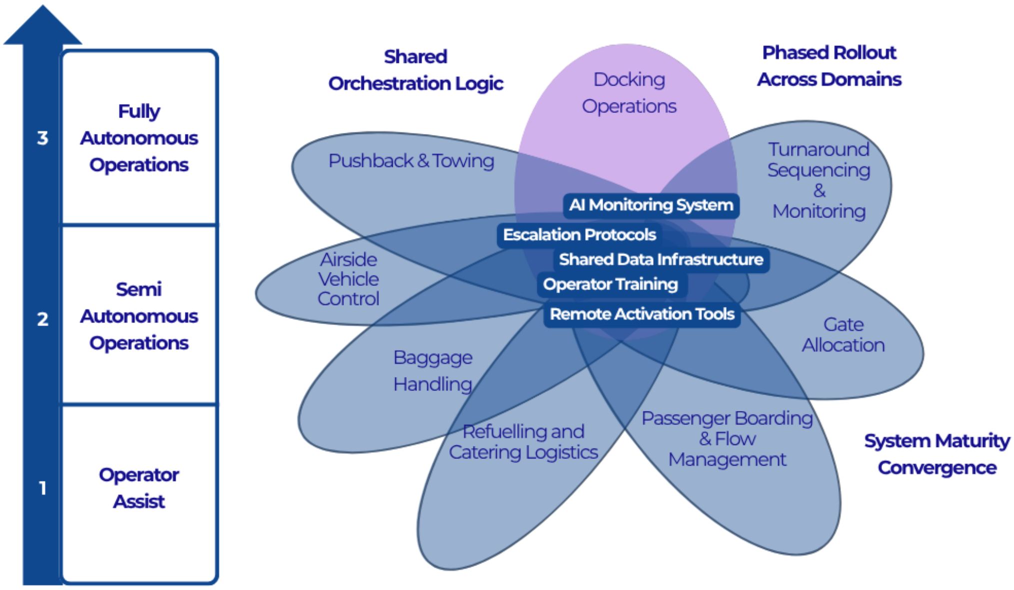

Figure 35: System Overlap in Interconnected Airside Systems

Notably, partial overlaps between domains allow shared capabilities to be reused or scaled across contexts. This results in a form of interconnected automation, where local innovations in one function may guide orchestration logic in others.

Accordingly, the roadmap generated for the orchestration of docking operations is more than a functional, specific guide, it stands as a basis for how maturity convergence can develop across Schiphol. In this view, autonomy is not a product of isolated upgrades, but of interconnected transformation across operational layers. The roadmap's relevance lies not just in automation, but in how it facilitates alignment across a intricate, dynamic, and increasingly intelligent ecosystem.

8.1.2 Non-Linear Progression and AI acceleration

A key insight emerging from both the roadmap and broader autonomous technology literature is that the evolution of system components is not inherently linear. Specifically, certain components, like AI-based orchestration, predictive analytics, and real-time monitoring, are progressing at a much faster rate than physical systems like autonomous docking tools [59, 15]. This concept causes a mismatch between what is technically possible and what is operationally feasible. It complicates the common assumption that sensing, vehicle control, and orchestration capabilities will all evolve at the same pace.

This is relevant in the context of Schiphol's airside, where the orchestration centre of the system, the top orchestration layer will reach deployment maturity significantly earlier than the autonomous on-site components. As highlighted by progress in airport remote tower operations, AI-based perception and decision-support systems are already enabling remote control and situational awareness without full autonomy of the physical systems they support [73]. Similarly, advanced smart monitoring and remote activation tools for apron operations can be developed and deployed quickly, whereas certifying and scaling up fully autonomous vehicles will likely require years due to regulatory, safety, and physical integration constraints.

While the orchestration AI may be technically ready, it should not be prematurely deployed in full autonomy mode. According to the maturity framework, a misalignment between technology readiness and ecosystem maturity can lead to failure. Thus, fast-moving technologies must sometimes "wait" for business processes and organisational readiness to catch up. In practice, this means AI orchestration should be introduced as a decision-support tool first, making recommendations that are verified and executed by humans, before being allowed to act autonomously.

In summary, AI orchestration capabilities will likely outpace physical autonomy systems, but the roadmap enforces a linear sequence for good reason. The goal is not to delay AI use, but to ensure that its deployment aligns with operational reality, legal structures, and human acceptance. While the roadmap appears linear, its logic is maturity-driven. Orchestration systems must adjust their pace to the physical readiness of the docking process and its implications (figure 36). Rather than viewing this as a constraint, it should be seen as a design principle: deploy cognitive autonomy at the same pace with social and procedural maturity, even if this means deferring full technical potential until the system as a whole is ready.

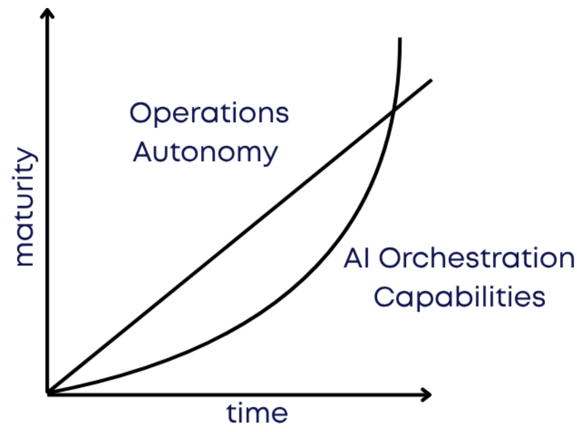

Figure 36: Maturity overview

8.1.3 Timeframes and Monitoring Alignment

Given the non-linear nature of AI advancement, the roadmap's linear timeframes may appear too structured. However, the roadmap follows a linear sequence because orchestration maturity must evolve together with the physical realities of airside operations. Autonomous decision-making in such safety-critical contexts cannot be accelerated recklessly, the orchestration system must be phased in carefully and aligned with both the technological maturity of the autonomous subsystems and the organisational and ecosystem capacity to manage and absorb this change.

AI-based monitoring systems offer a way to bridge this gap. These systems can be deployed early to provide oversight, risk detection, and system-level transparency even before the physical automation of assets is complete. This provides a protective layer during the transition, improving safety and building trust while full autonomy is still under development.

This also justifies the sequencing of the roadmap phases. For instance, in Horizon 1, human operators are supported by early-stage AI monitoring systems. In Horizon 2, the AI layer becomes more proactive, recommending optimised task sequences, flagging conflicts, or redistributing tasks across semi-autonomous assets. Only in Horizon 3 do orchestration systems act with minimal human input, once trust has been earned and the ecosystem is prepared. Contrarily, ecosystem readiness is neither necessarily a linear process. As concluded in the Enablers and Blockers analysis in figure 14, adoption challenges are not static. Resistance can transform into familiarity and embracement once users begin to experience the reliability and usefulness of AI systems in practice. This point of inflection, where scepticism turns into familiarity, can accelerate the acceptance of cognitive orchestartion systems and reinforce the maturity-asligned logic of the roadmap.

Accordingly, the timeframes in the roadmap are not determined milestones but readiness thresholds. Since readiness and maturity is usually not linear, and Thomson's framework highlights, deployment should be governed by maturity convergence: technology, organisation, and ecosystem [76]. As such, Phase 2 might be accelerated if AI oversight proves effective and stakeholders are aligned, or delayed if resistance, regulatory gaps, or process mismatches arise. This means that monitoring maturity should be the primary trigger for phase advancement.

Ultimately, while AI monitoring may mature ahead of autonomous operations, it serves as a precondition, not a replacement. Its role is not to accelerate the roadmap recklessly, but to protect and facilitate the safe progression of autonomy, ensuring each step is informed, trusted, and integrated in its organisational context. This reinforces the central insight from the literature and context research in which autonomy requires a synchronised evolution of systems, people, and organisations, not just fast technology.

8.1.4 Flexible Trajectories: Dynamic Over Static

Building on the previous discussion of orchestration maturity, AI acceleration, and readiness-aligned timeframes, it becomes clear that Schiphol's transition roadmap should not be interpreted as a fixed sequence of actions. Instead, it represents a dynamic structure, where progress is determined not by strict timelines, but by the convergence of system maturity across technological, organisational, and ecosystem domains.The influence of these dynamics is not fixed, the roadmap's development trajectory itself may shift over time depending on which forces dominate. The three figures below demonstrate how this non-linearity can turn out in different trajectories:

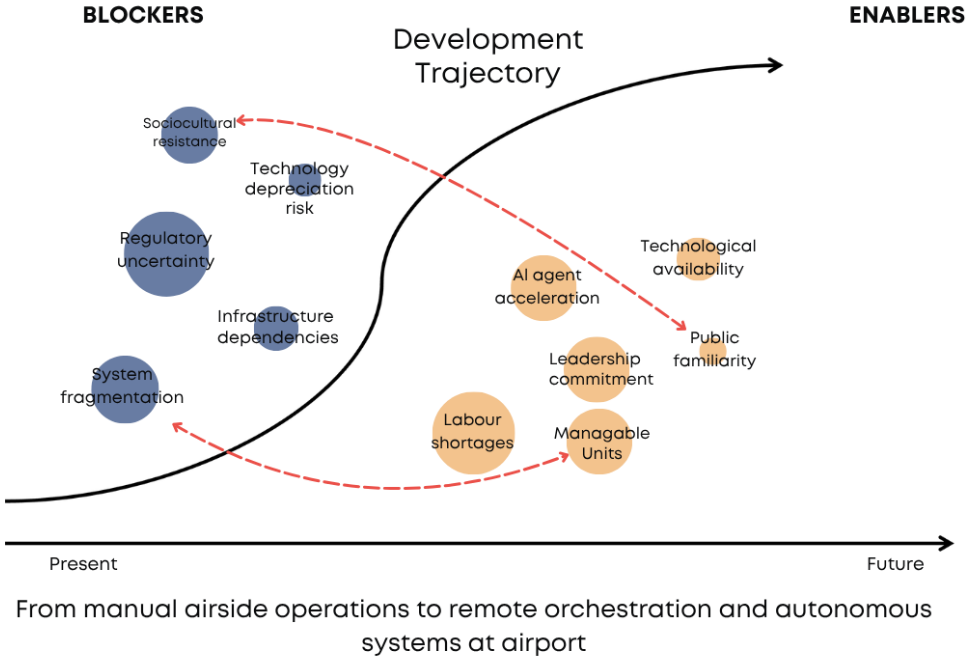

Figure 37: Baseline Trajectory: Standard S-curve showing balance between enablers and blockers

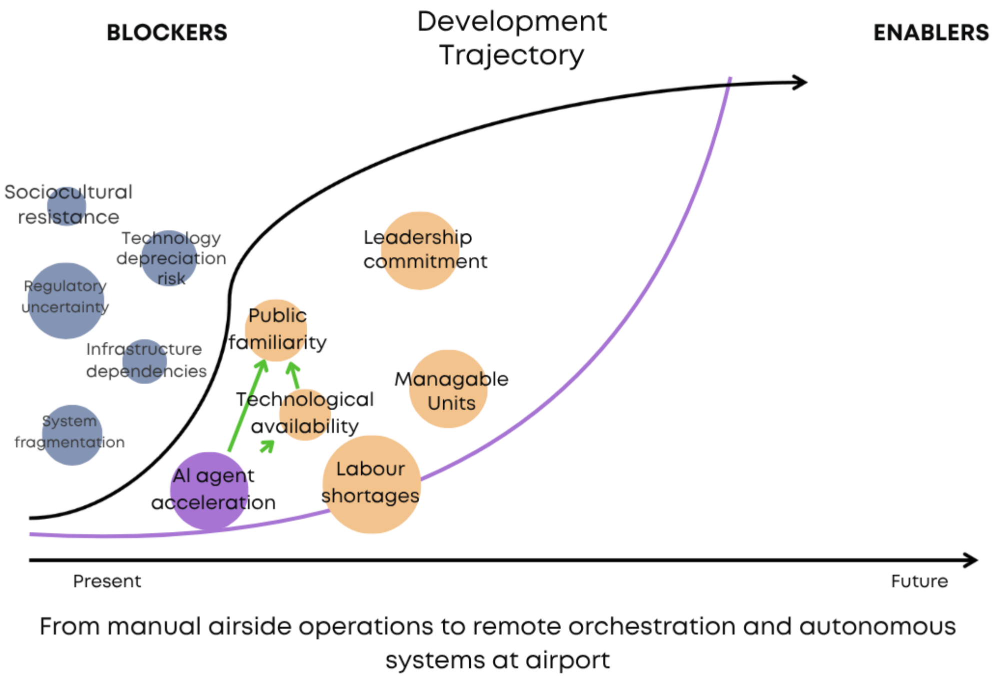

Figure 38: Accelerated Trajectory: Earlier inflection due to AI agent acceleration causing Technological Availability, leading to accelerated Public Familiarity

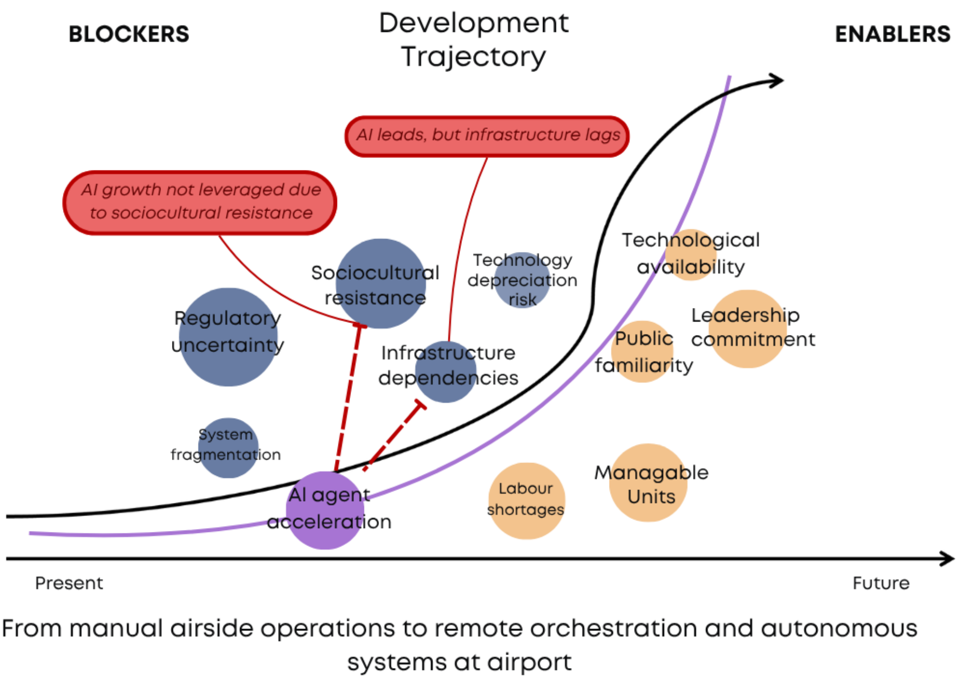

Figure 39: Delayed Trajectory: Friction emerges when AI agent acceleration outpaces infrastructure dependencies and sociocultural resistance

These scenarios illustrate that the trajectory toward autonomous operations is not predetermined. In some cases, strong reinforcement between enablers, such as AI agent acceleration, public familiarity, and technological availability, can create inflection points earlier than planned, enabling rapid acceleration. In others, bottlenecks such as regulatory uncertainty, sociocultural resistance, or infrastructure dependencies may slow adoption, even when AI capabilities are technically mature.

Therefore, the roadmap must be viewed as a dynamic, readiness-based structure rather than a fixed or linear sequence. These trajectory shifts highlight the system's sensitive nature to multi-directional forces, where enablers and blockers interact in fluctuating and sometimes amplifying ways. What appears as a smooth transition in planning may, in practice, be marked by acceleration windows, stagnation periods, or re-alignment efforts. This reinforces the importance of non-linear orchestration logic, one that actively monitors and responds to changing readiness across technological, organisational, and human domains. As also demonstrated in the maturity framework, maintaining alignment across these layers is crucial to avoid systemic mismatches. Ultimately, the roadmap should not be interpreted as a plan of milestones, but as an adaptive orchestration strategy that evolves through feedback, mitigation, and opportunity.

8.2 Limitations

While the roadmap for autonomous operations at Schiphol offers a structured and stakeholder-informed path, several limitations must be acknowledged regarding its scope, methodology, and assumptions.

8.2.1 Scope and Context

The roadmap is designed for Schiphol's unique operational environment, including its infrastructure, stakeholder landscape, and current plans. As such, while its phased logic and maturity framework are generalisable, the specific sequencing, technology selections, and timelines may not directly apply to other contexts. External conditions such as workforce setup, climate, goals or available capital may require significant adaptation. Though principles like human–AI collaboration and staged integration remain broadly useful, applications should adjust the roadmap to local operational context.

More specifically, this study focused on ramp and docking operations. It does not explore integration with terminal logistics, passenger handling, or en-route traffic systems. Full airport-wide autonomy will require alignment across all these layers, which remains a topic for future research and roadmap extension. From a single perspective this proves a limitations; however, drawing a parallel to the enablers and blockers analysis, this may also prove to be an enabler. Similar to the point of fragmentation, which is mostly seen as a blocker. Conversely, fragmentation as well as the scope of this project, allows for manageable, overseeable interventions. Smaller, self-contained domains act as proving grounds that reduce implementation risk and allow for controlled scaling toward wider transformation.

In addition, the sustainability dimension introduces a contradicting challenge. Although the roadmap supports Schiphol's ambitions by enabling shorter taxi times and reducing idle engine emissions, it also carries potential backlash. As highlighted in the futures wheel analysis, gains in operational efficiency could indirectly support higher flight throughput, which,if not regulated, might counteract carbon reduction goals. In addition, automation of the orchestration system and airside operations and excessive use of AI will result in new sources of significant emission. This underscores the importance of aligning autonomy initiatives with broader environmental policy, including emissions caps or offset strategies.

Ethical foresight is another area where the roadmap requires further elaboration. The futures wheel anticipates longer-term dilemmas related to labour transformation and system accountability. While the roadmap includes provisions for stakeholder co-development and upskilling, it does not yet detail governance mechanisms for managing job displacement, workforce equity, or liability in the event of AI failure. As automation expands in safety-critical environments, structured oversight will be necessary, particularly to address accountability gaps in autonomous decision-making and to ensure that the redistribution of roles is both fair and transparent.

8.2.2 Methodological and Technical Uncertainty

The roadmap was developed as a high-level strategic plan, informed by context research, literature, and stakeholder input, but currently lacks low-level technical trials or real-world pilots. This limits validation of assumptions related to system interoperability, or AI decision accuracy under real operating conditions. Empirical testing will be crucial to verify feasibility and surface integration challenges not evident in conceptual design.

Additionally, the timeframes used rely on projected developments in AI, autonomous systems, and supporting infrastructure. However, as past experience with self-driving technologies has shown, technological progress is often non-linear and unpredictable. Some capabilities may mature faster than expected (e.g., AI for coordination), while others (e.g., safe full process autonomy) may lag behind. Accordingly, the timeline should be viewed as indication rather than prescriptive, and adaptive mechanisms should be integrated to adjust based on actual readiness.

8.2.3 Human and Organisational Factors

While the roadmap integrates socio-technical considerations and stakeholder co-development, it does not detail operational strategies for change management, workforce training, or organisational transformation. Resistance to change, role uncertainty, and perceived unfairness remain potential barriers to adoption. Literature on adoption and resistance emphasises that even well-designed systems can fail without adequate user alignment [14, 50]. Future implementations will need to translate these insights into concrete programs prior to deployment.

8.3 Recommendations

Drawing on the findings of this study and the limitations outlined above, several strategic recommendations can be proposed for both practitioners and researchers working on autonomy transitions in aviation and comparable domains.

8.3.1 Adopt a Maturity-Layered Deployment Strategy

Organisations should treat technology readiness, organisational preparedness, and ecosystem coordination as equally necessary conditions for deploying autonomous solutions. The roadmap illustrates that even when AI orchestration tools are technically mature, their full deployment should be organised by social and operational maturity, in line with the Thomson et al. framework [76]. This means using AI first in decision-support roles, then gradually expanding autonomy as stakeholder trust, roles, and policies evolve in tandem.

8.3.2 Different Ownership Scenario's

As Schiphol prepares to scale autonomous airside operations, it must evaluate how ownership of orchestration tools and remote operations is distributed across stakeholders. The question of who owns the tools, data, and operational responsibilities has strategic implications for adoption, scalability, and system integration. Figure 40 illustrates three potential ownership pathways. Each presents a different balance of operational control and infrastructural responsibility:

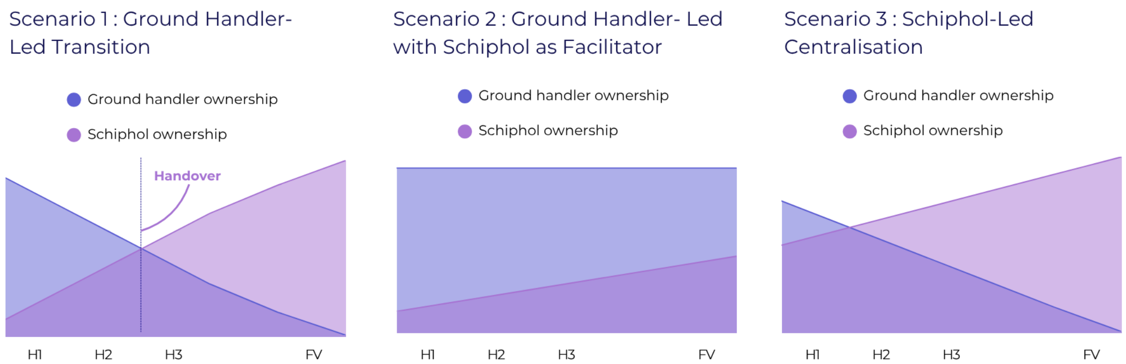

Figure 40: Three Different Ownership Scenarios

Scenario 1: Ground Handler-Led Transition

Ground handlers initially own and operate remote systems, with Schiphol facilitating the orchestration tools, before gradually transitioning control to Schiphol over time.

- *Ownership:* Dashboard tools and remote operations are facilitated by Schiphol but operated by ground handlers (each with their own remote operator).
- *Transition:* Decentralised start with a gradual shift toward Schiphol ownership.
- *Future Plan:* Remote operator roles may shift to Schiphol as automation consolidates.

Benefits: This scenario facilitates strong ecosystem buy-in by giving ground handlers early ownership of operations, which increases trust and willingness to adopt the system. Their familiarity with operational workflows improves initial tool usage and feedback quality. It also encourages rapid skill development and internal adaptation within handler organisations. By sharing early implementation responsibilities, the pressure on Schiphol is reduced during implementation.

Disadvantages: However, distributing control among multiple handlers increases the risk of fragmentation in tool usage and data interpretation. It also limits Schiphol's ability to establish unified, airport-wide oversight early on. Over time, transitioning operational ownership to Schiphol introduces additional complexity.

Scenario 2 : Ground Handler- Led with Schiphol as Facilitator Ground handlers retain operational ownership while Schiphol provides shared digital infrastructure and coordination tools, enabling a distributed yet connected system.

- *Ownership:* Schiphol supplies orchestration infrastructure; ground handlers control operations.
- *Transition:* Maintains decentralised operations throughout, with optional standardisation pathways.
- *Future Plan:* Long-term decentralised governance with evolving coordination protocols.

Benefits: In this model, ground handlers retain full operational control while Schiphol provides the digital infrastructure, enabling low-friction adoption with minimal disruption. This setup leverages the expertise and autonomy of each handler, supporting faster rollout and early value generation. A clear division of roles allows Schiphol to focus on infrastructure, while handlers concentrate on operational execution.

Disadvantages: At the same time, the lack of full Schiphol control makes airport-wide orchestration more difficult to achieve. Differences in standards and tool usage across handlers can lead to fragmented data, undermining the value of system-wide analytics. Moreover, the model relies heavily on consistent cooperation and may slow down strategic alignment as automation efforts mature.

Scenario 3 : Schiphol-Led Centralisation Schiphol owns and operates the orchestration tools and remote operations from the outset, integrating ground handler workflows under a centralised control model.

- *Ownership:* Full Schiphol ownership of tools and (remote) operations.
- *Transition:* Centralised from day one, with process alignment support.
- *Future Plan:* Fully Schiphol-operated autonomy layer.

Benefits: With Schiphol fully in control (as far as autonomous processes and remote intervention goes) from the start, this scenario enables early implementation of centralised oversight and consistent data-driven decision-making. It simplifies governance by clearly defining responsibilities and ensures all systems follow standardised protocols. This model is also well-aligned with Schiphol's long-term vision for scalable, integrated autonomous operations.

Disadvantages: On the downside, this approach is most disruptive to the current workflow and ground handlers may resist adopting externally imposed tools, potentially creating friction. Significant effort is required to align existing workflows, making coordination and change management more demanding. In early phases, Schiphol may also face gaps in operational knowledge that limit the effectiveness of system deployment.

The recommendation is to adopt a phased, adaptive ownership model, guided by orchestration maturity and stakeholder readiness, rather than committing to a rigid governance structure upfront.

8.3.3 Leverage AI Acceleration Safely and Strategically

Rather than resisting the non-linear acceleration of AI capabilities, organisations should proactively integrate these tools in early phases, especially for oversight, risk detection, and system diagnostics. Monitoring maturity can become a leading indicator of readiness, helping ensure that transitions are evidence-based and safe. However, AI deployment should remain synchronised with governance, training, and interoperability efforts to prevent mismatches between technical abilities and reality.

8.3.4 Integrate Co-Development and Change Management

To reduce resistance and align expectations, planners should integrate stakeholder involvement not only in design but throughout deployment. Roadmaps should be accompanied by explicit strategies for change communication, role redefinition, and fairness perception. Aligning automation gains with perceived benefits for employees increases support and long-term success.

8.3.5 Roadmaps as Adaptive Instruments

Given the uncertainties in technology development and external shocks, autonomy roadmaps should be treated as dynamic guidance rather than static plans. This requires periodic review and adjustment, performance-based triggers for progression, and the ability to adjust phases in response to empirical monitoring data. Governance operators or working groups can facilitate this flexibility by integrating technical, organisational, and stakeholder feedback.

9 Conclusion

This study explored how an adaptive orchestration system can facilitate Schiphol Airport in its long-term mission toward autonomous airside operations. Rather than focusing solely on technological feasibility, the study framed autonomy as a multi-dimensional challenge requiring synchronised progress across systems, organisations, and stakeholders. The central research question asked: *Design an adaptive orchestration model and maturity-aligned transition plan to facilitate the implementation of autonomous airside operations at Schiphol Airport, while ensuring stakeholder integration, scalable automation, and human–machine collaboration across multiple levels of autonomy*

The study originates from the need to facilitate a transition toward autonomous docking operations at Schiphol, addressing the immediate challenge of coordinating increasing automation without compromising safety, efficiency, or human engagement. The roadmap and orchestration framework presented here provide the foundational architecture to manage this transition, while deliberately designing for scalability. Though the current scope focused on the aircraft docking process, the principles and logic underlying the roadmap are extendable to wider airside functions such as baggage handling, pushback, and ground servicing. In this way, the proposed orchestration model serves not only current integration needs but also as a platform for long-term system-wide autonomy.

To address the central research question, the roadmap was developed not as a blueprint for automation itself, but as a coordination framework, sequencing interventions based on when they can be integrated into the wider ecosystem. Grounded in Thomson et al.'s three-level maturity model, the plan allows AI capabilities, semi-autonomous vehicles, and organisational processes to evolve in step rather than in isolation [76]. A core insight is that orchestration systems, though often the fastest to mature, must initially serve as decision-support tools to avoid premature delegation of control in safety-critical contexts.

The study also addressed three key sub-questions indirectly through the development of the orchestration roadmap and its implications. First, the architectural foundations for orchestrating different autonomy-level operations were derived from existing control principles and integrated into the roadmap's control area design. Second, the human–machine collaboration approach, featuring supervisory roles, AI-assisted decision-making, and upskilling trajectories, was reflected in the roadmap's maturity-aligned role evolution. Third, strategies for achieving real-time interoperability and coordination, such as instant and predictive problem solving and central oversight mechanisms, were incorporated into the proposed orchestration framework. Ultimately, standing as the foundation of the strategic roadmap development.

The study also revealed that resistance to AI is not static but can be reconfigured through meaningful involvement, reliability, and shared value creation. Early-phase adoption of assistive tools enables stakeholders to build familiarity, which becomes a cultural enabler in later, more autonomous phases. This human-centric logic is a critical factor for success, especially in environments with strong safety cultures and established roles.

Importantly, the roadmap acknowledges second-order effects. As the futures wheel analysis illustrated, operational efficiency gains may result in increased flight capacity, potentially counteracting environmental goals. Similarly, automation raises questions of job transformation, liability, and procedural fairness, issues that require governance mechanisms beyond engineering. This suggests that orchestration planning must be paired with policy frameworks to ensure ethical and environmental alignment.

Although tailored to Schiphol, the roadmap's principles are transferable to other domains undergoing automation. Sectors such as ports, logistics, and public infrastructure could adopt similar phased strategies using readiness thresholds and stakeholder co-development. Still, future work must validate the roadmap through technical pilots, cost modelling, and regulatory exploration.

In conclusion, automation transitions and orchestration models should not be treated as linear technological upgrades but as coordinated, maturity-sensitive evolutions of socio-technical systems. When orchestration, human factors, and organisational readiness are treated as equal design priorities, autonomy becomes not only implementable, but sustainable, inclusive, and future-proof.

References

[1] Federal Aviation Administration (FAA). *Air Traffic Control Decision Support Tool Design and Implementation Handbook*. Tech. rep. William J. Hughes Technical Center, 2019. URL: http://actlibrary.tc.faa.gov.

[2] Acceldata. "Orchestration Layer: Coordinating Data for Seamless Integration". In: (2024). Accessed: 2025-05-23. URL: https://www.acceldata.io/blog/orchestration-layer-explained-functions-tools-and-best-practices.

[3] D. Acemoglu and P. Restrepo. "Automation and new tasks: How technology displaces and reinstates labor". In: *Journal of Economic Perspectives* 33.2 (2019), pp. 3–30.

[4] Federal Aviation Administration. *Airport Ground Vehicle Operations: An FAA Guide*. Tech. rep. Washington, DC: FAA Office of Runway Safety, 2010.

[5] Federal Aviation Administration. *Principles for Human-Automation Teaming in Air Traffic Control*. Tech. rep. Washington, DC: U.S. Department of Transportation, 2019.

[6] P. Agbaje et al. "Survey of interoperability challenges in the Internet of Vehicles". In: *IEEE Transactions on Intelligent Transportation Systems* 23.12 (2022), pp. 22838–22861.

[7] Schiphol Airport. *Schiphol tests autonomous buses for passenger transfer*. Accessed: 2025-06-05. 2024. URL: https://www.schiphol.nl/en/innovation/blog/schiphol-tests-autonomous-buses-for-passenger-transfer/.

[8] Abdullah Alghadeir and Hasan Al- Sakran. "Smart Airport Architecture Using Internet of Things". In: *International Journal of Innovative Research in Computer Science & Technology (IJIRCST)* 4 (5 Sept. 2016). Available at SSRN: https://ssrn.com/abstract=3534138. ISSN: 2347-5552. URL: https://ssrn.com/abstract=3534138.

[9] B. Alix, C. Dupont, and A. Smith. "Rethinking human–machine collaboration in autonomous operations: A socio-technical approach". In: *International Journal of Aviation Management* 5.2 (2021), pp. 134–150.

[10] C. Alix et al. "Empowering adaptive human autonomy collaboration with artificial intelligence". In: *16th International System of Systems Engineering Conference (SoSE)*. Västerås, Sweden, 2021.

[11] Amsterdam Airport Schiphol. *The Predictive Power of Wilbur*. Feb. 2020. URL: https://www.schiphol.nl/en/innovation/blog/the-predective-power-of-wilbur/.

[12] Amsterdam Airport Schiphol. *Deep Turnaround Brochure*. Mar. 2024. URL: https://assets.ctfassets.net/biom0eqyyi6b/1VwKaoZAoN86AEI1rSaKaS/b85417249e52da4f632ea31456ccc009/deepturnaround_brochure.pdf.

[13] Jan Baert. "Remote Apron Management at Airports: Human Factors and Implementation Strategies". PhD thesis. University of Antwerp, 2023. URL: https://repository.uantwerpen.be/docman/irua/8f5b3bmotoMa5.

[14] Jack J. Baroudi, Margrethe H. Olson, and Blake Ives. "A Model of User Involvement for Information System Implementation". In: *MIS Quarterly* 10.1 (1986), pp. 125–143. URL: https://misq.umn.edu/a-model-of-users-perspective-on-change-the-case-of-information-sy.

[15] Simon Batzner et al. "EfficientAD: Accurate Visual Anomaly Detection at Millisecond-Level Latencies". In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* (2024). Accessed: 2025-06-04, pp. 128–137. URL: https://openaccess.thecvf.com/content/WACV2024/papers/Batzner_EfficientAD_Accurate_Visual_Anomaly_Detection_at_Millisecond-Level_Latencies_WACV_2024_paper.pdf.

[16] W. Beelaerts van Blokland et al. *Future airport turnaround ground handling*. 2008.

[17] David N. Bengston. "The futures wheel: A method for exploring the implications of social–ecological change". In: *Society & Natural Resources* 29.3 (2016), pp. 374–379.

[18] Simon Elias Bibri. "Backcasting in futures studies: a synthesized scholarly and planning approach to strategic smart sustainable city development". In: *European Journal of Futures Research* 6.1 (2018), p. 13.

[19] Boston Consulting Group. *AI Agents: What They Are and Their Business Impact*. Accessed: 2025-05-23. 2025. URL: https://www.bcg.com/capabilities/artificial-intelligence/ai-agents.

[20] S. Bouyakoub et al. "Smart airport: An IoT-based airport management system". In: *ACM ICFNDS '17*. Cambridge, United Kingdom, 2017.

[21] S. Bouyakoub et al. "Smart airport: An IoT-based airport management system". In: *Proceedings of the International Conference on Future Networks and Distributed Systems*. ACM, 2017.

[22] Arwin van Buuren. "Collaborative Problem Solving in a Complex Governance System: Amsterdam Airport Schiphol and the Challenge to Break Path Dependency". In: *ResearchGate* (2012). Accessed: 2025-05-23. URL: https://www.researchgate.net/publication/264760479_Collaborative_Problem_Solving_in_a_Complex_Governance_System_Amsterdam_Airport_Schiphol_and_the_Challenge_to_Break_Path_Dependency.

[23] CBS. *76 million airline passengers in 2024: 6 percent fewer than in 2019*. 2024.

[24] Yokogawa Electric Corporation. *Autonomous Operations – Digital Infrastructure Wiki*. Accessed: [Insert Access Date]. 2023. URL: https://www.yokogawa.com/eu/solutions/featured-topics/digital-infrastructure-wiki/general/autonomous-operations/.

[25] Alexandra Curry and Donald Hodgson. "Emerging Features of Future Systems". In: *Some Journal* 12 (2008), pp. 34–56.

[26] Daifuku ATEC. *Legacy Systems: The Shadow in Aviation Progress*. https://daifukuatec.com/blog/legacy-systems-the-shadow-in-aviation-progress?utm_source=chatgpt.com. Accessed: 2025-06-05. 2024.

[27] Clara Dias and Jorge Silva. "Unveiling the future: Smart airports - applications, advantages, strategies and technological challenges". In: *Journal of Airline and Airport Management* 14.2 (2024). DOI: 10.3926/jairm.419.

[28] Future Travel Experience. *Royal Schiphol Group building a roadmap towards a "fully autonomous airside operation" in 2050*. Accessed: [date you accessed it]. Mar. 2021. URL: https://www.futuretravelexperience.com/2021/03/royal-schiphol-group-building-a-roadmap-towards-a-fully-autonomous-airside-operation-in-2050/.

[29] Future Travel Experience. *How Fraport's innovative AI@FRA initiative is enhancing operational capabilities and passenger experience*. Accessed: [insert date you accessed it]. Sept. 2024. URL: https://www.futuretravelexperience.com/2024/09/how-fraports-innovative-aifra-initiative-is-enhancing-operational-capabilities-and-passenger-experience/.

[30] Federal Aviation Administration. *Autonomous Ground Vehicle Systems (AGVS) on Airports*. https://www.faa.gov/airports/new$_entrants/agvs_on_airports$. https://www.faa.gov/airports/new_entrants/agvs_on_airports?utm_source=chatgpt.com. 2024.

[31] Future Travel Experience. *Munich Airport introduces smart baggage trolleys for a more digital passenger experience*. Accessed: [insert date you accessed it]. Oct. 2023. URL: https://www.futuretravelexperience.com/2023/10/munich-airport-introduces-smart-baggage-trolleys-for-a-more-digital-passenger-experience/.

[32] Mario H. Castañeda Garcia et al. "A Tutorial on 5G NR V2X Communications". In: *IEEE Communications Surveys & Tutorials* 23.3 (2021), pp. 1972–2028. DOI: 10.1109/COMST.2021.3057017.

[33] R. Gödöllei and S. Beck. "Insecure or optimistic? Employees' diverging appraisals of automation and consequences for job attitudes". In: (2023).

[34] G. Gomez-Beldarrain et al. "Why does automation adoption in organizations remain a fallacy?" In: (2025).

[35] G. Gomez-Beldarrain et al. "Revealing the challenges to automation adoption in organizations: Examining practitioner perspectives from an international airport". In: *CHI EA '24, Honolulu, HI, USA*. ACM, 2024. DOI: 10.1145/3613905.3650964.

[36] G. Gómez-Beldarrain et al. "Revealing the challenges to automation adoption in organizations: Examining practitioner perspectives from an international airport". In: *Extended Abstracts of CHI '24*. ACM, 2024.

[37] Theodore J. Gordon and David Hayward. “Initial Experiments with the Cross-Impact Matrix Method of Forecasting”. In: *Futures* 1.2 (1968), pp. 100–116.

[38] ARC Advisory Group. *What are Autonomous Operations?* Accessed: [Insert Access Date]. 2023. URL: https://www.arcweb.com/industry-best-practices/what-autonomous-operations.

[39] Changi Airport Group. *CAG collaborates with Aurrigo to trial new autonomous baggage handling vehicle*. 2024. URL: https://www.changiairport.com/en/corporate/our-media-hub/newsroom/press-releases.cag-collaborates-with-aurrigo-to-trial-new-autonomous-baggage-handling-vehicle.2024.all.html.

[40] T. Hajam and P. John. “Reskilling and upskilling strategies in the era of automation”. In: (2024).

[41] T. Hardjono, A. Lipton, and A. Pentland. “Toward an interoperability architecture for blockchain autonomous systems”. In: *IEEE Transactions on Engineering Management* 67.4 (2020), pp. 1109–1121.

[42] B. Heerens. *KLM ground services workforce transformation for autonomous operations*. Tech. rep. Amsterdam: Royal Schiphol Group, 2025.

[43] IATA. “Generative AI and aviation: Finding crossroads for future implementation”. In: (2023).

[44] International Airport Review. *The Future is Here: Airports Enter the Age of Automation*. https://www.internationalairportreview.com/article/260776/the-future-is-here-airports-enter-the-age-of-automation/?utm_source=chatgpt.com. Accessed: 2025-06-05. 2024.

[45] M. Jaradat et al. “The internet of energy: Smart sensor networks and big data management for smart grid”. In: *Procedia Computer Science* 56 (2015), pp. 592–597.

[46] Abdullah Kormaz and Abdullah Okumus. “Public Acceptance towards Emerging Autonomous Vehicle Technology: A Bibliometric Analysis”. In: *Sustainability* 15.2 (2023). Accessed: 2025-05-23, p. 1566. DOI: 10.3390/su15021566. URL: https://www.mdpi.com/2071-1050/15/2/1566.

[47] D. Kovynyov and R. Mikut. “Digital transformation in airport ground operations”. In: (2024).

[48] I. Kovynyov and R. Mikut. “Digital technologies in airport ground operations”. In: *NETNOMICS: Economic Research and Electronic Networking* 20.1 (2019), pp. 1–30.

[49] Paolo Malighetti et al. “The turnaround tactic and on-time performance: Implications for airlines’ efficiency”. In: *Research in Transportation Business & Management* 46 (Jan. 2023), p. 100874. DOI: 10.1016/j.rtbm.2023.100874.

[50] M. Lynne Markus. “Power, Politics, and MIS Implementation”. In: *Communications of the ACM* 26.6 (1983). Accessed: 2025-06-04, pp. 430–444. DOI: 10.1145/358141.358148. URL: https://doi.org/10.1145/358141.358148.

[51] B. Mietkiewicz et al. “Enhancing control room operator decision-making”. In: (2024).

[52] Akshra Narasimhan Ramakrishnan. “Design and Evaluation of Perception System Algorithms for Semi-Autonomous Vehicles”. Master of Science Thesis. Department of Electrical and Computer Engineering: Ohio State University, 2020. URL: http://rave.ohiolink.edu/etdc/view?acc_num=osu1595256912692618.

[53] NASA. “Adjustable Autonomy for Human-Centered Autonomous Systems”. In: *NASA Technical Reports* (2020). Accessed: 2025-06-04. URL: https://www.nasa.gov/adjustable-autonomy-human-centered.

[54] Fedja Netjasov and Milan Janic. “A review of the research on risk and safety modelling in civil aviation”. In: *Proceedings of the 3rd International Conference on Research in Air Transportation (ICRAT 2008)*. Fairfax, VA, USA, June 2008.

[55] M. D. Neto et al. “Assessment method of fuel consumption and emissions during taxiing”. In: (2022).

[56] Debabrata Pal, Anvita Singh, and Abhishek Alladi. “SIRT: Machine-Learning-based Selective Intensity Range Thresholding for Aircraft Visual Docking Guidance Refinement and Interpretation”. In: *2022 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECCT)*. 2022, pp. 1–6. DOI: 10.1109/CONECCT55679.2022.9865711.

[57] J. Pretlove and S. Royston. “Towards autonomous operations: A six-level automation maturity model”. In: *Proceedings of the 2023 Offshore Technology Conference*. Houston, TX, 2023.

[58] ProRail. *Operationeel Controle Centrum Rail (OCCR) in major disruptions*. n.d. URL: https://www.prorail.nl/veelgestelde-vragen/wat-is-de-rol-van-het-operationeel-controle-centrum-rail-occr-bij-grote-verstoringen.

[59] Krish Ramachandran. *Latest Advances in Agentic AI Architectures and Frameworks*. Accessed: 2025-06-04. LinkedIn. 2025. URL: https://www.linkedin.com/pulse/latest-advances-agentic-ai-architectures-frameworks-ramachandran-ry3fe.

[60] Resume Now. “AI Disruption: 9 in 10 Workers Fear Job Loss to Automation”. In: (2025). Accessed: 2025-05-23. URL: https://www.prweb.com/releases/ai-disruption-9-in-10-workers-fear-job-loss-to-automation-302371514.html.

[61] Martijn Ringelberg. “Airline strategies’ impact on gate occupancy at Schiphol airport”. Full-text available. MA thesis. Delft University of Technology, June 2018.

[62] RIVM. “Footprint of aviation at Schiphol on air quality”. In: (2021).

[63] Bruce Rogers. *Why 84% Of Companies Fail At Digital Transformation*. Accessed: 2025-06-04. 2016. URL: https://www.forbes.com/sites/brucerogers/2016/01/07/why-84-of-companies-fail-at-digital-transformation/.

[64] G. de Rooij, A. B. Tisza, and C. Borst. “Flight-Based Control Allocation: Towards Human–Autonomy Teaming in Air Traffic Control”. In: *[Journal Name - If Available]* (2024). Accessed: [date you accessed it]. URL: http://resolver.tudelft.nl/uuid:add59c4b-0dee-4c4c-9e4c-c68d48fa972f.

[65] M. San Emeterio de la Parte et al. “Breaking down IoT silos: A semantic interoperability support system for the Internet of Things”. In: *Proceedings of the 3rd International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME 2023)*. IEEE, 2023, pp. 441–446.

[66] Schiphol. *An autonomous airport in 2050*. Accessed: [date you accessed it]. Feb. 2020. URL: https://www.schiphol.nl/en/innovation/blog/an-autonomous-airport-in-2050/.

[67] Schiphol. *Volgende fase van testen met autonoom bagagevoertuig*. Accessed: [date you accessed it]. Feb. 2025. URL: https://www.schiphol.nl/nl/innovatie/blog/volgende-fase-van-testen-met-autonoom-bagagevoertuig/.

[68] Schiphol Airport. *Autonomous Airside Operations*. Accessed: 2025-06-05. 2024. URL: https://www.schiphol.nl/en/innovation/blog/autonomous-airside-operations/.

[69] ShinMaywa Industries, Ltd. *World’s First Fully Remote-Controlled Passenger Boarding Bridge Installed at Changi Airport*. https://www.shinmaywa.co.jp/english/products/paxway/news/20230825_01.html?utm_source=chatgpt.com. Accessed: 2025-06-05. 2023.

[70] Statista. “Chart: Fatal Accidents Damage Trust in Autonomous Driving”. In: (2018). Accessed: 2025-05-23. URL: https://www.statista.com/chart/13450/perceived-safety-of-self-driving-cars/.

[71] Statista. *Leading airports in Europe by passenger transport*. 2023. URL: https://www.statista.com/statistics/1118486/air-passenger-transport-leading-airports-europe/.

[72] Diego Alonso Tabares, Felix Mora-Camino, and Antoine Drouin. “A multi-time scale management structure for airport ground handling automation”. In: *Journal of Air Transport Management* 90 (Jan. 2021), p. 101959. DOI: 10.1016/j.jairtraman.2021.101959.

[73] Andy Taylor. “Digital Towers and Artificial Intelligence”. In: *International Airport Review* (2022). Accessed: 2025-06-04. URL: https://www.internationalairportreview.com/article/178144/digital-towers-artificial-intelligence-air-traffic-management/.

[74] F. Tener and J. Lanir. “Devising a high-level command language for the teleoperation of autonomous vehicles”. In: *International Journal of Human–Computer Interaction* (2024). Advance online publication.

[75] Linus Thomson et al. “A maturity framework for autonomous solutions in manufacturing firms: The interplay of technology, ecosystem, and business model”. In: *International Entrepreneurship and Management Journal* 18.1 (2021), pp. 125–152. DOI: 10.1007/s11365-020-00717-3.

[76] Linus Thomson et al. “A maturity framework for autonomous solutions in manufacturing firms: The interplay of technology, ecosystem, and business model”. In: *International Entrepreneurship and Management Journal* 18.1 (2022), pp. 125–152.

[77] T. Truitt. “Electronic flight data in airport traffic control towers”. In: (2006). URL: https://your-source-link.com.

[78] J. Urso. “Autonomous Operations Maturity Model”. In: *Honeywell User Group (HUG) 2021 Conference*. Phoenix, AZ, June 2021.

[79] Vertical Magazine. “Certification challenges for autonomous aircraft systems”. In: (2019). Accessed: 2025-05-23. URL: https://verticalmag.com/opinions/certification-challenges-for-autonomous-aircraft-systems/.

[80] Wolters Kluwer. “Future-proofing internal audit: Tackling emerging technological disruptions”. In: (2025). Accessed: 2025-05-23. URL: https://www.wolterskluwer.com/en/expert-insights/future-proofing-internal-audit-tackling-emerging-technological-disruptions.

[81] Zembla. *Episode 24 (Ziek van Schiphol)*. Television broadcast. Zembla: Season 30, Episode 24. Dec. 2021.

A Original Version Swimlane

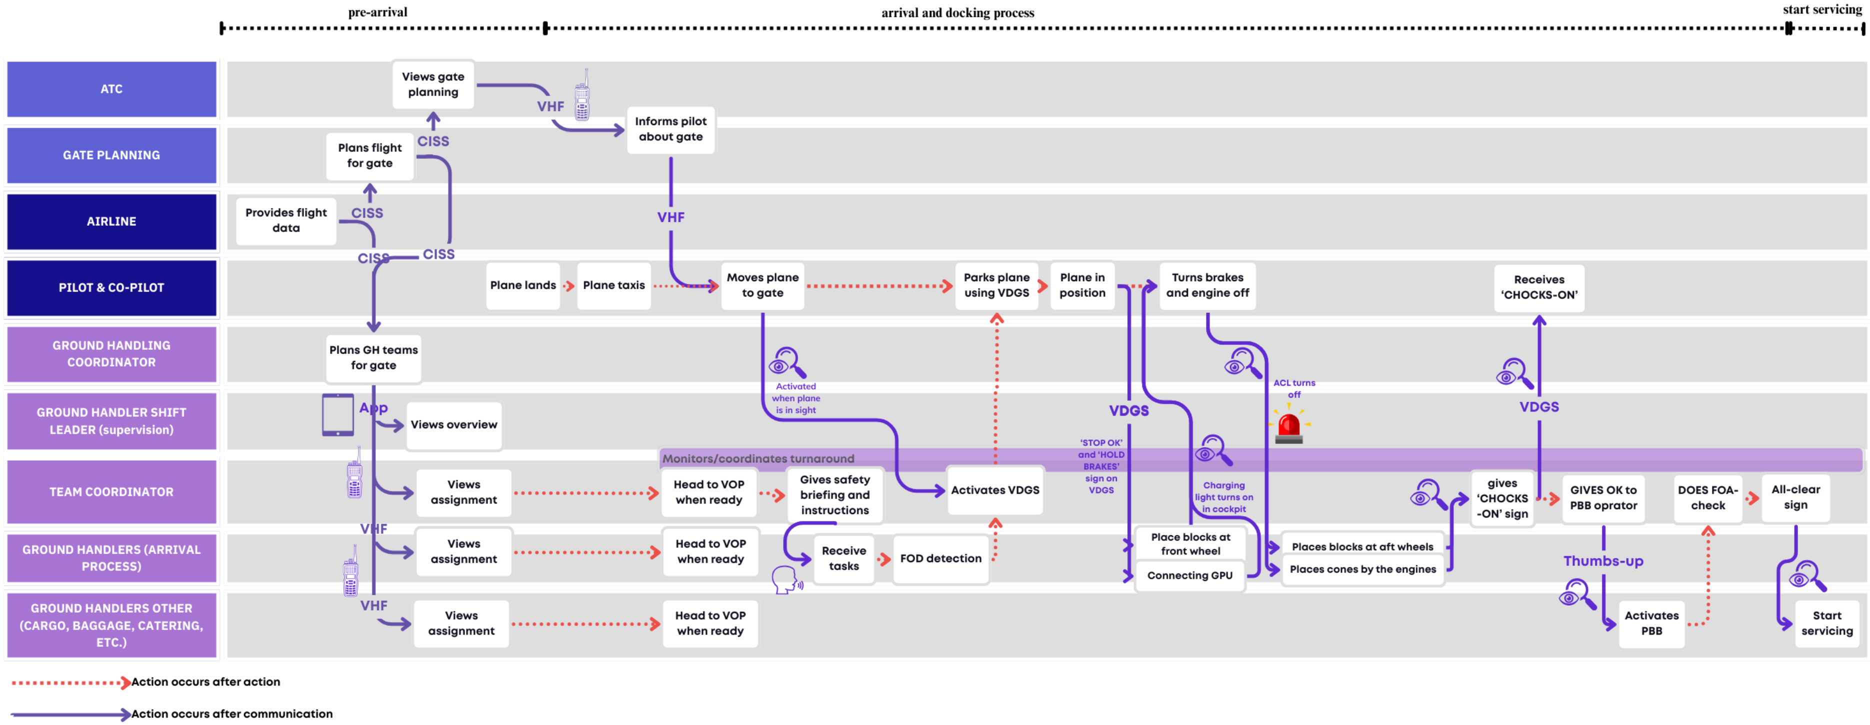

B Backcasting Method

Backcasting

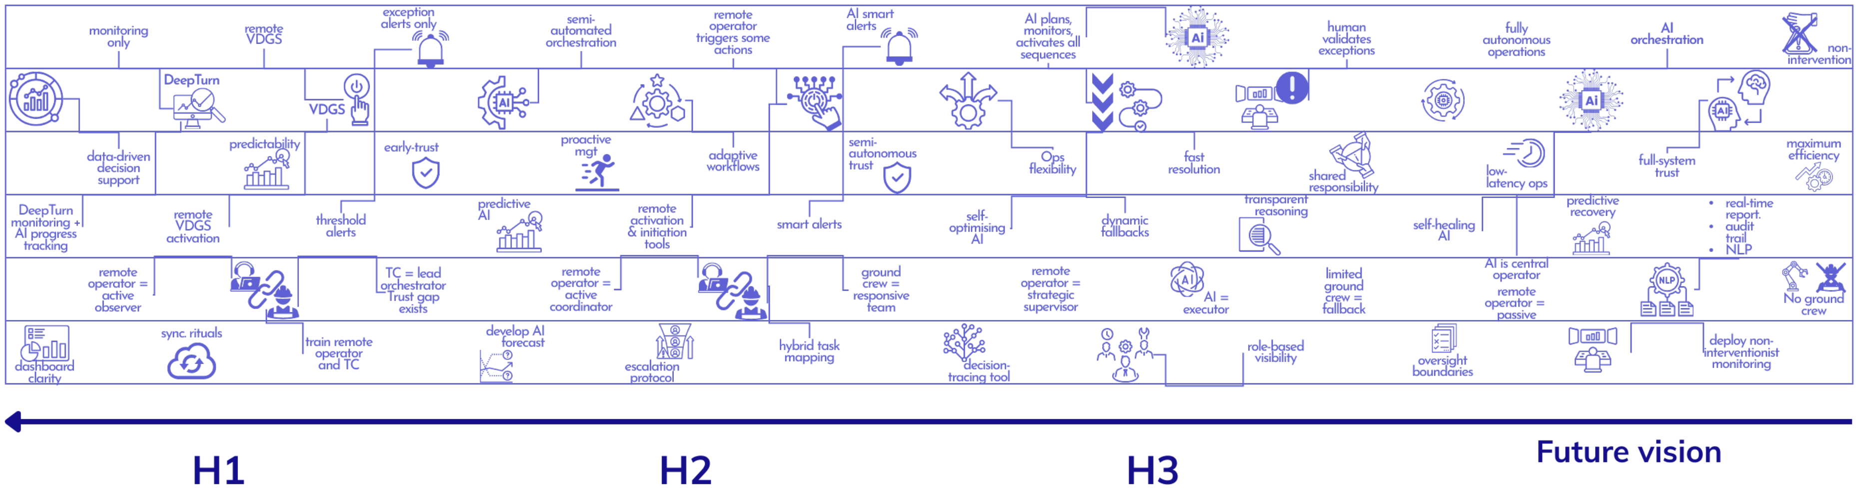

Future vision- Fully autonomous operations
This is the ultimate vision in which airside operations at Schiphol are fully autonomous, self-orchestrating, and self-healing. All ground processes, from docking guidance to GPU connection and start servicing, are carried out by AI-driven systems and autonomous service vehicles. Human oversight is limited to passive oversight, triggered only in exceptional cases or emergencies via exception-based alerts. These systems operate in a coordinated, uninterrupted chain, supported by predictive obstruction detection, multi-tier fallback protocols, and real-time audit trails. A remote orchestration dashboard provides full transparency and traceability of AI decisions, ensuring compliance and trust. Achieving this states requires an established legal and regulatory framework, proven reliability across all automated tiers, strong human-machine trust, and fully integrated airside infrastructure.

Horizon 3- Remote and Resilient Operations
Moving backward, the conditions for full autonomy are set in Horizon 3, which introduces resilient, semi autonomous orchestration. AI handles most planning, execution, and adaptation in real-time, while the remote operator acts as a strategic supervisor, validating or adjusting AI logic when necessary. Ground handlers are largely reoriented into supporting roles, providing manual backup during exceptions. this phase focuses on transparency, with explainable AI interfaces allowing humans to trace decisions and simulate exception cases. The orchestration system in centralised, and fallback mechanisms are in place to ensure the system can proceed in case of obstruction. the organisation em-

braces a culture of AI-led operations, supported by training programs and shared oversight principles. Regulations start shifting toward the authorisation of systems, rather than humans.

Horizon 2- Intelligent Operations

Prior to the remote and resilient orchestration, Horizon 2 represents the transition into AI-supported intelligent and remote operations. While much of the ground equipment is now autonomous or semi autonomous, it is still initiated and monitored by remote human operators via a centralised dashboard. Predictive AI supports proactive coordination, like alerting the operator about issues before they escalate and suggesting interventions. Automated processes like FOD detection, VDGS activation, and GPU con nection may be triggered remotely. This phases requires reliable connectivity, secure system integration, and robust exception management workflows. Operators are actively involved in process activation and coordination, bridging that gap between automation and oversight. Training programs are launched to turn the remote operators into orchestration experts, and Schipholo shift its operational logic to shared human-AI decision-making.

Horizon 1- Controlled Operations

At the starting point of the transition, Horizon 1 stands for the initial digitisation and remote visibility phase. Human ground handlers are still in operational control, with new digital systems introduced to enhance their work. Technologies like Deepturn provide live monitoring of apron activity, and remote VDGS activation is introduced under strict conditions. Remote operators begin to observe the operation, but only engage during exceptions. This phase is essential for building trust, collecting operational data. and familiarising ground staff with remote collaboration. Small-scale pilots of automation are being trialed under controlled conditions. There are infrastructure upgrades which includes sensor deployment, dashboard development, and secure communication channels. By the end of Horizon 1, Schiphol has a functioning monitoring system, limited remote activations, and foundational organisational awareness of the upcoming shift toward autonomy.

C Forecasting Method

Forecasting

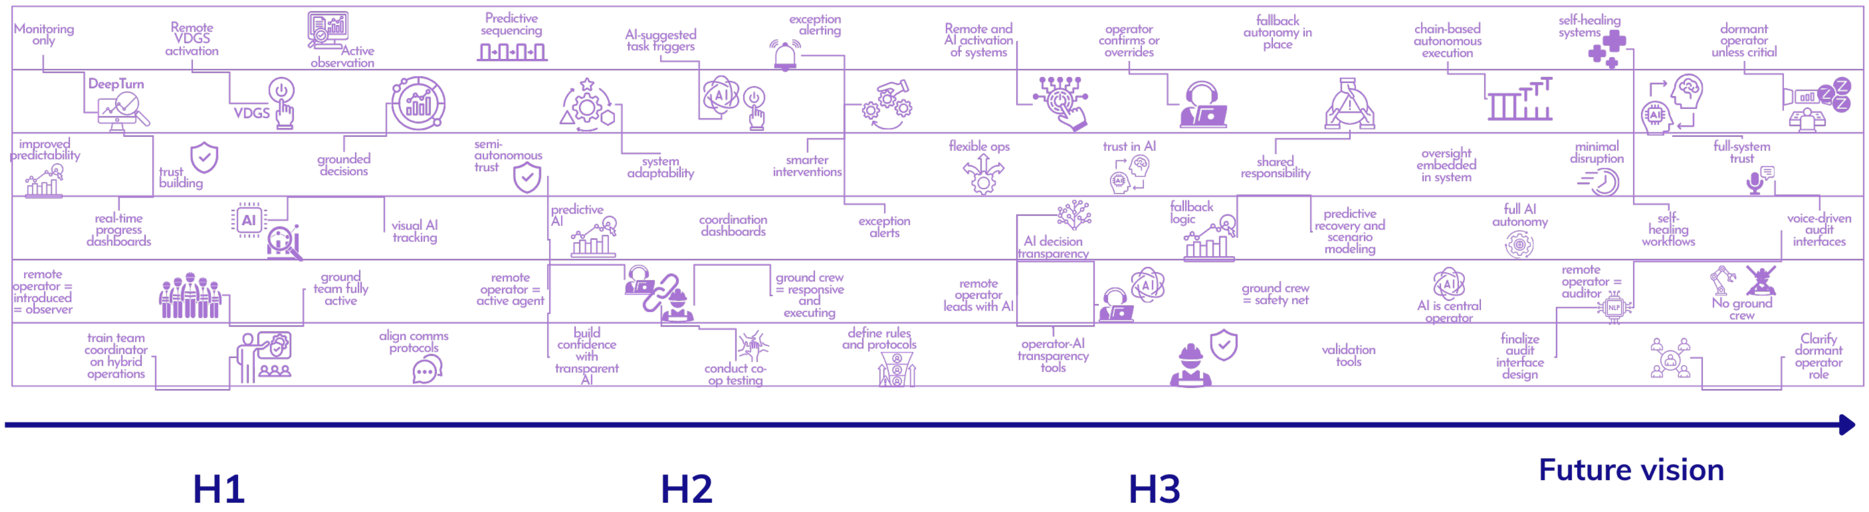

Horizon 1- Controlled Operations

In the current and near phase, the airside environment is still predominantly human-operated. The emphasis lies in building trust in remote monitoring technologies like DeepTurnaround, which provide real-time visibility of operations. Remote operators are introduced in an observational role, monitor ing progress through dashboards and activating systems like VDGS under strict safety measures. The forecasting exercise recognises that building trust is essential, ground handlers for now remain in full control, and any AI involvement is strictly supportive. Staff training, transparent communication, and early testing programs are crucial in creating confidence. These pilots may offer PoCs and generate an initial round of data-driven insights, enabling more informed decisions in later stages.

Horizon 2- Intelligent Operations

In this transition phase, the focus shifts to collaborative orchestration between AI systems and human operators. Forecasting anticipates the rise of predictive sequencing, exception alerts, and remote initiation of tasks. The remote operator becomes an active agent, initiating AI-suggested workflows via coordination dashboards. Ground crews begin functioning as responsive teams, executing tasks based on system-set assignments while adjusting to changing roles. The forecasting emphasises the importance of boundaries, transparent AI interfaces, and iterative feedback loops, all features designed to build human-machine trust. New tools and training program help redefine workflows, while confidence in automation grows through performance validation and routine use. At this stage, human-AI collaboration is implemented in a deliberate, visible and relational manner.

Horizon 3- Remote and Resilient Operations

As predictive systems and fallback logic develop further, Schiphol moves into a stage of resilient, mostly autonomous operations. The remote operator now serves as a strategic validator who supervising the AI's chain-based task orchestration and stepping in solely for exceptions or ambiguous cases. The ground crew shifts into a fallback and safety role, involved only when physical presence is required. Technology is not only autonomous, it's transparent, with clear decision-making dashboards, simulation modes, and voice-driven audit tools ensuring clarity and oversight. Emphasised is the need to maintain trust by embedding resilience, refining operator-AI interfaces, and sustaining shared accountability. Human input remains critical in the governance, training, and trust assurance, even as AI becomes the operational standard.

Future Vision- Fully Autonomous Operations

In the forecast's final state, all apron processes are conducted by fully autonomous systems, coordinated through a central AI operator. The remote operator assumes a dormant auditor role that is engaged only when system confidence is low or escalation is triggered. The forecast highlights the emergence of self-healing workflows, real-time reporting, and predictive recovery protocols. The interface becomes a supervisory tool, not so much a control panel, focused on validation and compliance (whilst still main taining override functions). The autonomous operations are not only technical development, but also an organisational change in which it is embedded into procedures, culture, and oversight frameworks. The forecast highlights the need for clear ethical boundaries, transparent decision logs, and final operator override options that are all crucial for sustaining validity and legitimacy in a fully autonomous environment.

D Backcasting and Forecasting Overlay

Merged

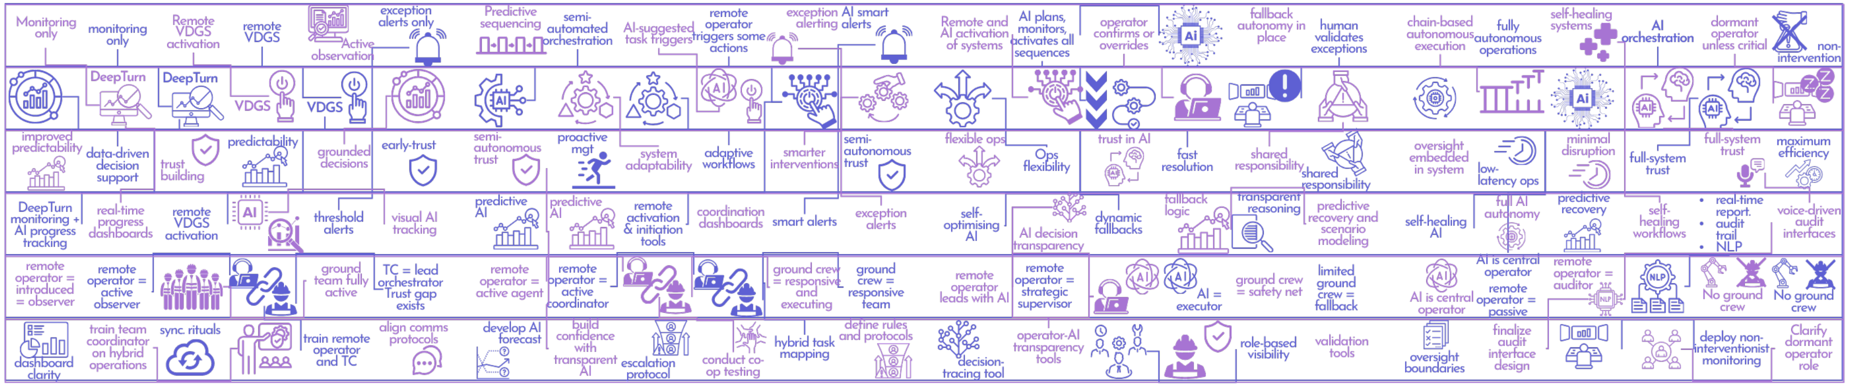

E Interactive Roadmap Link

https://flightpath-explorer.lovable.app/

F Dashboard Horizon 1

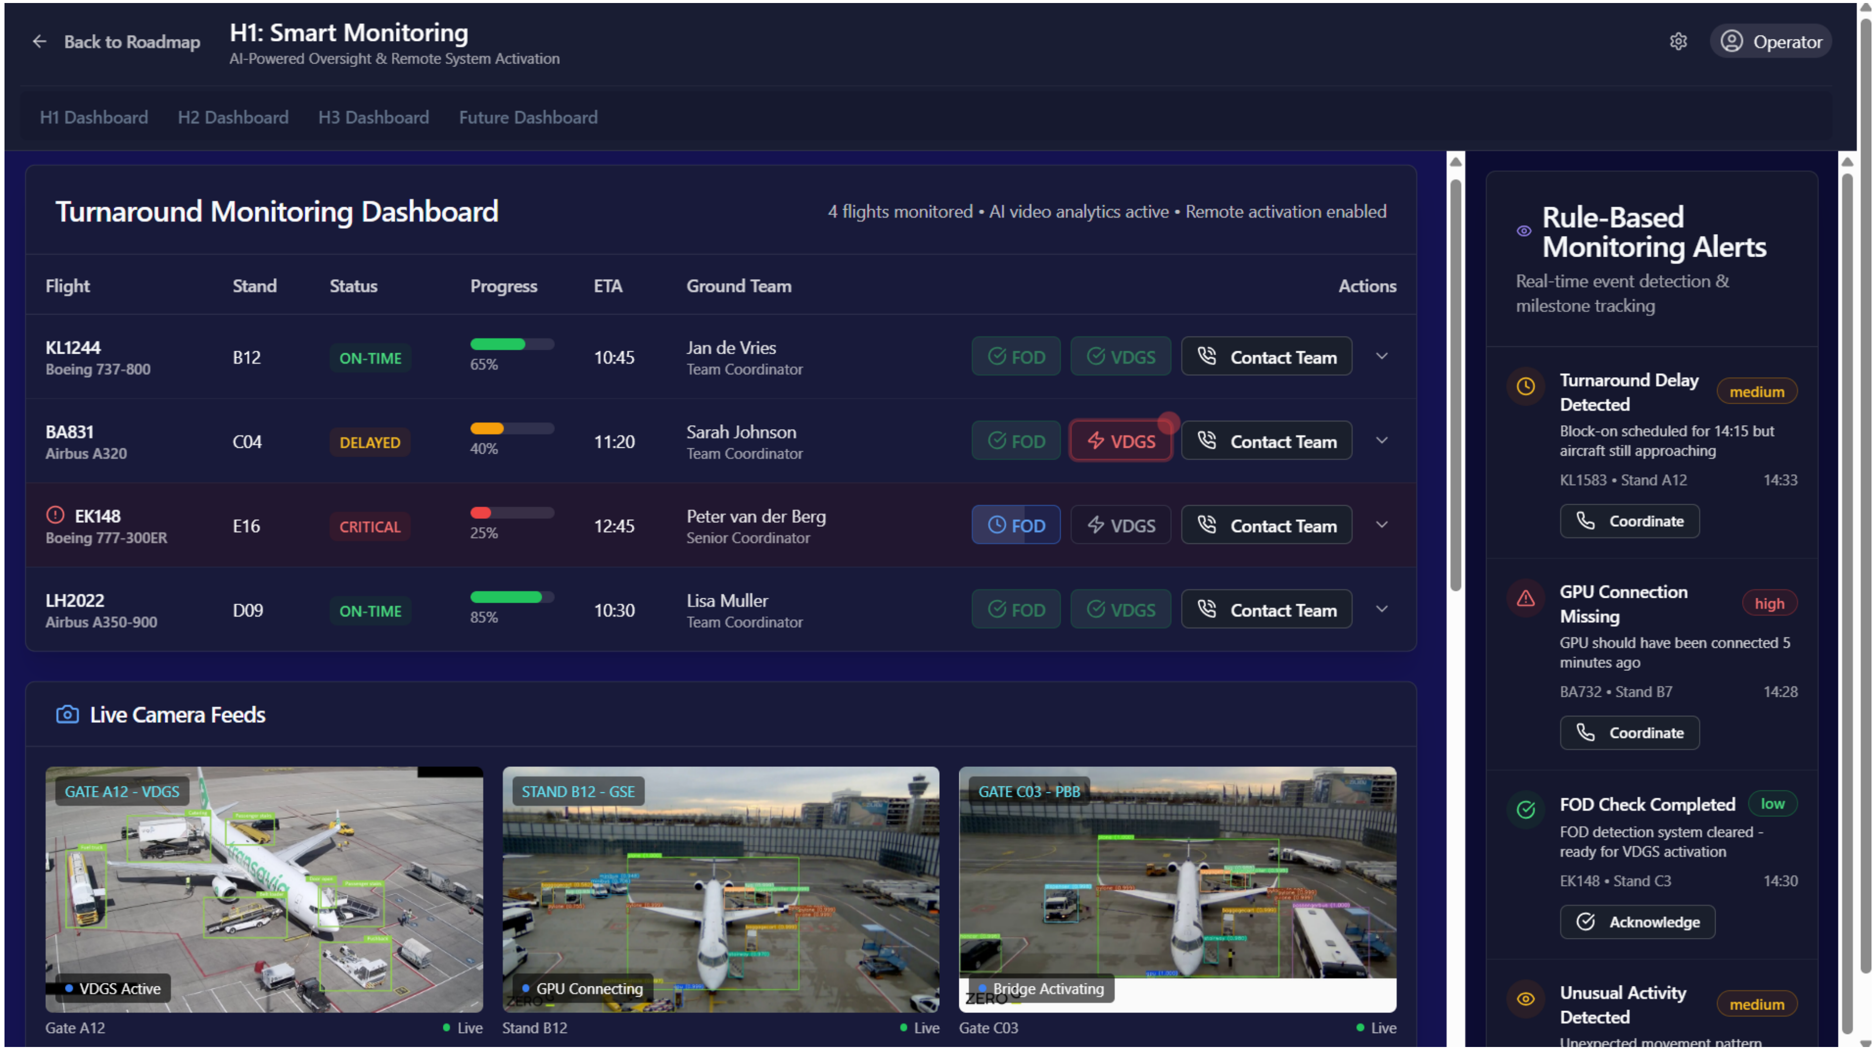

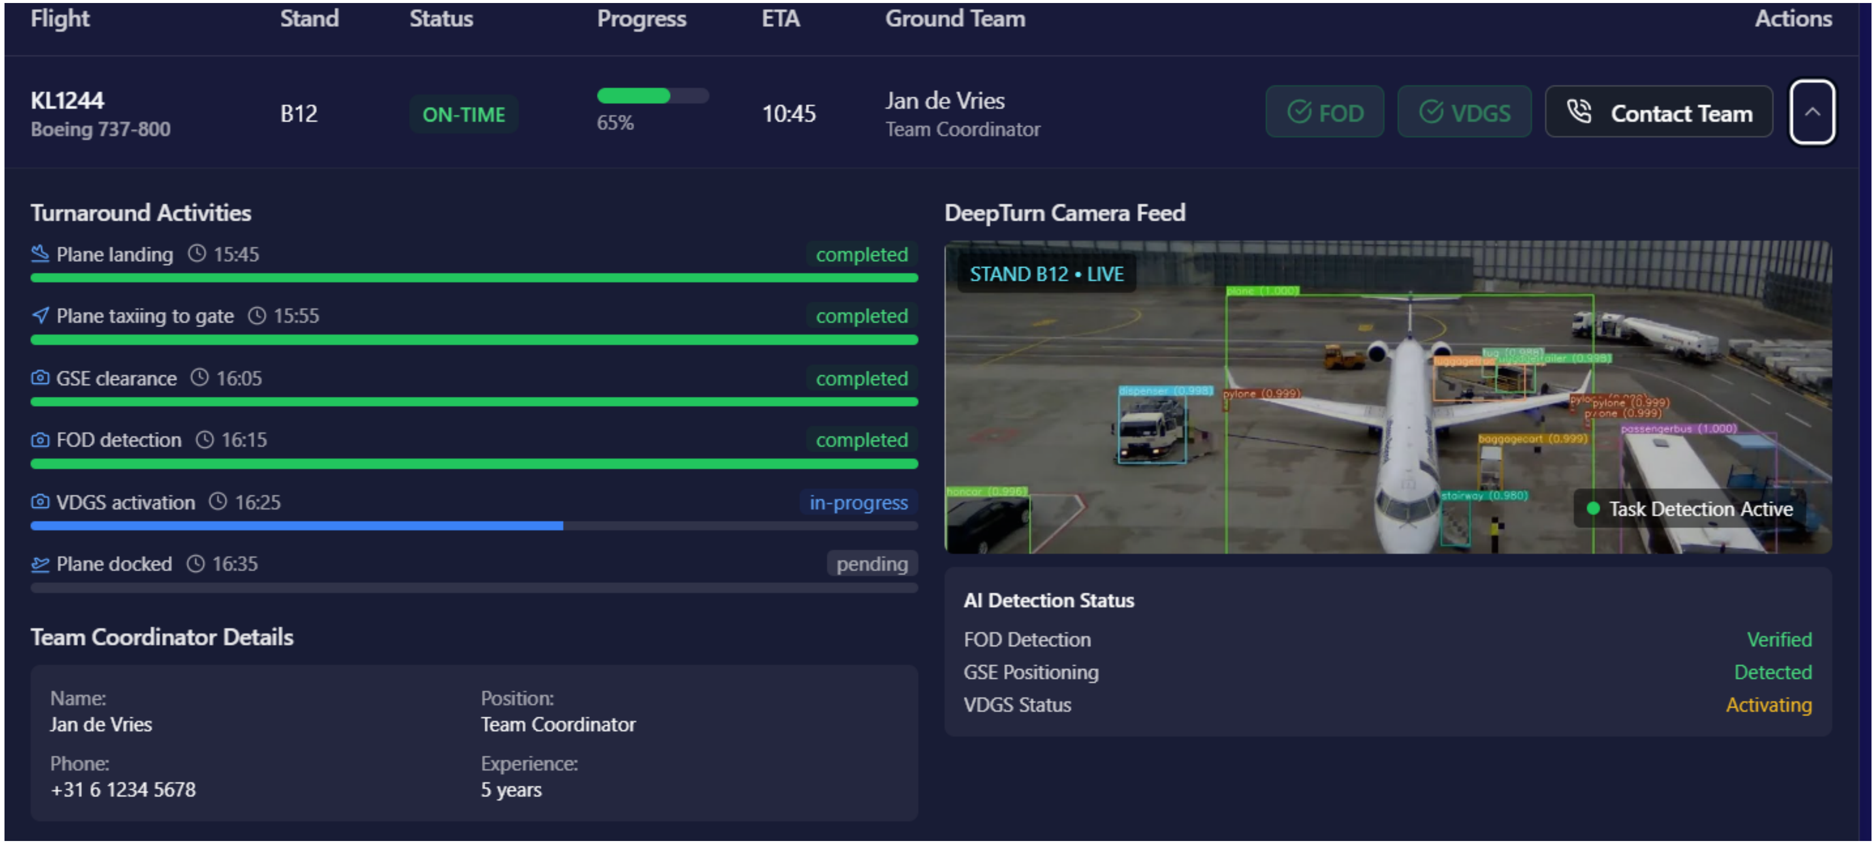
Flight
Stand
Status
Progress
ETA
Ground Team
Actions
KL1244
Boeing 737-800
B12
ON-TIME
65%
10:45
Jan de Vries
Team Coordinator
FOD
VDGS
Contact Team
Turnaround Activities
Plane landing 15:45 completed
Plane taxiing to gate 15:55 completed
GSE clearance 16:05 completed
FOD detection 16:15 completed
VDGS activation 16:25 in-progress
Plane docked 16:35 pending
Team Coordinator Details
Name:
Jan de Vries
Position:
Team Coordinator
Phone:
+31 6 1234 5678
Experience:
5 years
DeepTurn Camera Feed
STAND B12 • LIVE
Task Detection Active
AI Detection Status
FOD Detection Verified
GSE Positioning Detected
VDGS Status Activating

G Dashboard Horizon 2

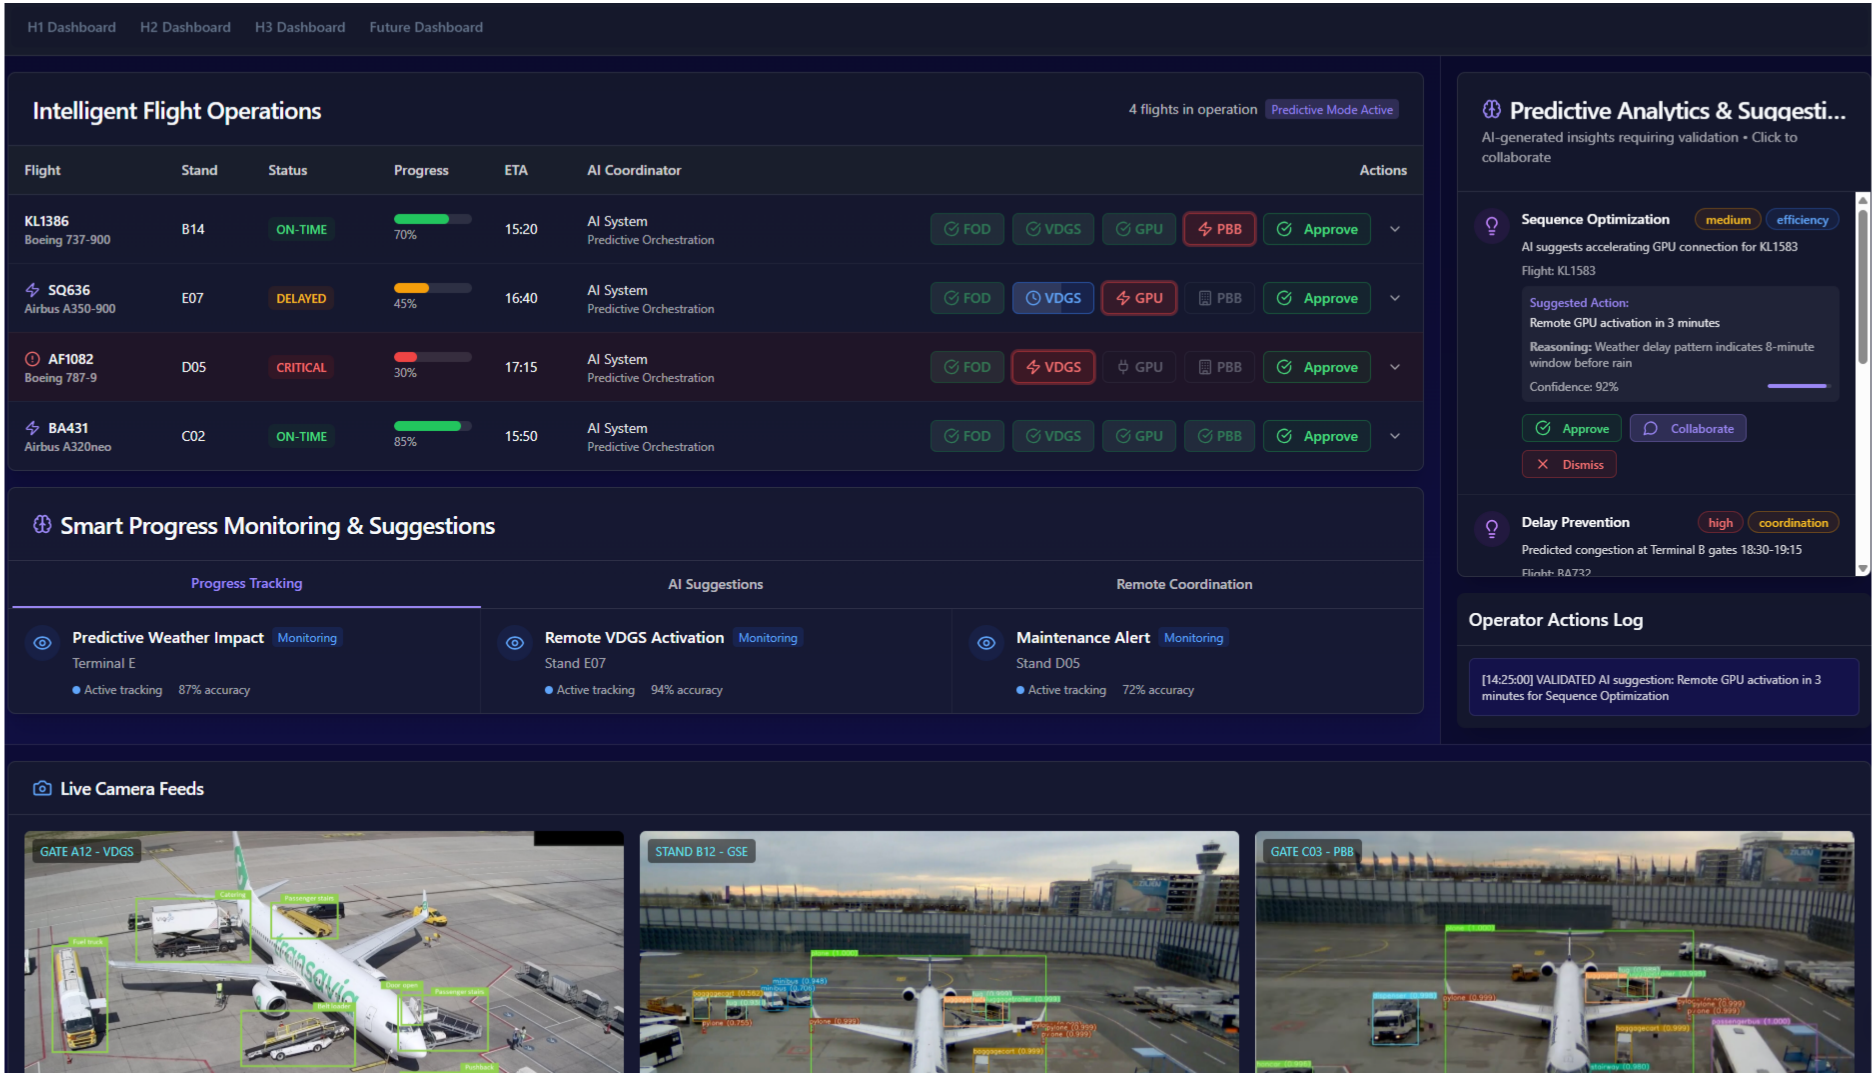

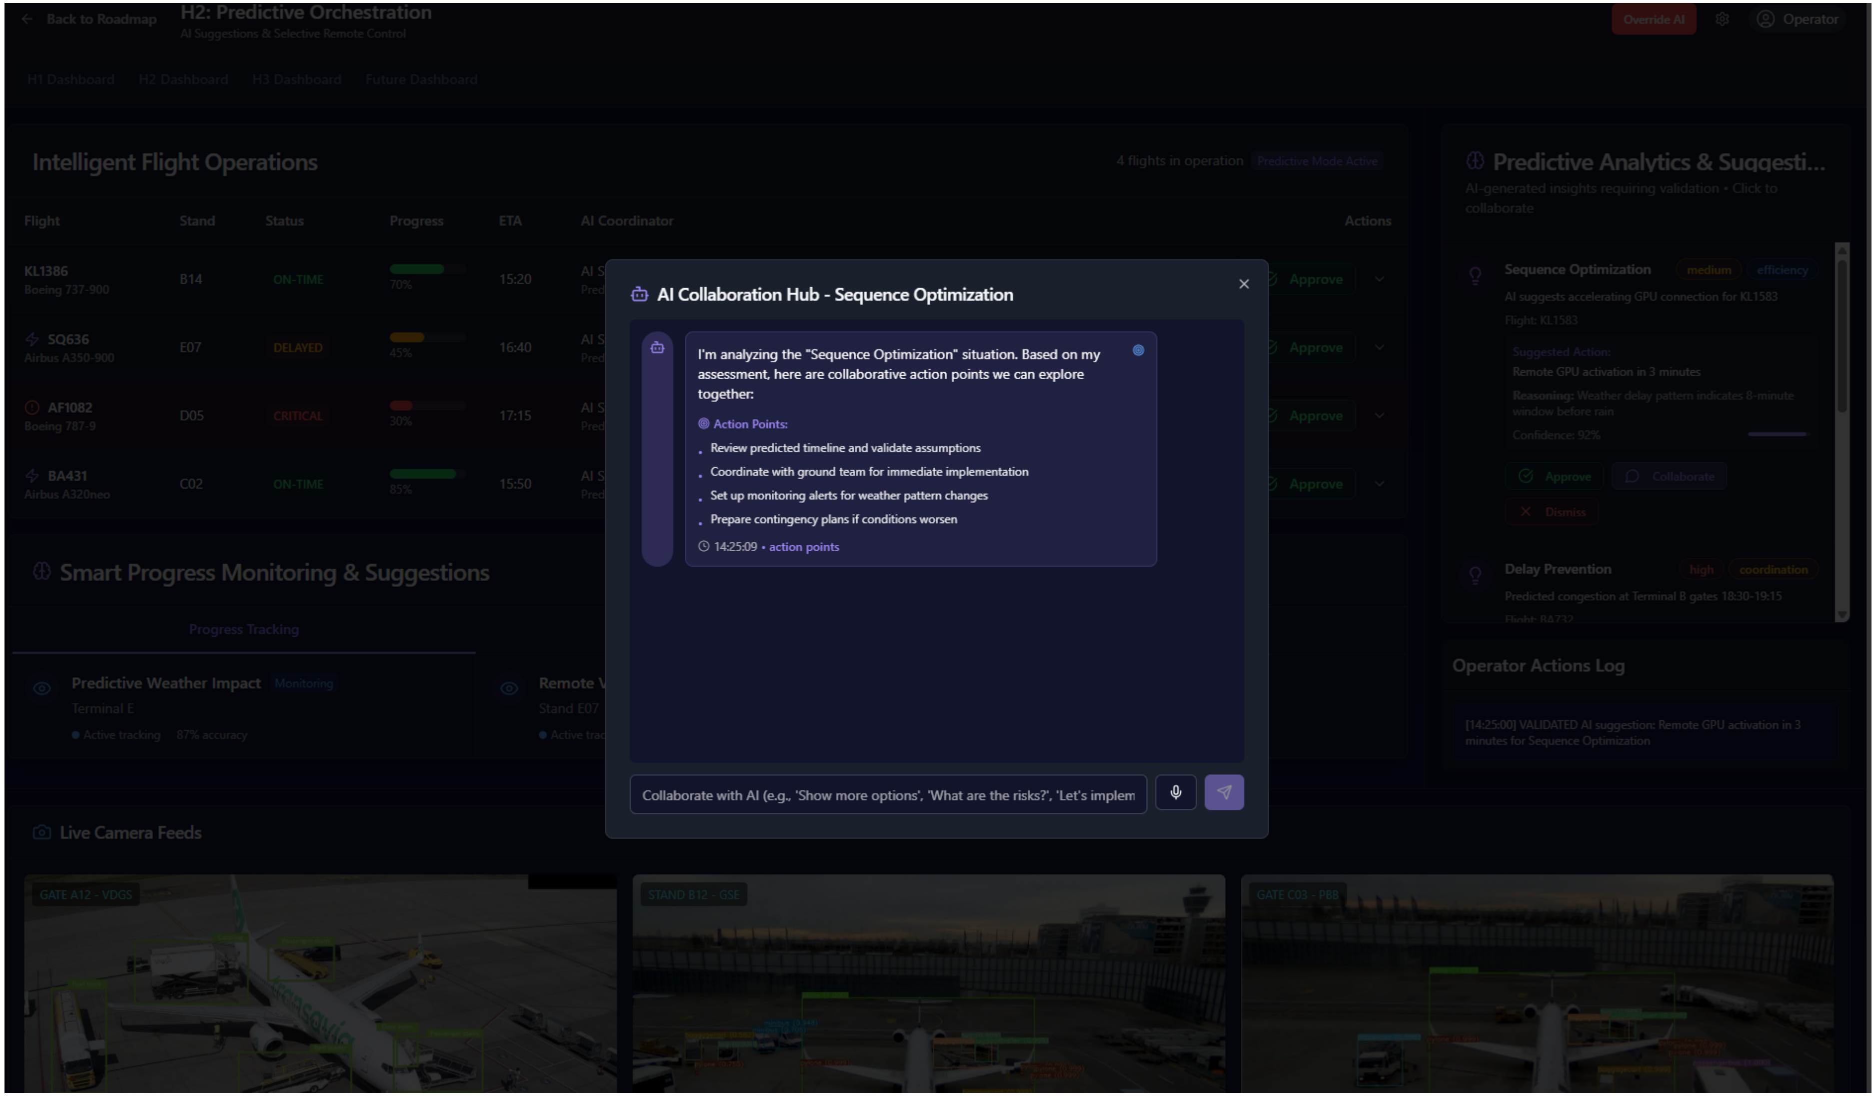
Back to Roadmap
H2: Predictive Orchestration
AI Suggestions & Selective Remote Control
Override AI
Operator
H1 Dashboard
H2 Dashboard
H3 Dashboard
Future Dashboard
Intelligent Flight Operations
4 flights in operation
Predictive Mode Active
Flight
Stand
Status
Progress
ETA
AI Coordinator
Actions
KL1386
Boeing 737-900
B14
ON-TIME
70%
15:20
SQ636
Airbus A350-900
E07
DELAYED
45%
16:40
AF1082
Boeing 787-9
D05
CRITICAL
30%
17:15
BA431
Airbus A320neo
C02
ON-TIME
85%
15:50
Approve
AI Collaboration Hub - Sequence Optimization
I'm analyzing the "Sequence Optimization" situation. Based on my assessment, here are collaborative action points we can explore together:
Action Points:
Review predicted timeline and validate assumptions
Coordinate with ground team for immediate implementation
Set up monitoring alerts for weather pattern changes
Prepare contingency plans if conditions worsen
14:25:09 • action points
Collaborate with AI (e.g., 'Show more options', 'What are the risks?', 'Let's implem
Predictive Analytics & Suggesti...
AI-generated insights requiring validation • Click to collaborate
Sequence Optimization
medium
efficiency
AI suggests accelerating GPU connection for KL1583
Flight: KL1583
Suggested Action:
Remote GPU activation in 3 minutes
Reasoning: Weather delay pattern indicates 8-minute window before rain
Confidence: 92%
Approve
Collaborate
Dismiss
Delay Prevention
high
coordination
Predicted congestion at Terminal B gates 18:30-19:15
Smart Progress Monitoring & Suggestions
Progress Tracking
Predictive Weather Impact
Monitoring
Terminal E
Active tracking
87% accuracy
Stand E07
Operator Actions Log
[14:25:00] VALIDATED AI suggestion: Remote GPU activation in 3 minutes for Sequence Optimization
Live Camera Feeds
GATE A12 - VDGS
STAND B12 - GSE
GATE C03 - PBB

H Dashboard Horizon 3

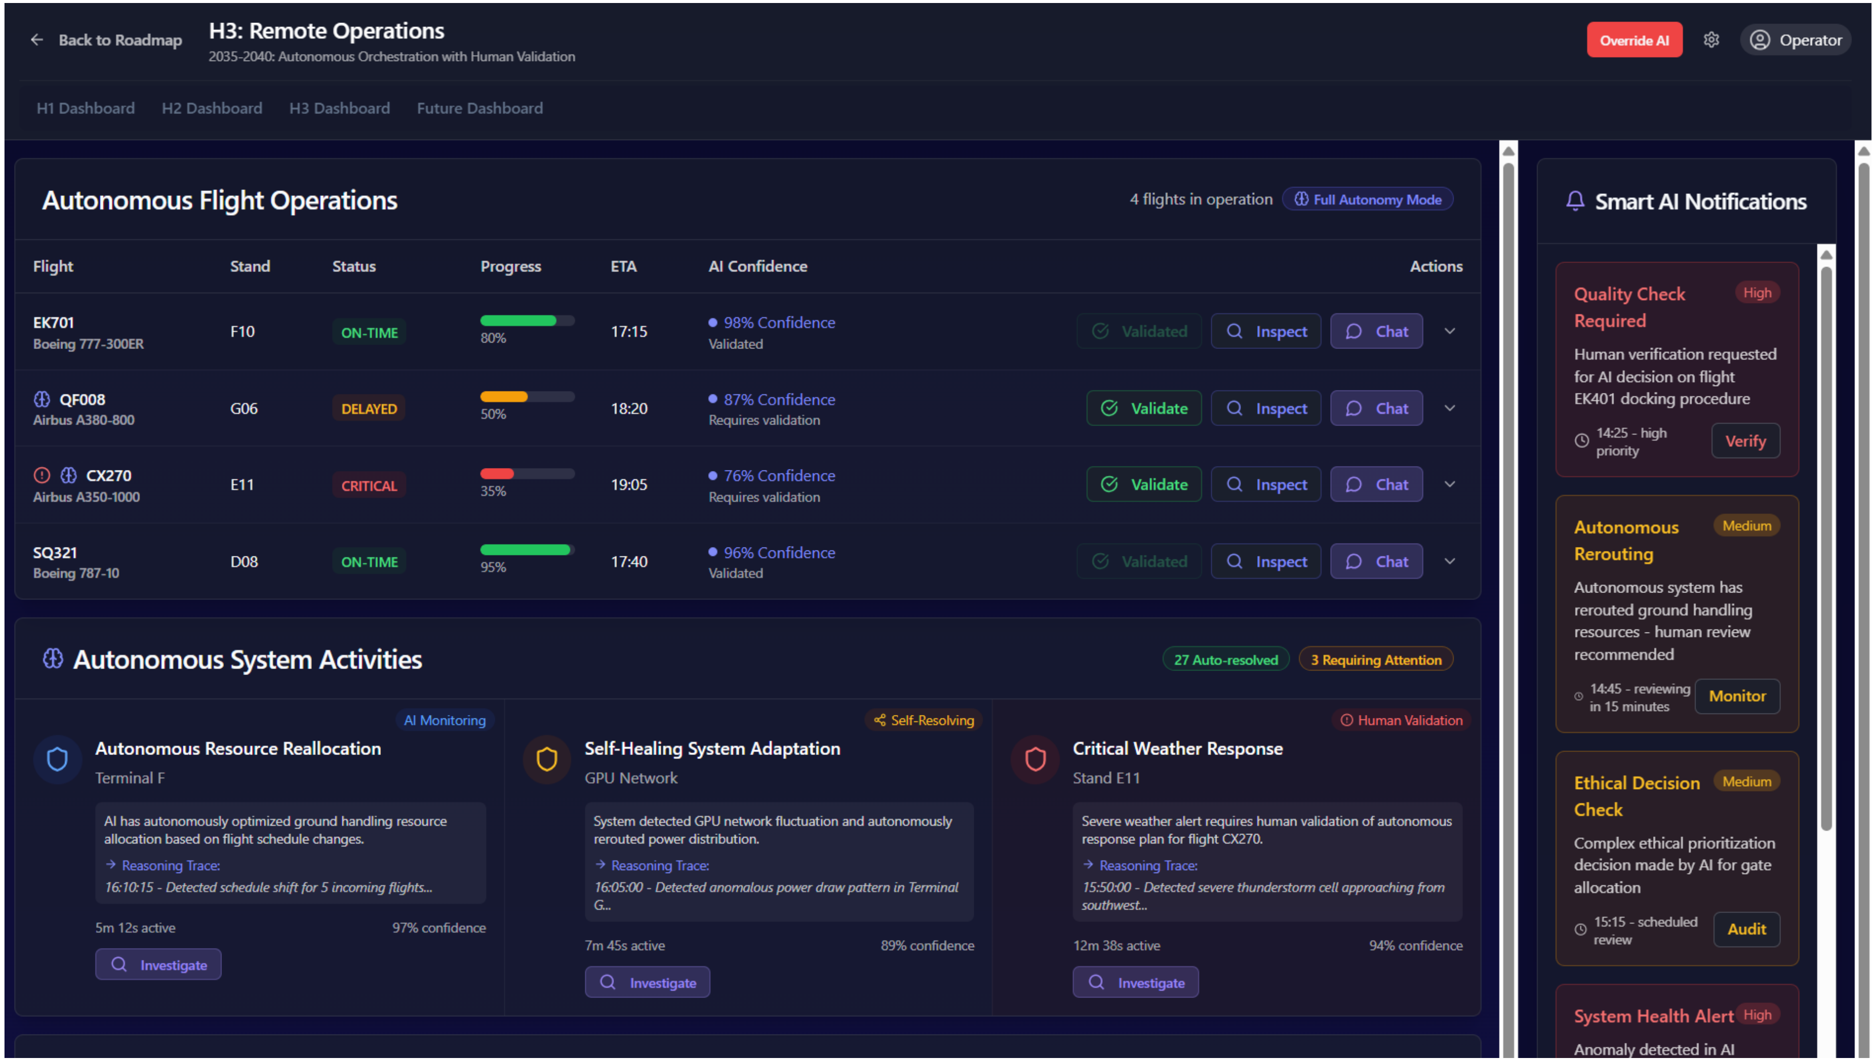

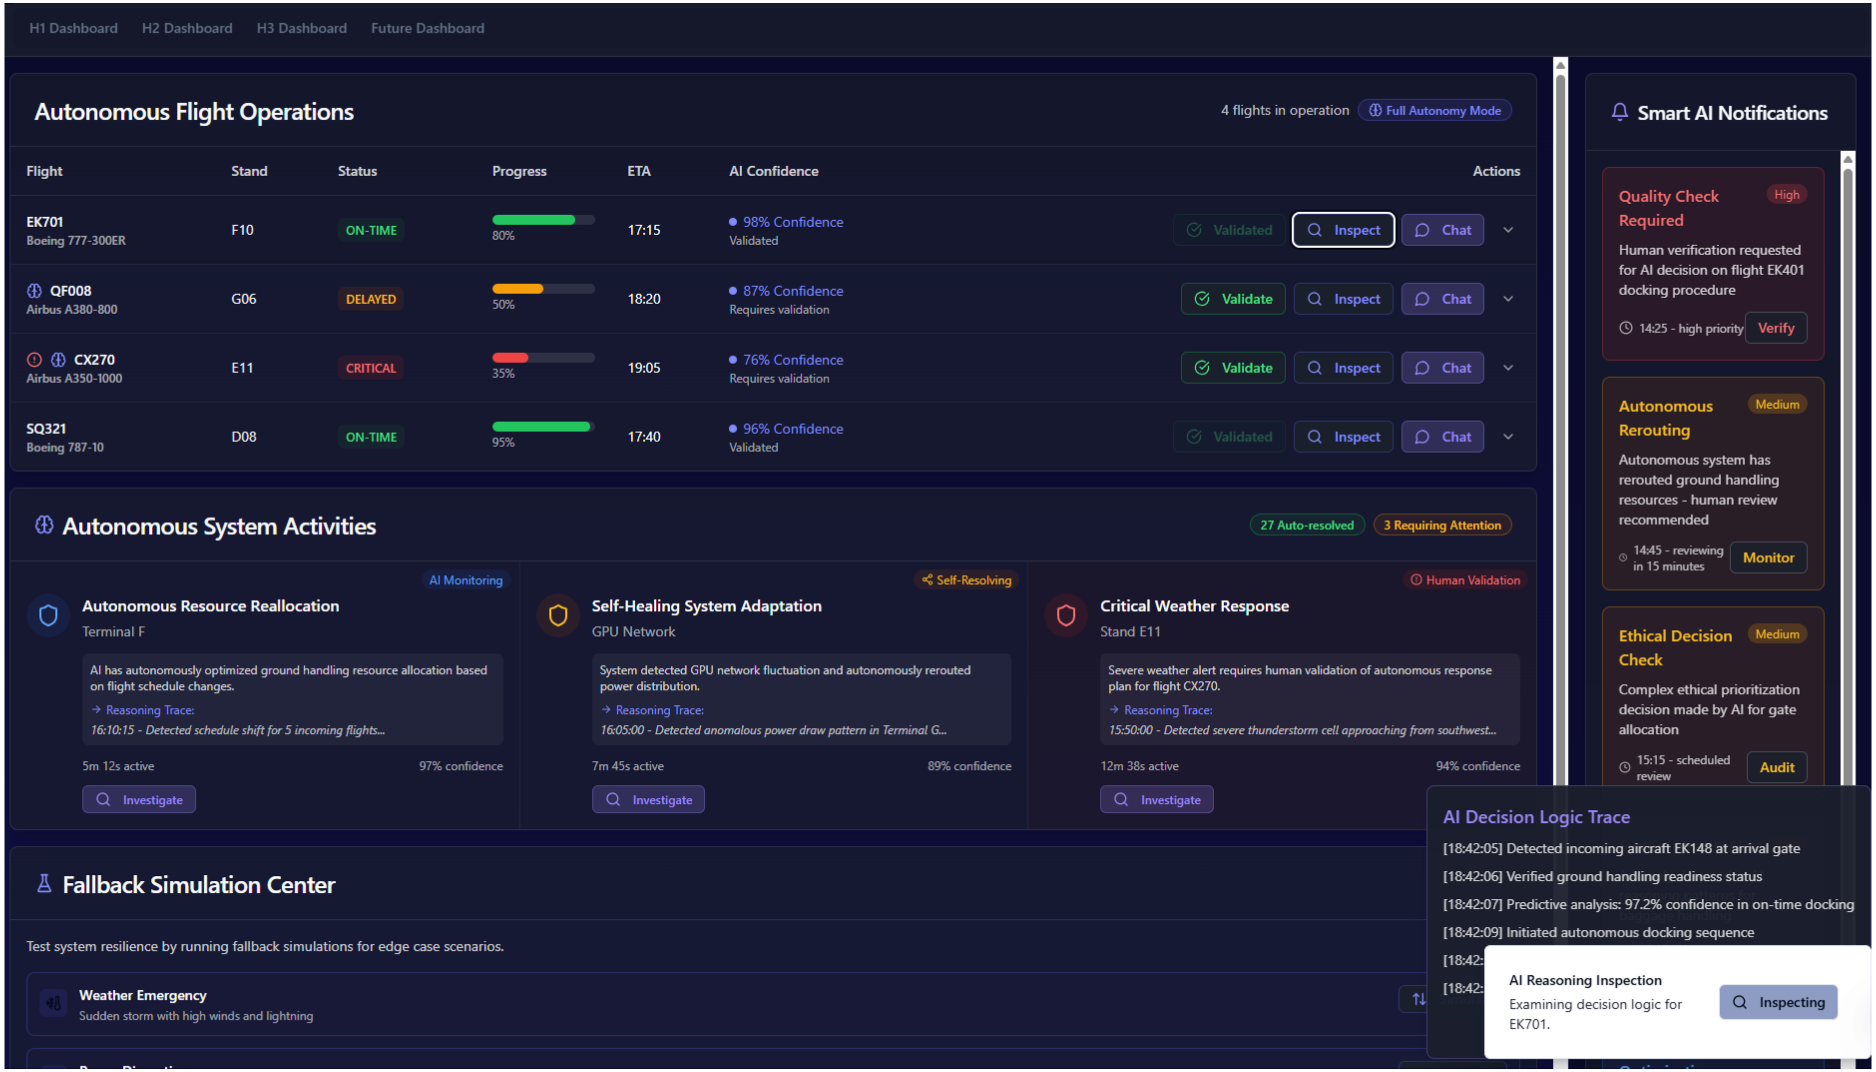
H1 Dashboard
H2 Dashboard
H3 Dashboard
Future Dashboard
Autonomous Flight Operations
4 flights in operation
Full Autonomy Mode
Flight
Stand
Status
Progress
ETA
AI Confidence
Actions
EK701
Boeing 777-300ER
F10
ON-TIME
80%
17:15
98% Confidence
Validated
Validated
Inspect
Chat
QF008
Airbus A380-800
G06
DELAYED
50%
18:20
87% Confidence
Requires validation
Validate
Inspect
Chat
CX270
Airbus A350-1000
E11
CRITICAL
35%
19:05
76% Confidence
Requires validation
Validate
Inspect
Chat
SQ321
Boeing 787-10
D08
ON-TIME
95%
17:40
96% Confidence
Validated
Validated
Inspect
Chat
Autonomous System Activities
27 Auto-resolved
3 Requiring Attention
AI Monitoring
Autonomous Resource Reallocation
Terminal F
AI has autonomously optimized ground handling resource allocation based on flight schedule changes.
Reasoning Trace:
16:10:15 - Detected schedule shift for 5 incoming flights...
5m 12s active
97% confidence
Investigate
Self-Resolving
Self-Healing System Adaptation
GPU Network
System detected GPU network fluctuation and autonomously rerouted power distribution.
Reasoning Trace:
16:05:00 - Detected anomalous power draw pattern in Terminal G...
7m 45s active
89% confidence
Investigate
Human Validation
Critical Weather Response
Stand E11
Severe weather alert requires human validation of autonomous response plan for flight CX270.
Reasoning Trace:
15:50:00 - Detected severe thunderstorm cell approaching from southwest...
12m 38s active
94% confidence
Investigate
Fallback Simulation Center
Test system resilience by running fallback simulations for edge case scenarios.
Weather Emergency
Sudden storm with high winds and lightning
Smart AI Notifications
Quality Check Required
High
Human verification requested for AI decision on flight EK401 docking procedure
14:25 - high priority
Verify
Autonomous Rerouting
Medium
Autonomous system has rerouted ground handling resources - human review recommended
14:45 - reviewing in 15 minutes
Monitor
Ethical Decision Check
Medium
Complex ethical prioritization decision made by AI for gate allocation
15:15 - scheduled review
Audit
AI Decision Logic Trace
[18:42:05] Detected incoming aircraft EK148 at arrival gate
[18:42:06] Verified ground handling readiness status
[18:42:07] Predictive analysis: 97.2% confidence in on-time docking
[18:42:09] Initiated autonomous docking sequence
AI Reasoning Inspection
Examining decision logic for EK701.
Inspecting

Future Vision Dashboard

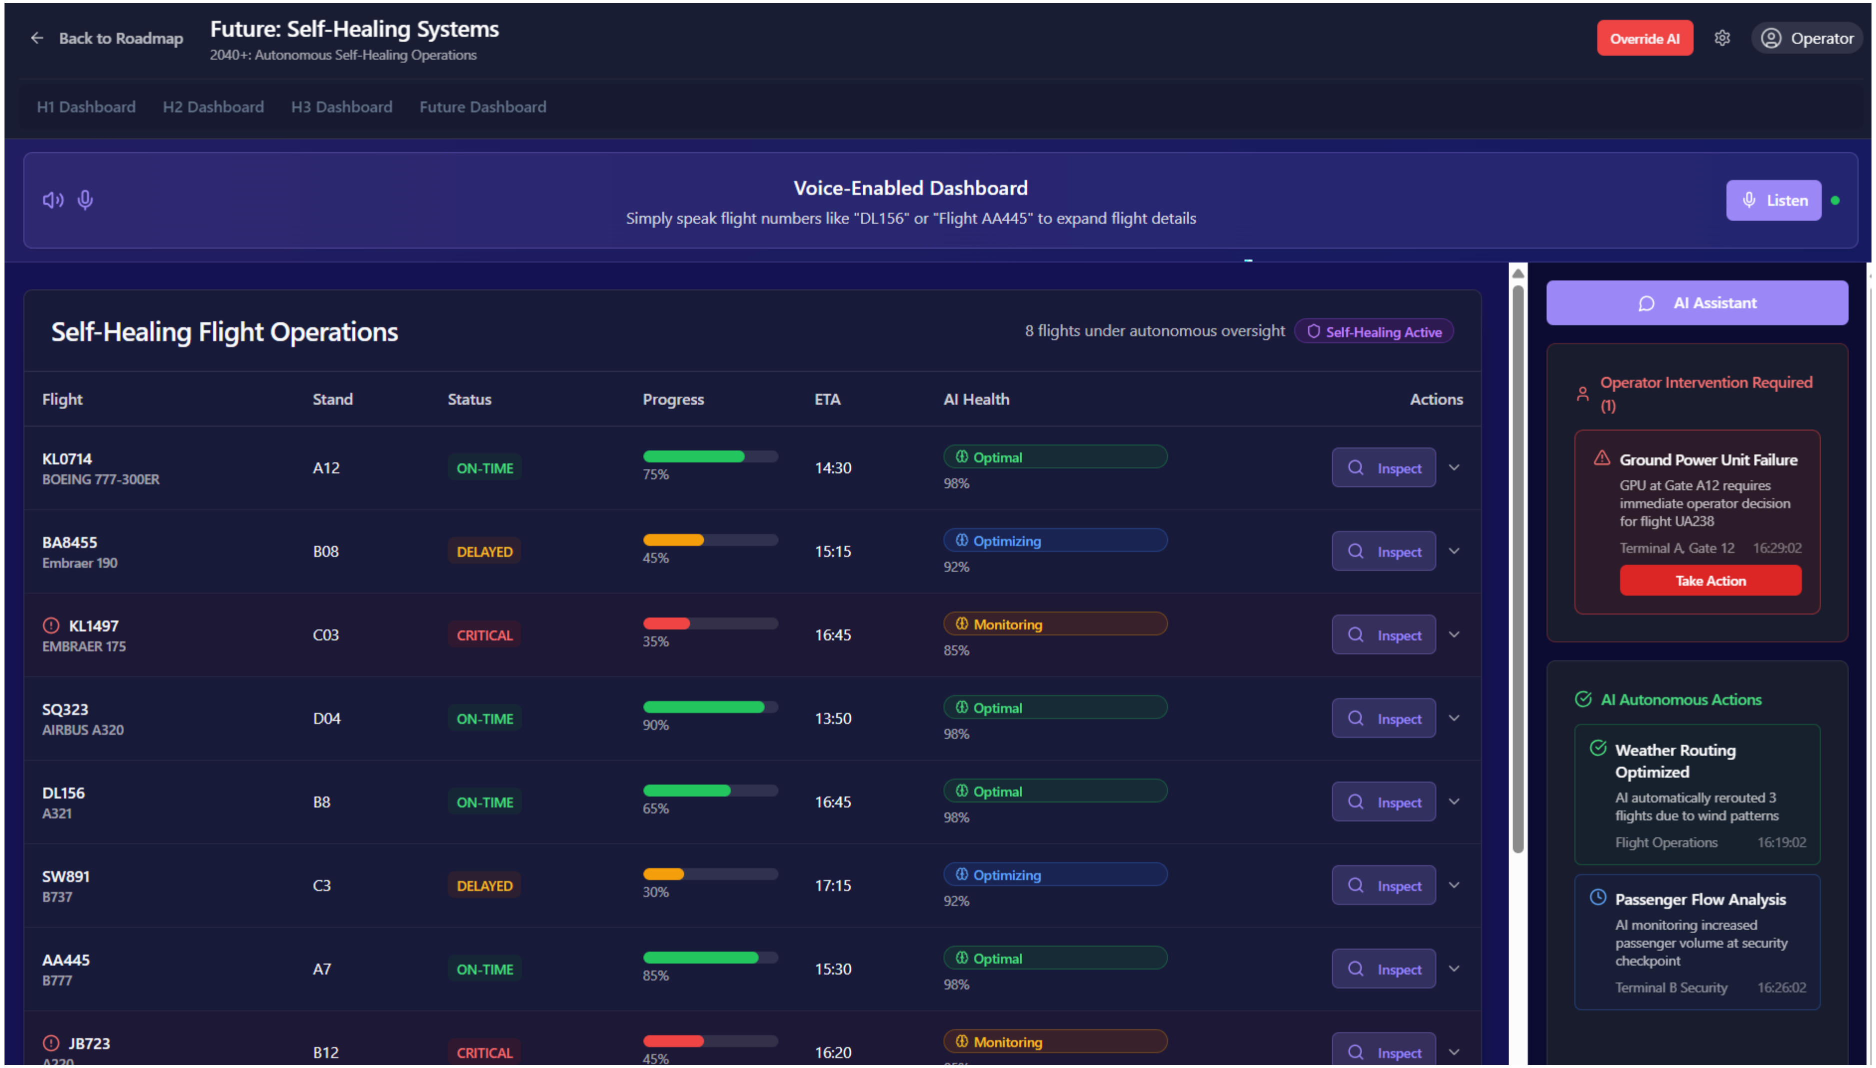


Voice-Enabled Dashboard
Simply speak flight numbers like "DL156" or "Flight AA445" to expand flight details
Listen
Self-Healing Flight Operations
8 flights under autonomous oversight
Self-Healing Active
Flight
Stand
Status
Progress
ETA
AI Health
Actions
KL0714
BOEING 777-300ER
A12
ON-TIME
75%
14:30
Optimal
98%
Inspect
Self-Healing Turnaround Activities
Overall Progress: 76%
Plane landing 14:17 completed
Plane taxiing to gate 14:22 completed
GSE clearance 14:27 completed
FOD Robot 14:32 completed
VDGS Display 14:37 completed
Plane docked 14:42 completed
Wheel Chocks and Cones 14:47 completed
GPU 14:52 in-progress
PBB 14:57 in-progress
Passenger bridge activated 15:02 pending
Passenger bridge connected 15:07 pending
Recent AI Decisions
14:23
Optimized turnaround sequence for weather conditions
98%
14:18
Predicted passenger flow and adjusted gate resources
92%
System Health Status
Ground Power
Optimal
Baggage Systems
Auto-Optimized
Fuel Systems
Monitored
Catering
Scheduled
BA8455
Embraer 190
B08
DELAYED
45%
15:15
Optimizing
92%
Inspect
KL1497
EMBRAER 175
C03
CRITICAL
35%
16:45
Monitoring
85%
Inspect
SQ323
AIRBUS A320
D04
ON-TIME
90%
13:50
Optimal
98%
Inspect
AI Assistant
Operator Intervention Required (1)
Ground Power Unit Failure
GPU at Gate A12 requires immediate operator decision for flight UA238
Terminal A, Gate 12
14:12:10
Take Action
AI Autonomous Actions
Weather Routing Optimized
AI automatically rerouted 3 flights due to wind patterns
Flight Operations
14:02:10
Passenger Flow Analysis
AI monitoring increased passenger volume at security checkpoint
Terminal B Security
14:09:10

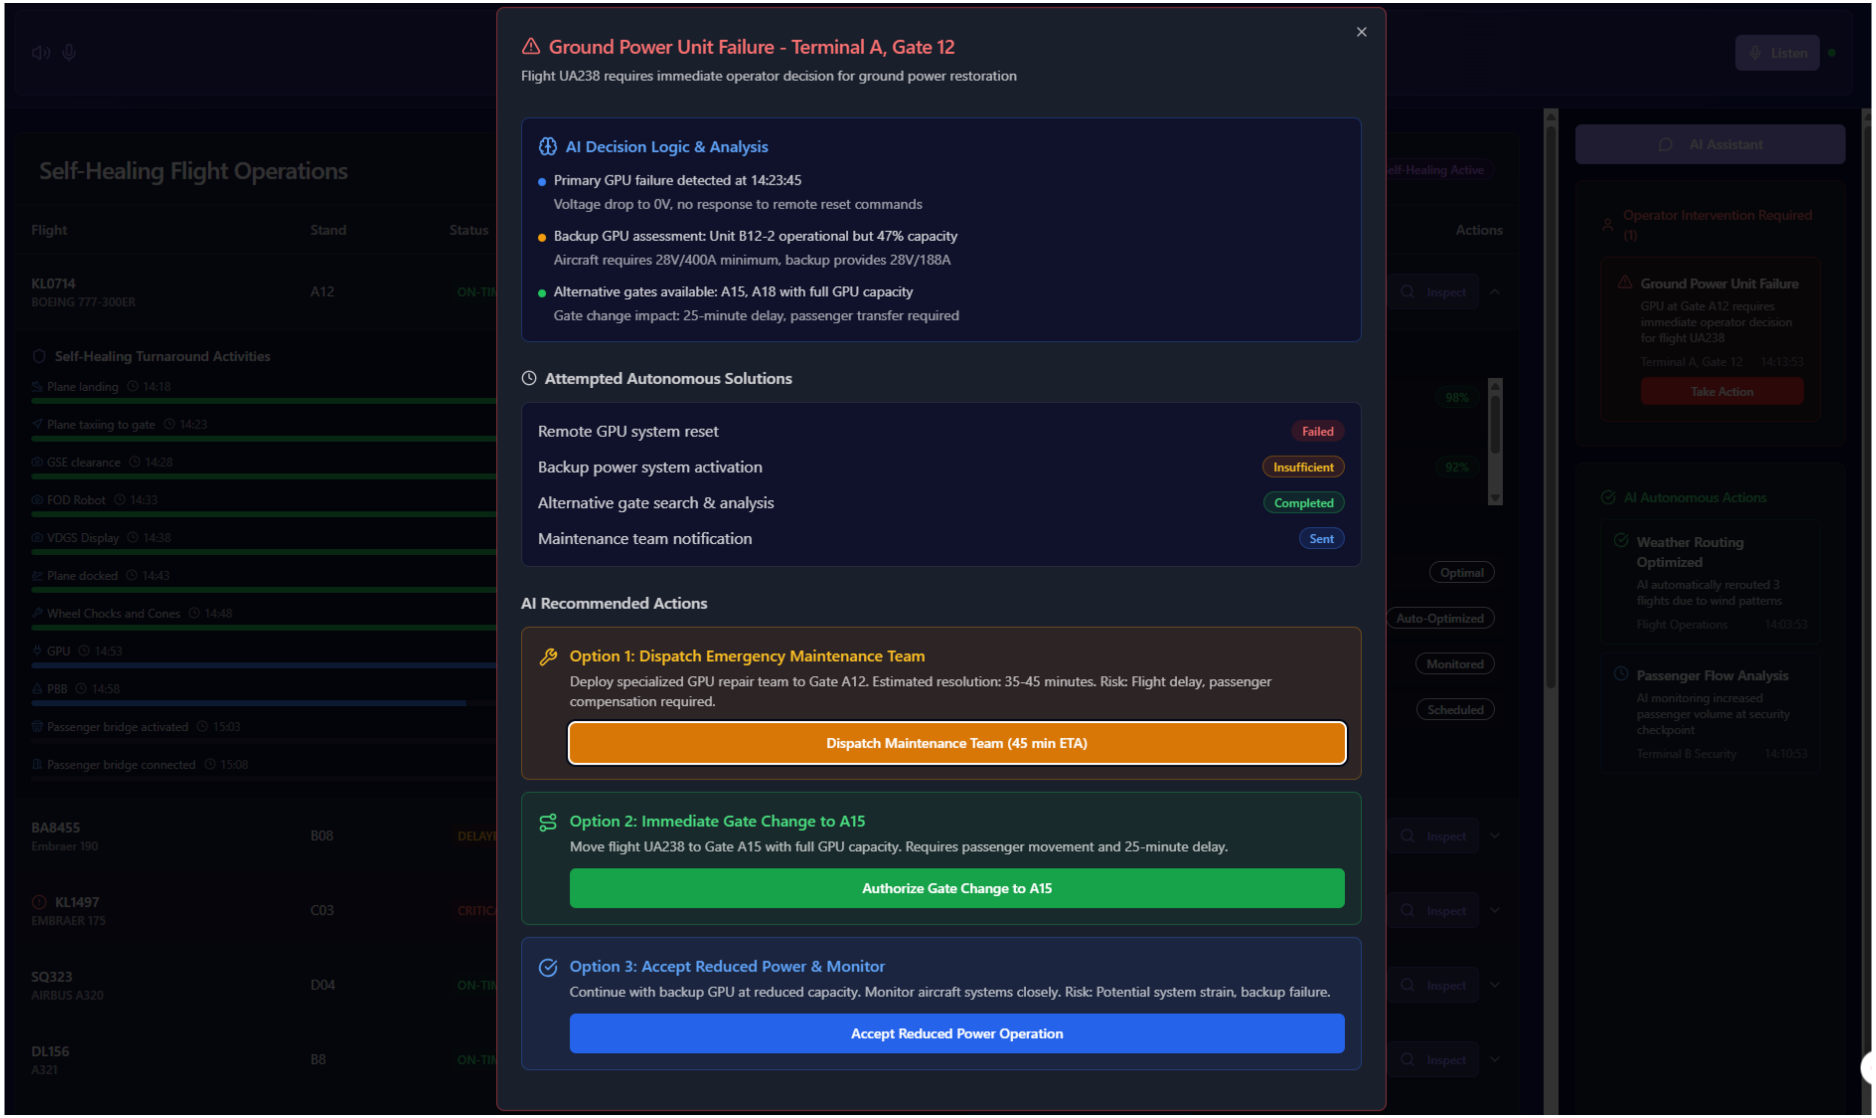
Ground Power Unit Failure - Terminal A, Gate 12
Flight UA238 requires immediate operator decision for ground power restoration
AI Decision Logic & Analysis
Primary GPU failure detected at 14:23:45
Voltage drop to 0V, no response to remote reset commands
Backup GPU assessment: Unit B12-2 operational but 47% capacity
Aircraft requires 28V/400A minimum, backup provides 28V/188A
Alternative gates available: A15, A18 with full GPU capacity
Gate change impact: 25-minute delay, passenger transfer required
Attempted Autonomous Solutions
Remote GPU system reset
Failed
Backup power system activation
Insufficient
Alternative gate search & analysis
Completed
Maintenance team notification
Sent
AI Recommended Actions
Option 1: Dispatch Emergency Maintenance Team
Deploy specialized GPU repair team to Gate A12. Estimated resolution: 35-45 minutes. Risk: Flight delay, passenger compensation required.
Dispatch Maintenance Team (45 min ETA)
Option 2: Immediate Gate Change to A15
Move flight UA238 to Gate A15 with full GPU capacity. Requires passenger movement and 25-minute delay.
Authorize Gate Change to A15
Option 3: Accept Reduced Power & Monitor
Continue with backup GPU at reduced capacity. Monitor aircraft systems closely. Risk: Potential system strain, backup failure.
Accept Reduced Power Operation
Self-Healing Flight Operations
Flight
Stand
Status
Actions
KL0714
BOEING 777-300ER
A12
Self-Healing Turnaround Activities
BA8455
Embraer 190
B08
KL1497
EMBRAER 175
C03
SQ323
AIRBUS A320
D04
DL156
A321
B8
AI Assistant
Operator Intervention Required (1)
Ground Power Unit Failure
Take Action
AI Autonomous Actions
Weather Routing Optimized
Passenger Flow Analysis

J MoSCoW Priorities and Actions

Horizon	Urgency	Priority	Action
H1	Must have	Smart monitoring with AI-enabled video analytics and real-time dashboards,	Deploy AI video systems at select high-traffic gates and integrate dashboards for situational awareness
		Human-machine coordination through remote activation (VDGS, FOD)	Pilot remote activation tools with operator-in-the-loop workflows
		Ecosystem engagement and training alignment	Establish stakeholder workforce for co-development
	Should have	Structured anomaly detection and alerting logic	Define escalation paths and alert thresholds
		Integrated dashboard covering multiple stands	Design and deploy a live dashboard prototype with basic anomaly visualisation
	Could have	Long-term data readiness for AI refinement	Outline data labelling and storage strategy for future model training
		Preliminary automation planning	Document constraints and opportunities for gradual airside automation
H2	Must have	Predictive analytics and contextual alerts	Train forecasting models using Horizon 1 data and integrate predictions into dashboards
		Operator intervention capabilities with AI recommendations	Deploy interfaces that allow operators to accept/reject AI-generated action suggestions
		Expanded remote task execution	Scale up remote activations to include multiple turnaround functions
	Should have	System interoperability with airport and third-party systems	Define data and API standards for integration with CISS and equipment control
		Operator override protocols for hybrid decision-making	Develop rules for manual overrides and automated handover during exceptions
	Could have	Fallback logic for partial autonomy	Introduce early-stage fallback scenarios and test operational impact
		Trend analysis for persistent anomalies	Begin long-term logging and classification of recurring event patterns
H3	Must have	Delegated AI decision-making with human override thresholds	Define and enforce system autonomy boundaries across orchestration tasks
		Automated fallback and error recovery	Implement tiered fallback strategies that escalate only when needed
		Operator role transition to orchestration governance	Redefine KPIs and responsibilities to support supervisory orchestration
	Should have	Shadow-mode simulations to test AI reliability	Run full AI orchestration in parallel with human oversight and compare performance
		Transparent diagnostics for trust and accountability	Develop audit logs and explanation tools for AI decision-making
	Could have	Advanced metrics for orchestration effectiveness	Create analytic models to measure efficiency, stability, and error handling
Future Vision	Must have	Self-healing orchestration with minimal human input	Develop fully autonomous orchestration logic with fallback AI and anomaly handling
	Should have	Operator shift to 'governer'	Define training, audit and oversight roles for exception-based intervention
	Could have	Adaptive orchestration ethics framework	Establish ethical guidelines for decision transparency and accountability

K Implementation Priorities and Action Points

To ensure effective and scalable deployment of autonomous airside operations, several components of the roadmap demand differentiated prioritisation based on their function, criticality, and timing. Priority elements are those that:

- Create foundational infrastructure needed in later stages,
- Impact many downstream activities,
- Or serve as early trust- or capability-building blocks.

Horizon 1 Priorities: Foundational Monitoring and Engagement System Logic Overview:

In Horizon 1, the primary focus lies in deploying smart monitoring systems and progress-tracking dashboards. Operators supervise operations remotely via AI-enabled video analytics and receive rule-based alerts that prompt them to activate functions like VDGS or FOD. While oversight spans multiple turnarounds, operators cannot directly intervene, they rely on communication with on-site crew (shift leaders or team coordinators) to address anomalies detected through delayed or missing events. Forecasting tools, similar to current APOC models, support situational anticipation.

Critical Priorities:

- **Deployment of AI-enabled turnaround monitoring systems:** This includes smart video analytics and integration with operational dashboards. These are the "eyes" of the orchestration layer and form the backbone of real-time decision support.
- **Remote process activation tooling (e.g., VDGS, FOD clearance):** This marks one of the first remotely executable ground operations and creates a test case for human-in-the-loop safety assurance and intervention design.
- **Stakeholder engagement and co-design:** Early involvement is essential for process adaptation, training needs, and trust in automation. This has foundational effects for all subsequent phases.

Important Priorities

- **Initial development of anomaly classification and alert thresholds:** Provides the basis for structured, explainable system notifications, increasing operator effectiveness.
- **Live dashboard design for multi-stand awareness:** Offers a centralised tool for visualising multiple concurrent turnarounds and supports early situational judgment.

Less Urgent Priorities

- **Advanced data labelling and long-term storage protocols:** While critical for future AI model refinement, these can evolve over time.
- **Full airside operation automation:** Requires long-term orchestration maturity and broad ecosystem readiness.

Actions:

- **Deploy AI video analytics on selected gates:** Begin with high-volume gates to test performance and gather baseline data.
- **Set up an integrated live dashboard:** Consolidates real-time visual, operational, and alert data in one interface.
- **Form a multi-stakeholder working group for interface and workflow design:** Ensures tools reflect diverse operational needs.

- **Run workshops to define anomaly escalation and communication channels:** Creates shared understanding and trust in alert logic.
- **Pilot remote VDGS activation:** Trials the supervised execution of core remote tasks under operational conditions.
- **Build operator training modules:** Develop competence and familiarity with new tools and workflows.

Horizon 2 Priorities: Assisted Orchestration and Predictive Coordination

System Logic Overview: In Horizon 2, orchestration becomes more proactive. Operators begin receiving condition-based alerts, contextual recommendations, and predictive analytics from the dashboard. The system suggests interventions based on real-time and forecasted data, with the operator validating or approving actions. New remote control functions (e.g. remote activations) are introduced, enabling further remote intervention without needing on-site execution. Coordination with ground staff remains vital for manual steps.

Critical Priorities

- **Predictive analytics and contextual alerting:** Enables foresight into upcoming delays or conflicts, improving orchestration foresight and response time.
- **Operator intervention capabilities:** Grants human operators limited control over automated systems, essential for building trust and enabling hybrid workflows.
- **Remote activation expansion:** Supports efficiency by enabling the operator to perform more remote tasks safely and effectively.

Important Priorities

- **Interoperability standards for system connectivity:** Facilitates seamless data sharing across equipment, IT, and scheduling systems.
- **Definition of remote override protocols:** Clarifies when and how operators can or must intervene in automated processes.

Less Urgent Priorities

- **Gradual development of fallback logic:** Lays groundwork for resilience but is iterative and can co-evolve with deployment.
- **Deeper analytics on anomaly recurrence and root causes:** Supports long-term process optimisation but is not immediately required for deployment.

Actions:

- **Train predictive models on Horizon 1 data:** Use historical monitoring data to develop delay and anomaly prediction algorithms.
- **Pilot recommendation interfaces for remote sequencing decisions:** Allow operators to test AI-suggested interventions in low-risk contexts.
- **Scale up remote intervention capabilities:** Add more remote activation functionalities to increase operator leverage.
- **Expand dashboard with delay forecasting layers:** Build integrated awareness of current status and future projections.
- **Establish early fallback protocols and handover thresholds:** Define escalation paths and decision authority between system and operator.
- **Monitor operator trust and role adaptation:** Track how staff engage with AI suggestions and intervene during anomalies.

Horizon 3 Priorities: Semi-Autonomous Sequencing and Scalable Resilience System

Logic Overview: In Horizon 3, the orchestration system operates with partial autonomy and dynamic fallback mechanisms. Processes are automatically sequenced. When anomalies occur, the AI attempts resolution independently and only involves the operator for validation or edge-case handling. The operator is either prompted to investigate or simply informed of the actions taken and their outcomes. Thresholds must be defined to determine which decisions the AI can execute autonomously versus when operator oversight is required.

Critical Priorities

- **Threshold-based decision delegation:** Defines when AI can act independently versus when human approval is needed, crucial for scalable autonomy.
- **Dynamic fallback orchestration:** Enables safe recovery from disruptions without human input unless escalation is triggered.
- **Operator role redefinition:** Supports the shift from manual control to supervisory orchestration through new responsibilities and KPIs.

Important Priorities

- **Shadow-mode simulation to build confidence:** Runs parallel AI and human control to assess reliability and identify edge-case failures.
- **System auditability and diagnostic transparency:** Ensures explainability of AI decisions and post-event investigation.

Less Urgent Priorities

- **Full robotisation in controlled areas:** Only feasible once orchestration and fallback systems prove robust in semi-automated zones.
- **Development of advanced metrics for orchestration efficiency:** Enables fine-tuned performance evaluation but can follow initial rollout.

Actions:

- **Simulate full orchestration in shadow mode:** Allow systems to operate in parallel with human oversight to validate autonomy.
- **Define delegation thresholds for system autonomy:** Set clear boundaries between human and AI decision rights.
- **Deploy fallback strategies with multi-tier logic:** Introduce stepwise fallback procedures and escalation conditions.
- **Introduce operator KPIs and orchestration governance roles:** Align performance indicators with new orchestration tasks.

Future Vision:

System Logic: The Future Vision envisions fully autonomous orchestration. All tasks are executed in sequence, anomalies are handled through self-healing AI and multi-tier fallback strategies, and operator input is limited to rare edge cases where system resolutions are insufficient. Recommendations are provided as needed, but human intervention is the exception rather than the rule.