iMERMAID

D5.3: Roadmap

Task 5.4: Creating roadmap/action plan

WP5: Multistakeholder roadmaps for uptake and scalability of the innovative solutions

Author: Antti-Jussi Tahvanainen, Emma Peltonen, Maaria Nuutinen & Arto Wallin (VTT)

Date: 31 March 2026

Funded by the European Union

GRANT AGREEMENT NUMBER	101112824
ACRONYM/ FULL TITLE	iMERMAID/ Innovative solutions for Mediterranean Ecosystem Remediation via Monitoring and decontamination from Chemical Pollution
START DATE	01 June 2023
END DATE	31 May 2026
PROJECT URL	www.imermaid.eu
DELIVERABLE TITLE	Roadmap for replicating and scaling up the demonstrated solutions
WORK PACKAGE	WP5 – Multistakeholder roadmaps for uptake and scalability of the innovative solutions
CONTRACTUAL DATE OF DELIVERY	31/3/2026
ACTUAL DATE OF DELIVERY	M34 (31 March 2026)
NATURE	Report (R)
DISSEMINATION LEVEL	Public (PU)
LEAD BENEFICIARY	VTT
RESPONSIBLE AUTHOR	Antti-Jussi Tahvanainen, VTT
CONTRIBUTORS	Emma Peltonen, Arto Wallin, Maaria Nuutinen, Antti-Jussi Tahvanainen, VTT; IRIS; WF; EDEN; AIG; APCL
ABSTRACT	This deliverable presents a generalised roadmap for replicating and scaling the monitoring and remediation solutions developed in iMERMAID, integrating scenario-driven roadmapping with system-dynamics modelling and co-creation with three participating companies. The process combined company-specific futures (built via morphological analysis and AI-assisted consistency checks) with strategy workshops that distilled scenario-independent opportunities, threats, and strategic measures, followed by a cross-company synthesis into a single, time-sequenced plan aligned with iMERMAID's aims. The roadmap emphasises modularity and phased commissioning, evidence and certification as currency, mixture-robust/low-waste performance, and service-centric deployment models that link revenues to verified outcomes, each anchored to risks and leverage points identified by the modelling in Task 5.3.

	A three-horizon structure translates insight into action: 0–1 years to prove, package and protect (productise modularity; hard-wire monitoring and pilot-to-operation conversion; hedge delivery risks); 1–3 years to scale, certify and service (partnership-first market access; outcome-linked service; quality-assured manufacturing; collective compliance tooling); and 3+ years to serialise, integrate and diversify (framework-ready retrofit kits; regional manufacturing/service; integration with incumbents; institutionalised recovery and traceability). The roadmap's priorities and sequencing follow the Task 5.3 models' high-leverage variables—installation delays, evidence gaps, and threshold effects in certification/traceability—so that small structural changes shift adoption trajectories system-wide. The result is a living, policy-aligned plan that can be updated iteratively and used by technology developers, utilities, regulators, financiers, and integrators to convert demonstration results into replicable, bankable deployment pathways across diverse operating contexts.

Document Revision Log

Version	Description of changes	Author	Role in the project
0.1	Task description, headings and bullets	Maaria Nuutinen (VTT)	Beneficiary
0.2	First full draft	Antti-Jussi Tahvanainen (VTT)	Beneficiary
0.3	Internal revision draft	Marcela Permann-Doubková (AIG), Maria Zafeiropoulou (APCL), Arto Wallin (VTT), WF, IRIS & EDEN	Beneficiary
0.4	Final draft ready for submission	Antti-Jussi Tahvanainen, Maaria Nuutinen	Beneficiary
1.0	Final review and formatting	Rodrigo Sedano, ITCL	Project Coordinator

Executive Summary

This report provides a stakeholder-ready roadmap for replicating and scaling iMERMAID solutions, built from company-level scenario work, structured strategy design, and Task 5.3's system-dynamics analysis. It is written for the organisations that will adopt, integrate, regulate, finance, or operate the solutions—technology developers, utilities, industrial users, regulators and public laboratories, system integrators, financiers/insurers, and circular-economy off-takers. The scope and structure follow the project's Task 5.4 mandate: to convert alternatives assessment, foresight, and modelling insight into a normative, actionable plan with milestones and preconditions, and to do so in a way that remains iterative as conditions evolve.

Three steps underpinned the analysis. First, company-specific futures were developed using morphological scenario analysis supported by AI-assisted consistency checks, and transformed into narratives that clarified operating logics and strategic stakes. Second, co-creative strategy sessions mapped opportunities and threats in each future, then isolated those that generalise across futures and translated them into strategic measures. Third, a qualitative cross-company synthesis reconciled differences at the level of strategic intent, resulting in a consolidated measure set that was aligned with insights emerging from other iMERMAID work packages. The entire pipeline was then read against Task 5.3's models to expose dependencies, delays, thresholds, and leverage points that determine whether adoption compounds or stalls.

Across the full material, the most durable opportunities concentrate where buyers can deploy modular upgrades fast, prove performance continuously, and keep operations stable under variability and mixtures; the most persistent threats arise where procurement serialises and favours scale, where verification shifts liability onto suppliers, where administrative fragmentation adds rework and delay, and where supply-chain/production bottlenecks turn urgent tenders into missed windows. In this configuration, evidence becomes currency, low-waste residuals become a differentiator, and traceable recovery becomes resilience only when quality classes and product status are clear.

The roadmap's centre of gravity is practical: modular systems commissioned fast, evidenced continuously, specialised where selectivity and circularity create value, and supported by service and financing models that turn proof into revenue. This thesis is not opinion; it follows the Task 5.3 models, which show that adoption hinges on reducing installation delays, institutionalising evidence loops that close the gap between desired and realised quality, and crossing thresholds in certification/traceability that unlock compounding effects (e.g., offtake acceptance for recovered outputs).

The roadmap translates insights into three time windows with clear intent and roles:

0–1 years: Prove, package, and protect: Finalise productised modularity and standard tie-ins to shorten commissioning; make operational transparency and lab-aligned monitoring part of the offer so approvals, finance, and service models price against verified outcomes; require pilot-to-operation terms before any pilot to stop "pilot-to-nowhere"; and hedge delivery risk via qualified suppliers, small strategic inventories, and regional assembly options. These moves change system structure quickly by increasing installation inflows and replacing uncertainty with auditable performance, which the models indicate is the fastest path to shift adoption trajectories.

1–3 years: Scale, certify, and service: Make partnership-first access the default (utilities and integrators), professionalise outcome-linked service with insurer-aligned documentation and predictive maintenance, scale manufacturing with experienced partners and harmonised QA to avoid quality drift, and amplify low-waste and recovery value with traceability and emerging standards so circular value clears procurement. At the same time, reduce administrative drag through collective compliance packs aligned

with risk-management specifications for reuse. These steps transform scattered wins into a programmatic roll-out and dampen oscillations from price and permit variability.

3+ years: Serialise, integrate, and diversify: Embed serialised retrofit kits into framework tenders, regionalise manufacturing and service to meet surge windows reliably, integrate with incumbents as a specialist module provider to avoid channel exclusion, and institutionalise recovery with quality-assured outputs and offtakes aligned to circular-materials benchmarks. The long-term posture preserves distinctiveness as the market consolidates around platforms and frameworks.

The roadmap is a living instrument. A cross-partner coordination function should own the update cadence, with clear roles for evidence stewardship, strategic integration, and decision authority. Update triggers should be tied to threshold signals—regulatory enactments affecting desired quality, certification/traceability milestones, performance drifts that widen the lab–sensor gap, or surge windows that justify regional duplication—so revisions respond to observed system behaviour rather than opinion. This mirrors the Task 5.3 logic that small changes at the right leverage points (delays, evidence, thresholds) shift many downstream decisions at once.

If implemented as sequenced, the outcome is a replicable, bankable pattern: fast modular retrofits that meet compressed windows, continuous evidence that lowers approval and financing friction, service discipline that makes outcome-linked revenues safe, manufacturing capacity that keeps promises during surges, and traceable circular outputs that move from engineering wins to tradable value. In short, it is a credible path from pilots to serialised deployment across contexts where policy is tightening, water stress is recurrent, and stakeholders increasingly pay for what can be proven.

Note to the reader

This is a milestone report related to Task 5.4:

T5.4 Creating roadmap/action plan (M1-M6); [Lead: VTT; Participants: SOCAMEX, ITCL, HP, ECSA, UNIFI, IRIS, HCMR, APCL, ZEN, ENIG, UA, CMMI, WF, EDEN]:

The goal of the task is to create a roadmap and action plan for replicating and scaling up the demonstrated solutions. The roadmap builds the assessment of alternatives and insights from foresight and SD modelling. The aim is to devise an actionable roadmap for relevant stakeholders who want to commercialise developed solutions. The roadmap will be a normative guideline for scaling up, identifying crucial milestones and preconditions which must be met to meet the criteria of building viable business. However, since the future is highly uncertain, the roadmap follow-up is designed to be an iterative process. The roadmap is structured in a way that enables responding to market feedback and unforeseeable events as effectively as possible.

Table of Contents

List of Figures

List of Tables

Acronyms

AI	Artificial Intelligence
AOP	Advanced Oxidation Process
APCL	Anaptixiaki Pliroforiki Crete SA
AIG	Aristotle Institute of Geosciences
CAPEX	Capital Expenditures
CECs / CoEC / CoECs	Contaminants of Emerging Concern
CMMI	Cyprus Marine & Maritime Institute
CRMA	Critical Raw Materials Act
CSR	Corporate Social Responsibility
DWD	Drinking Water Directive
EC / EU	European Commission / European Union
ECSA	European Cluster for Sustainable Aquaculture
EPR	Extended Producer Responsibility
GA	Grant Agreement
HCMR	Hellenic Centre for Marine Research
HP	Hellenic Petroleum
ICT	Information and Communication Technology
IRIS	IRIS srl
KOLs	Key Opinion Leaders
LP	Leverage Point
MSFD	Marine Strategy Framework Directive
Miro	Miro online collaboration platform
OPEX	Operational Expenditures
OEM	Original Equipment Manufacturer
PESTEL	Political, Economic, Social, Technological, Environmental, Legal
PFAS	Per- and Polyfluoroalkyl Substances
PDP	Pulsed Discharge Plasma
POPs	Persistent Organic Pollutants

PaaS	Product-as-a-Service
R&D	Research and Development
REACH	Registration, Evaluation, Authorisation and Restriction of Chemicals
ROI	Return on Investment
SD modelling / SD model / SD	System Dynamics Modelling / System Dynamics
SME / SMEs	Small and Medium-Sized Enterprises
SOCAMEX	Socamex S.A.
T	Task
TRL	Technology Readiness Level
UA	University of Alicante
UWWTD	Urban Wastewater Treatment Directive
VTT	VTT Technical Research Centre of Finland Ltd.
WF	Weeefiner
WP	Work Package

1.0 Introduction

1.1 Purpose and scope

This deliverable presents the roadmap and action plan prepared under Task 5.4 of Work Package 5, which aims to guide the replication and scale-up of the solutions demonstrated in iMERMAID. The task focuses on translating the project's technical, methodological, and analytical findings into a structured pathway that supports successful deployment in diverse contexts. The scope of the deliverable follows the task description in the project proposal, which states that Task 5.4 produces a normative roadmap based on the assessment of alternatives, foresight, and the insights of system dynamics modelling, and identifies the milestones and preconditions necessary for building viable business around the iMERMAID solutions. The roadmap is planned as an iterative mechanism that can react to changing market conditions and unexpected events.

1.2 Target stakeholders and intended use

The roadmap is intended for stakeholders who will consider adopting, integrating, or commercialising the solutions developed in iMERMAID. According to the proposal, the project will demonstrate monitoring and remediation solutions for chemical pollution across different stages of the water system, from sources such as agriculture and pharmaceuticals to wastewater treatment plants and marine environments. This implies a stakeholder group that spans technology developers, water utilities, industrial users, regulators, policy bodies, and other actors who operate in regions exposed to contaminants of emerging concern or responsible for their management. The roadmap supports these actors by providing practical direction for commercial replication and scale-up, informed by evidence and learning generated throughout the project. It also contributes to broader efforts to reinforce the uptake of the Marine Strategy Framework Directive and the Water Framework Directive, which shape regulatory expectations for chemical pollution monitoring and reduction in the Mediterranean basin.

1.3 Relation to other project results

The roadmap builds directly on the analytical and strategic components of Work Package 5, as described in the proposal. The assessment of chemical and non-chemical alternatives under Task 5.1, the foresight activities under Task 5.2, and the system dynamics modelling under Task 5.3 provide the foundation for identifying pathways, strategic measures, and conditions that support scale-up. Task 5.4 integrates the outputs of these tasks into a single coherent plan for future deployment. The roadmap also draws on evidence generated across the project, including the monitoring technologies developed in Work Package 2, the remediation solutions advanced in Work Package 3, and the pilot demonstrations carried out in Work Package 4, which validate performance and applicability in real settings across the Mediterranean region. Finally, the established roadmap is an effective vehicle for communicating for action across stakeholder groups, rendering the roadmap a tool for disseminating results in the dissemination task of the project. These links ensure that the roadmap reflects both the technical feasibility and the operational environments of the iMERMAID solutions.

1.4 Deliverable structure

The deliverable follows the structure given in the report template. After the introduction, the methodology and process chapter presents the analytical and procedural foundations that guided the roadmap creation. The subsequent chapter describes the demonstrated solutions and the broader context for their replication and scale-up. The scenario chapter presents future developments that influence deployment prospects. The roadmap chapter outlines the sequence of actions, milestones, and responsible actors. The final chapter sets out the mechanism for iterative updates, which follows the requirement in the proposal that the roadmap responds to market feedback and emerging conditions, and concludes with recommendations derived from the analysis.

2.0 Methodology and process

2.1 Methodology

2.1.1 Framework

This project applies a scenario-driven roadmapping framework inspired by Siebelink, Halman and Hofman (2016), adapted to the objectives of T 5.4 and the solution landscape of iMERMAID. Their work addresses a gap in conventional roadmapping: standard methods struggle with uncertainty and often sacrifice the communicative and directive strengths of a single roadmap when they try to accommodate multiple futures. Scenario-driven roadmapping tackles this by first exploring distinct future environments and then distilling their shared implications into one robust roadmap that preserves clarity while coping with uncertainty. In their demonstration, the approach reduced environmental uncertainty and supported prioritised innovation choices without multiplying documents or fragmenting guidance.

Siebelink, Halman and Hofman's framework rests on three principles that are material to iMERMAID. First, uncertainty is not noise around a forecast but a structural feature that spans state, effect and response dimensions. A roadmap that ignores these dimensions risks brittleness when conditions change. Scenario-driven roadmapping treats uncertainty as a design input and tests strategic direction against a spread of external developments before selecting the most robust areas to anchor the roadmap. Second, strategic communication matters as much as analytical completeness. Using a single roadmap avoids the diffusion of attention that occurs when each scenario spawns its own plan, which earlier integrations of scenarios and roadmaps have struggled to prevent. The value is a document that is actionable for decision-makers and legible to operational teams. Third, robustness is earned through disciplined synthesis. Common implications across scenarios identify focus areas and preconditions that hold across divergent futures, allowing a firm or consortium to commit to milestones while reserving options on elements that remain contingent.

The academic record supports this orientation. Prior work combining scenarios and roadmapping shows that scenarios broaden the horizon and expose gaps in "business-as-usual" plans, but if left inside the roadmap they can undermine communicative clarity through parallel trajectories or multiple competing roadmaps. The case of clean production in metal manufacturing illustrates how scenarios set context and stress-test portfolios, while the resulting roadmap remains a single artefact for action (Sarita & Aylen, 2010). This strengthens portfolio design and policy linkages without overcomplicating delivery. Similarly, later methodological contributions argue for scenario-driven roadmapping that uses scenarios first to define plausible external states and then converts them into a roadmap with clear transition signals, helping organisations connect today's choices with pathway-specific triggers for adjustment (Hussain, Tapinos & Knight, 2017). At the same time, foundational work on technology roadmapping highlights the need to retain the roadmap's role as a visual and communicative device linking markets, products and capabilities over time, and cautions that process complexity should not overwhelm that role (Phaal, Farrukh & Probert, 2004; Probert, Farrukh & Phaal, 2003).

These design choices fit the iMERMAID context. The project faces heterogeneous solution types, regional variation in pressures and permitting, evolving standards under the Water Framework and Marine Strategy Framework Directives, and demand signals that can shift with policy and funding. A scenario-driven approach enables a single, stakeholder-ready roadmap that aligns solution owners, pilot hosts, and adopters, while ensuring that the roadmap remains valid under different regulatory cadences, financing windows and market receptivity. It creates a structure to reduce state uncertainty by scanning drivers that shape water quality regulation and adoption, to reduce effect uncertainty by assessing how those drivers change deployment economics and evidence needs, and to reduce response uncertainty by identifying the strategic measures that perform across futures.

The framework also aligns with broader research that links robust planning to decision quality under deep uncertainty. Studies of risk-aware or scenario-based roadmapping emphasise that most roadmaps fail when they do not integrate uncertainty at the design stage, and propose processes that embed scenarios while keeping ownership, milestones and responsibilities clear (Ilevbare, Probert and Phaal, 2014). This line of work reinforces the need for a single roadmap with explicit decision points, rather than multiple scenario-specific artefacts. In parallel, operations research and foresight literature demonstrate that combining exploratory views of the future with structured roadmapping improves plan resilience and stakeholder alignment and helps organisations manage the trade-off between guidance and flexibility (Zumbrunn, 2023; OECD, 2025; Samson-Onuorah & Bakare, 2025). Finally, scenario planning scholarship underlines the role of scenarios in building strategic flexibility and robust strategies (Lindgren & Bandhold, 2009), which is the intent of T 5.4 when it converts heterogeneous project evidence into an actionable scaling pathway.

In this deliverable, the framework therefore serves a pragmatic purpose. It grounds the roadmap in a defensible academic lineage, yet it is adapted to iMERMAID's multi-solution, multi-region setting so that the output remains a single, clear pathway with defined milestones, preconditions and update rules. The next subsection details the project-specific process that operationalises this framework for T 5.4.

2.1.2 Process

The methodology followed in T 5.4 rests on a structured, scenario-driven roadmapping process inspired by Siebelink, Halman and Hofman (2026), and adapted to the specific needs of the companies participating in iMERMAID. The process combined facilitated expert engagement, AI-enabled analytical support, and iterative refinement. Its purpose was to ensure that the roadmap draws on a wide exploration of possible operating environments, captures implications that persist across uncertainty, and translates these into measures that support future adoption and scaling of the solutions developed in the project.

Building scenarios: The process began with dedicated **scenario-building workshops** held separately with the three commercial companies involved in T 5.4. Each workshop followed the logic of a **morphological scenario analysis** in which participants first identified **change drivers and uncertainty factors** relevant to their business. They then produced **alternative future states** for each factor and created **futures tables** that captured the structure of their operating environment. These discussions were guided by a **facilitated dialogue** that encouraged the participants to **refine and prioritise the factors** as their understanding deepened. All workshops were facilitated by VTT through the **ScenAIrios -platform**, an AI-enabled scenario building system under development at VTT. ScenAIrios widens the range of uncertainty factors and reduces bias, generates richer arrays of alternative outcomes, and conducts consistency assessments that would be difficult to complete manually. Through this, experts could concentrate on **reflection, interpretation and framing of the strategic meaning** of the futures produced. The platform supported the breadth and structure of the morphological analysis, while the workshops ensured that the scenarios retained clear links to each company's context.

After the workshops, the futures tables and the alternative states within them were transferred to ScenAIrios for an **exhaustive consistency analysis**. This allowed VTT to evaluate all combinations of the alternative future states and identify those combinations that form rich, internally **consistent and plausible scenarios**. ScenAIrios applies automated consistency checks at a scale not feasible by hand and enables the construction of scenarios that are diverse across multiple dimensions of the future. VTT then transformed the consistent combinations into **scenario narratives** for each company. These narratives provided a vivid and structured description of future business environments and offered each company a firm basis for strategic analysis.

Identifying opportunities & developing strategic measures: The scenario narratives were sent to the companies for familiarisation before the next stage of the work. Separate **strategy design sessions** were then held with each company on Miro Boards. These sessions followed the principles of the approach proposed by Siebelink, Halman and Hofman, adapted to the purpose of T 5.4. At the start of each session, participants revisited the scenarios and confirmed that no changes were needed. They then identified **business opportunities and threats** for each scenario. This created a broad set of implications grounded in the characteristics of each future environment. Once this pool was complete, the participants examined it to determine which implications appeared **across all scenarios**. These shared implications represent elements of the future that hold independent of how the environment evolves and therefore form the foundation of a robust roadmap. From these implications, the participants developed **strategic measures** to address or exploit them. They distinguished measures that build on **existing strengths** from those that compensate for **internal weaknesses**. After generating these measures, the participants assessed them through facilitated discussion and **prioritised them by their impact and urgency**. The prioritised measures were then placed along a **chronological structure to form a company-specific strategic roadmap** linked to the corresponding opportunities and threats.

Synthesizing results: The outcomes of the workshops were documented, reviewed and approved by each company. Once this material was available for all companies, VTT conducted a **qualitative synthesis**. This synthesis followed established practice in strategic foresight work. The company-specific results were compared to identify common patterns in opportunities, threats and strategic measures. Themes that appeared across companies were grouped to form consolidated categories. Differences in detail were resolved by focusing on the strategic intent of the measures rather than the company-specific formulation. Through iterative analysis, these categories were shaped into a high-level roadmap that represents the collective insights of the three companies. This generalised roadmap was aligned with the broader aims of iMERMAID by linking the findings to the results emerging from the project's other work packages. Through this, the roadmap was enriched with evidence from monitoring technologies, remediation solutions and the demonstration activities, and positioned as a project-level roadmap for replication and scaling.

Advocating for action: The final step of the work was the formulation of conclusions and recommendations. These were based on the aggregated material and the workshop discussions. The conclusions highlight the strategic implications of the scenarios and measures, while the recommendations offer direction for the continuation of work in the project and the potential adoption of the solutions after the project.

2.1.3 Co-creation

Co-creation was a central working principle in T 5.4. The roadmap emerged from a structured collaboration between VTT and the three participating companies, rather than from an external analytical exercise. This approach ensured that the strategic content reflects the lived business realities of the companies and that the roadmap connects directly to organisational priorities, capabilities and constraints. The scenario-building workshops and the strategy sessions formed a shared learning process in which participants shaped the material through discussion, interpretation and sense-making, while VTT facilitated the structure and analytical coherence of the work.

The value of co-creation in this context lies in the nature of scenario-driven roadmapping. The method depends on the identification of uncertainty factors that matter for strategic choices and on the recognition of implications that hold across divergent futures. These elements cannot be produced solely through desk research, since they require knowledge of internal processes, markets and competitive positions that only the companies themselves can provide. Co-creation also helps reduce the risk of producing a roadmap that is detached from the decision environments in which it must operate. The involvement of company representatives in every stage of the work builds shared ownership of the outcomes, which increases the likelihood that the roadmap will inform strategic choices after the project.

The presence of VTT in a facilitation and analytical role brought structure, methodological consistency and the ability to incorporate AI-enabled support into the process. The workshops combine expert judgement with systematic analysis. The participants shape the content and validate the emerging insights, while VTT ensures the robustness of the process and the alignment with the aims of T 5.4. This balance between expert input and methodological rigour is at the heart of the co-creative approach and is necessary for a roadmap that is both strategically relevant and grounded in the operational realities of each company.

2.2 Inputs and evidence base

The work conducted in T 5.4 drew on inputs from the company-specific workshops as well as the wider evidence generated across the project. The material produced through the scenario-building sessions and the strategy design workshops formed the central input for the roadmap. These workshops provided insight into uncertainty factors, plausible future states, and strategic implications that each company considers relevant for the development and uptake of the solutions. The scenario narratives, the opportunities and threats identified across futures, and the measures prioritised by the participants created a structured base for the roadmap.

The task also relied on the analytical outputs of ScenAIrios. The platform expanded the range of uncertainty factors, introduced alternative outcomes and provided consistency analysis for the scenario set. These elements added structure and ensured the internal coherence of the scenarios used in strategy formulation. The combination of expert input and AI-supported analysis produced a broad and coherent set of future environments that informed the strategy design and the resulting measures.

In addition to the task-specific material, the roadmap was informed by findings that emerge across the project. Work on societal drivers, policy frameworks and public perception in work package 1 contributed insight into the regulatory and societal conditions that shape the uptake of monitoring and remediation solutions. Work in work package 2 generated evidence on sensor performance, monitoring capability and data management, which influences the feasibility and scalability of the monitoring solutions. Work package 3 provided results on the performance and operational conditions of the remediation technologies. Work package 4 contributed evidence from the demonstration activities in real settings, which grounds the future pathways in operational and organisational realities. Work package 6 offered understanding of stakeholder environments and potential routes for dissemination and uptake. Together, these inputs ensure that the roadmap reflects the broader context of iMERMAID and is aligned with project-wide aims.

The evidence base brought together in T 5.4 therefore combines insights from company workshops, AI-supported scenario analysis and findings emerging from the project's technical, societal and demonstration work. This creates a foundation for a roadmap that is directly connected to the companies' strategic concerns and consistent with the broader evidence generated in the project.

2.3 Limitations and assumptions

The approach used in task 5.4 follows a scenario-driven roadmapping logic that strengthens the treatment of uncertainty but also carries inherent limitations. These arise from both the methodological foundations described in the literature and the practical choices made in the project.

A central assumption of the approach is that a structured exploration of uncertainty can reveal implications that remain stable across different future environments. Siebelink, Halman and Hofman (2016) argue that robust implications can be identified when multiple scenarios are examined side by side and when participants search for elements that recur across divergent futures. This assumes that the scenario set is sufficiently broad to capture the major drivers that influence strategic choices. It also assumes that participants are able to recognise implications that hold across the scenarios without overlooking factors that may be important but less visible in workshop settings (Lindgren & Bandhold, 2009).

The method presumes that co-creative workshops surface organisational insights that cannot be captured through desk research alone. This is supported in the literature, which notes that internal knowledge of markets, processes and constraints is essential for identifying relevant uncertainties and evaluating what they mean for strategy. This depends on the willingness and ability of participants to contribute openly, and on the facilitator's ability to guide discussion. It also assumes that the participating organisations represent a broad enough range of perspectives to avoid systematic gaps (Lindgren & Bandhold, 2009).

The use of the ScenAIrios -platform introduces assumptions linked to AI-supported analysis. The platform increases the breadth of uncertainty factors, expands the range of possible outcomes and automates consistency analysis. Scenario-based roadmapping studies note that automated support can increase structure and reduce bias, but results remain dependent on the quality of the underlying data, the framing of uncertainty factors and the interpretative work performed by experts during workshops. AI-generated outputs still require human validation, and the method assumes that participants can judge which outputs are meaningful and which require adjustment (Cheng et al., 2016).

Scenario-driven roadmapping also involves practical limitations identified in prior work. Research shows that scenario work demands time and interpretative effort from participants, and that the quality of the roadmap depends on their ability to work with abstraction and long-term thinking (Lindgren & Bandhold, 2009). The process assumes that workshop participants share enough familiarity with their organisation's context to evaluate opportunities, threats and strategic measures. If internal views differ significantly, identifying shared implications may be difficult.

Another limitation arises from the structure of the roadmap itself. A single robust roadmap retains communicative clarity, but this choice means that less emphasis is placed on scenario-specific strategies. Studies of scenario-based roadmapping note that the decision to focus on a single roadmap trades scenario-specific nuance for simplicity and usability. The assumption is that the roadmap's robustness compensates for the reduced detail, and that scenario-specific considerations can be integrated later in implementation (Siebelink, Halman & Hofman, 2016).

Finally, the qualitative synthesis of company-specific material introduces interpretative steps. The process assumes that insights from separate firms can be generalised into shared categories without losing essential context. Scenario-based methods recognise this as a common challenge, since organisations differ in constraints, resources and strategic intent. The roadmap therefore reflects a balanced aggregation rather than a precise representation of each company.

Taken together, these limitations and assumptions do not undermine the validity of the approach. They reflect the conditions under which scenario-driven roadmapping operates and the methodological choices that support a clear and practical outcome.

3.0 Future scenarios

This section presents the original scenarios developed for each participant company separately.

The scenarios developed in T 5.4 follow a common structure to ensure clarity, coherence and analytical depth. Each scenario describes a distinct future operating environment for the companies participating in the project. The structure was designed to support strategic sense-making, highlight causal mechanisms and offer a consistent basis for identifying opportunities, threats and strategic measures across divergent futures.

Each scenario begins with an introduction that presents the core logic of the future environment. This opening section establishes the overarching conditions that shape the world in which the company operates. It highlights the regulatory, market or technological forces that define the future setting and frames the central tension or strategic question that the scenario raises. The introduction provides readers with a clear anchor before moving into finer-grained analysis.

This is followed by an exploration of change drivers and dynamics. This section explains the forces that push the scenario forward and describes how regulation, technology, markets or supply chains evolve. It shows how these shifts interact and how they reshape the company's conditions for action. Its purpose is to clarify the underlying structure of the scenario so that strategic implications become traceable to specific drivers rather than to general background trends.

A day in the future describes the scenario through a concrete, narrative moment set in 2030. This short vignette translates the driver logic into lived experience. It shows how the scenario plays out in practice and demonstrates how decisions, risks and operational realities emerge in the world defined by the scenario. This narrative format helps participants recognise the practical meaning of the scenario rather than viewing it as an abstract projection.

The trajectory section outlines how the scenario unfolds between the present and 2030. It shows how early signals may develop, how system behaviour evolves and how the company's strategic position changes over time. This timeline supports the identification of strategic turning points, pressures and windows of opportunity that shape long-term planning.

The outcomes section describes the strategic consequences for the company if the scenario becomes reality. It sets out the conditions under which the company succeeds or struggles and clarifies the strategic stakes. This section links the scenario directly to the roadmapping work, since the universal opportunities and threats identified later in the process draw from the logic presented here.

Each scenario ends with a short reflection. This section raises a question or highlights a tension that the scenario exposes. Its purpose is to support critical thinking, ensure that the scenarios challenge prevailing assumptions and open space for strategic insight. Reflections help participants consider what the scenario means for innovation, competition or governance without prescribing answers.

This structured format ensures that all scenarios provide the same types of insight, which allows meaningful comparison across the set. It also enables the identification of shared implications, which is essential for scenario-driven roadmapping. By maintaining a consistent structure, the scenarios support a clear link between exploratory futures and the strategic measures that form the roadmap.

Table 1: Summary table of scenarios.

Scenario	Core future logic	Main trigger	Winning logic
The second supply chain	Selective pollutant focus plus trade volatility turns recovery into resilience; targeted add-ons scale under contract-driven proof.	Security of supply + targeted compliance.	Modular add-ons, traceability and certified recovery outputs, rapid deployment near sources.
The compliance squeeze	Near-zero limits and reputational penalties force rapid retrofits; hazardous residual classification raises the cost of getting recovery wrong.	Enforcement acceleration and public scrutiny.	Retrofit speed, low-waste pathways, auditable operation, validated performance under dense streams.
The quiet spotlight	Numeric limits stall, but data-led risk scoring and certification determine scrutiny, funding and market access.	Visibility and risk scoring.	Evidence-as-product, certification strategy, stable performance in well-characterised streams, partnerships to scale.
Islands of pressure	Regulation fragments into “must-act” pockets; pilots multiply but scaling stays local due to thin post-pilot pathways and weak recovered-product standards.	Region-specific enforcement and water stress.	Focus on strict markets, strong pilot-to-operation conversion, QA/standards for outputs.
The retrofit reflex: Compliance under constraint	Capital scarcity and fragmented standards favour low-risk retrofits; innovation survives in niches where fit and OPEX discipline dominate.	Cost-constrained compliance.	Brownfield modularity, low OPEX, low disruption installs, clear residual handling.
Innovation in the slow lane	Enforcement stays light and funding goes to visible resilience; advanced treatment scales mainly when bundled and risk is capped.	Long-run intent, weak short-run urgency.	Endurance strategy, bundling with reuse/resilience, credible pilots, capped liability.
No place to hide	Proactive enforcement and EPR-funded tenders drive urgent demand; continuous monitoring and supplier-heavy guarantees push risk onto suppliers.	Strong enforcement and broadened pollutant scope.	Guarantee-ready service models, continuous evidence, load-flex performance, robust supply chains.
Proof at scale	EU-wide certification and grants accelerate modular deployment in niches; centralised hybrids set high benchmarks at large plants.	Regulatory clarity plus funded pathways.	Certification readiness, fast commissioning, replacement logistics, niche targeting where modularity wins.
Reuse becomes routine infrastructure	Bankable mid-sized reuse projects scale; reclaimed-water pricing strengthens economics; credibility depends on phased commissioning and continuous assurance.	Routine reuse investment cycles.	Modular delivery, evidence bundles, phased commissioning, energy-aware operation.

Evidence-driven and performance-linked markets	Drought and salinity tighten margins; payments and financing depend on delivered quality and uptime; monitoring and service become integral.	Reliability economics under water stress.	Outcome-linked service, 24/7 monitoring, lab–sensor alignment, resilient procurement.
Fragmented markets and uneven regulatory pace	Tariffs and approvals vary widely; momentum clusters in a few jurisdictions; compliance relies more on averaged lab evidence than dashboards.	Local economics and uneven approvals.	Adaptable compliance packs, cost discipline, service capability, strategies for bundlers and local regulators.
Stop–start surges and rapid-response permitting	Heatwaves trigger "procure now" bursts and temporary permits, followed by tighter rules; compliance becomes insurable and data-dependent.	Restriction-driven urgency.	Fast modular commissioning, upgrade-ready designs, regional delivery resilience, insurer-compatible evidence.

3.1 Company-specific scenarios: Weeefiner

3.1.1 Scenario 1: The second supply chain

3.1.1.1 Introduction

By 2030, Europe's quaternary wastewater treatment market no longer moves in a single regulatory wave. Instead, it advances in careful steps, tightening limits, but selectively. Policymakers focus first on a short priority list of micropollutants: PFAS and a small set of metals. For operators, this means a shift in mindset. The first investments are not massive retrofits, but smarter monitoring, followed by targeted add-on treatment exactly where problems appear.

At the same time, another pressure quietly reshapes decisions. Trade restrictions on critical materials expand, licensing slows, and import prices swing unpredictably from month to month. What once looked like an optional innovation, recovering metals and nutrients from wastewater, starts to resemble something else entirely: a backup supply line. For smaller and newer entrants offering modular, near-source solutions, recovery becomes less about sustainability branding and more about resilience.

3.1.1.2 Change drivers and dynamics

The first driver is procurement instability. Broad trade restrictions and slow licensing turn critical materials into a planning risk. Deliveries are delayed, prices jump without warning, and long-term contracts lose their reassuring certainty. In response, EU security-of-supply funding for recovery remains steady through 2030. Subsidized modular units near industrial and municipal sources are no longer experimental—they are treated as insurance against supply disruption.

The second driver reshapes how compliance is proven. Inspections still exist, but they matter less than before. Increasingly, compliance is enforced through contracts. Buyers, banks, and insurers demand verified discharge data as a condition for financing, coverage, or long-term agreements. As PFAS and selected metals rise to the top of the priority list, expectations sharpen: monitor these substances first, prove performance continuously, or risk losing trust.

The third driver is quieter but decisive: market infrastructure finally catches up. EU guidance clarifies by-product status for clean concentrates, reducing legal ambiguity. Regulated quality classes and market-access labels signal that recovered metals and nutrients meet defined standards. Together, these rules smooth the path from pilot projects to routine transactions, turning recovery outputs into tradable goods rather than regulatory headaches.

3.1.1.3 Day in the future

It is a Tuesday morning in 2030, and the procurement team at a mid-sized industrial site is reviewing bids. The question on the screen is not "Do we need quaternary treatment?"—that debate ended years ago. Instead, the discussion turns on specifics: Which targeted add-on gives us verified removal of PFAS and nickel, and how quickly can it be deployed? "We can't afford gaps anymore," one manager says, tapping a clause in their insurance contract. It requires continuous, verified discharge data—not annual reports, not estimates. Their lender has similar language. Waiting for the next inspection is no longer an option.

Outside the meeting room, water prices continue their slow climb. The site has already invested in pre-treatment at source to cut fees and manage risk locally. Central treatment plants now receive lower average loads, but with sharper peaks and unpredictable variability. For operations teams, modular, near-source units feel safer. Capacity can be added where the problem actually appears, not where it used to be.

Recovered outputs, once hard to place, now move more easily. With "approved for use" labels attached, clean concentrates fetch premium prices, especially when imports are delayed and global markets jittery. Recovery does not solve every problem, but it takes the edge off uncertainty.

3.1.1.4 Trajectory

From the mid-2020s to 2030, micropollutant regulation tightens in stages. The first response is almost universal: upgraded monitoring. Soon after, targeted add-ons appear, initially as pilots, then as contractual necessities. As compliance-by-contract spreads, the burden of proof shifts decisively. Verified discharge data becomes a condition of doing business, not a regulatory afterthought.

Meanwhile, trade volatility becomes the background noise of procurement. Policymakers respond with sustained funding for recovery, framing it as a security measure rather than an environmental bonus. Modular recovery units proliferate near sources, especially as more pre-treatment shifts loads away from central plants and makes system-wide flows harder to predict.

Market formation follows. Clear by-product rules reduce friction in transport and resale. Quality classes and labels improve buyer confidence. Incumbents respond by forming structured partnerships, integrating modular entrants into service and sales channels. Scaling does not come from dramatic breakthroughs, but from repetition: dependable performance, documented results, and accepted standards.

3.1.1.5 Outcomes

By 2030, demand concentrates on targeted quaternary add-ons that can prove removal of priority micropollutants while enabling recovery where economics allow. Performance documentation becomes a primary competitive differentiator, often more decisive than headline removal rates.

Economically, stable recovery subsidies and volatile imports tilt decisions toward modular deployments near sources. The value is not just lower treatment costs, but reduced exposure to supply shocks and price swings. Recovery is justified as risk management as much as resource efficiency.

Technologically, progress remains incremental. Selective media last longer thanks to better cleaning and backwash routines. These modest gains matter when multiplied across hundreds of installations, especially within partnership models where utilities and OEMs share revenue and responsibility. The result is a market that scales quietly, through integration rather than disruption.

3.1.1.6 Reflection

If wastewater becomes Europe's "second supply chain," who should be responsible for maintaining it: utilities, industrial operators, or the public sector, and what risks emerge if that responsibility remains unclear?

3.1.2 Scenario 2: The compliance squeeze

3.1.2.1 Introduction

By 2030, Europe's quaternary wastewater treatment market is shaped by a tightening grip. On one side, regulators move fast and visibly. Near-zero PFAS limits arrive with little grace period, and after several high-profile pollution scandals, penalties escalate sharply. For operators, especially those under public or investor scrutiny, quaternary add-ons are no longer a strategic option. They are a deadline.

At the same time, recovery is no longer framed as a hedge against blocked imports. Trade flows ease, but expectations harden. Recycling rules grow stricter, sourcing claims must be proven, and traceability becomes non-negotiable. Recovery systems are judged not just on what they extract, but on how well they document every step—while navigating a paradox: tighter controls now classify more concentrates and spent media as hazardous waste, raising the cost of getting recovery wrong.

3.1.2.2 Change drivers and dynamics

Enforcement sets the tempo. Between 2028 and 2030, emergency PFAS limits force rapid retrofits. Following widely publicized pollution cases, fines increase and offender lists go public. Compliance delays now carry reputational damage, not just regulatory risk. Demand accelerates through procurement urgency rather than long-term planning.

Economics narrow the solution space. Public funding does not disappear, but it changes character. Grants shift away from pilots toward scaled, proven retrofits. At the same time, hazardous waste classification expands. Disposal costs rise, cross-border shipment becomes harder, and residual management moves from a footnote to a core design concern. Meanwhile, industrial sites maintain output while tightening water loops: volumes fall, concentrations rise, and selective treatment becomes unavoidable.

Technology stabilizes, competition intensifies. Electrochemical interfaces become more reliable and controllable. Energy use does not drop dramatically, but consistency improves, critical in a world of traceability audits and grant reporting. Incumbents respond aggressively, rolling out in-house modules bundled with long service contracts and sharp pricing. For smaller players, modularity alone is no longer enough.

3.1.2.3 Day in the future

In 2030, a plant manager in a chemical corridor scans the morning news before heading into a compliance call. Another facility, two regions over, has just been fined millions. The story names executives. Investors are asking questions.

"Our permit review is in six months," she tells her team. "If PFAS shows up above the limit, we don't get a warning—we get a headline."

Her plant recycles more water than ever. Internal loops are tight, volumes are down—but the remaining wastewater is dense, difficult, and unforgiving. A quaternary unit has to be installed fast. Procurement focuses less on peak removal rates and more on proof: continuous monitoring, traceable outputs, predictable residuals.

Then comes the trade-off. One option recovers more, but produces a concentrate that now counts as hazardous waste. Disposal routes are limited and expensive. Another removes slightly less, but generates cleaner residuals that are easier to handle. The "best" technology is no longer obvious. It is a risk calculation.

3.1.2.4 Trajectory

Momentum builds through the late 2020s as PFAS limits tighten. Early movers retrofit first, but acceleration comes when enforcement hardens. After publicized cases, several member states raise fines and publish offender lists. Waiting becomes visibly dangerous.

Financing reshapes the market next. Support programs pivot from experimentation to deployment. To qualify, suppliers must demonstrate repeatable performance, standardized monitoring, and clear reporting. Traceability standards align with this shift, rewarding solutions that can continuously substantiate compliance and recycled content claims under stricter recycling rules.

Operational pressures complicate scale-up. Water reuse intensifies concentrations, increasing both treatment difficulty and residual risk. Hazardous classification spreads, forcing operators to confront disposal logistics early in design. Technological reliability improves, but competition tightens as incumbents lock in customers with bundled offerings.

3.1.2.5 Outcomes

By 2030, quaternary treatment is a compliance reflex, not a strategic debate. PFAS limits and penalty spikes dominate decision-making, driving a fast, retrofit-heavy market where speed, auditability, and risk reduction matter as much as performance.

Recovery plays a dual role. It supports recycling and sourcing expectations in a more regulated trade environment, but its value is constrained by hazardous residual rules. Operators prioritize solutions that minimize problematic concentrates and offer clear handling pathways, even if recovery yields are modest.

Market power concentrates. Incumbents leverage pricing and service contracts, while SMEs must differentiate through verifiable performance, traceability readiness, and fit with grant-driven retrofit programs. With production tied to water budgets and reuse tightening loops, demand persists for selective, modular quaternary systems that can operate reliably under high-strength conditions.

3.1.2.6 Reflection

When compliance pressure leaves little room to experiment, does Europe risk locking in "good enough" solutions, or does the squeeze finally force wastewater systems to become transparent, accountable, and resilient by design?

3.1.3 Scenario 3: The Quiet Spotlight

3.1.3.1 Introduction

By 2030, Europe's quaternary wastewater treatment market is shaped less by new limits than by a new kind of visibility. Regulators hesitate to impose additional numeric micropollutant thresholds, constrained by cost and political fatigue. Instead, they double down on monitoring. Data flows more freely, inspections become risk-based, and shared digital reporting systems quietly decide where attention falls.

Operators describe the shift as unsettling. There is no single deadline to prepare for, no headline limit to meet. Yet the pressure is constant. Risk scores determine who gets inspected, who qualifies for funding, and who draws scrutiny from insurers, buyers, and local authorities. In this environment, "being able to prove control" matters almost as much as control itself.

Against this backdrop, a small, recovery-focused player looks for traction. Industrial reshoring concentrates wastewater in European hubs. Water stress returns each summer. And although critical materials trade restrictions are targeted rather than sweeping, predictable imports still carry modest price swings—enough to keep recovery projects economically interesting, especially when compliance and resilience overlap.

3.1.3.2 Change drivers and dynamics

Regulation shifts from thresholds to exposure. Micropollutant limits stall, but enforcement intensifies. Data-led, risk-based systems sort sites continuously, flagging those that appear unstable or poorly documented. Operators invest selectively, not to meet new numbers, but to reduce their risk profile and avoid the spotlight.

Funding and market access become conditional. Subsidies remain available, but competition increases. Grants favour projects with fast payback, measurable pollutant reductions, and recoverable outputs. Smaller operators struggle to qualify. At the same time, recovered metals and nutrients gain product status through simple quality checks, but only from certified plants. Certification, not technology alone, becomes the gateway to commercialization.

Demand clusters and technology pathways narrow. Reshoring concentrates industrial wastewater in EU hubs. Recurring summer droughts push local reuse, tightening water loops. Selective media improve incrementally for key metals like copper, nickel, and zinc, but mostly in stable streams. Mixed wastewater still requires pretreatment, steering activity toward well-characterized industrial flows.

3.1.3.3 Day in the future

In a drought-prone industrial hub in 2030, a site engineer checks the regional risk dashboard before approving a maintenance shutdown. The plant's score is "amber", not alarming, but close enough to trigger questions.

"We're not breaking any limits," she says, half to herself. "But the data doesn't like uncertainty."

The site has invested in monitoring and reporting upgrades. Not because regulators demanded them outright, but because funding applications, insurance renewals, and customer audits now ask for the same thing: proof. Meanwhile, reshoring has increased metal-bearing wastewater on site. Prices for recovered copper fluctuate just enough to make recovery modules attractive, if the outputs can be certified.

Certification is the sticking point. One recovery option produces a clean, sellable stream; other risks falling into a grey zone, complicating QA and contracts. The decision is less about maximum yield and more about predictability. In a system governed by scores, surprises are dangerous.

3.1.3.4 Trajectory

From the mid-2020s to 2030, Europe leans into measurement. Numeric limits stall, but digital oversight accelerates. A "soft cap" discourages blanket upgrades, while a "hard spotlight" focuses scrutiny on high-risk sites identified by shared data systems.

Industrial hubs created by reshoring become testing grounds. Here, quaternary add-ons that can demonstrate measurable reductions and recovery potential attract funding, especially when projects are designed to meet grant criteria and achieve certification. Recovery is framed not as speculation, but as disciplined risk management.

Technology and market structure evolve unevenly. Incremental selectivity gains expand recovery in stable streams. Mixed wastewater remains challenging. Incumbents wait, then act: after multi-site proof, large OEMs acquire a handful of proven solutions and bundle them into retrofit packages. Validation becomes harder, but once achieved, scale comes quickly.

3.1.3.5 Outcomes

By 2030, market pull concentrates where visibility and vulnerability overlap: reshored industrial hubs, drought-stressed regions, and sites flagged by risk-based enforcement. Even without tighter numeric limits, pressure to demonstrate control is real, making monitored, auditable quaternary steps defensible and often necessary.

Recovery is more feasible but still uneven. Product status for recovered metals and nutrients rewards certified plants, turning certification into a strategic bottleneck. Sector-led specifications spread slowly, creating a patchwork of acceptance and pricing.

Competition favours those who can prove performance, then let others scale it. Selective acquisitions reward validated technologies but compress differentiation once bundled. For SMEs, opportunity lies in stable, well-characterized metal streams, where incremental selectivity gains align with a funding environment that prizes certainty over ambition.

3.1.3.6 Reflection

When regulation relies more on visibility than limits, who benefits most—the innovators, the incumbents, or the best-documented—and how might this quiet spotlight reshape trust, fairness, and innovation in water systems?

3.1.4 Scenario 4: Islands of Pressure

3.1.4.1 Introduction

By 2030, Europe's quaternary wastewater treatment market does not move as one. Instead, it fragments into islands of urgency and zones of delay. Micropollutant limits tighten, but only for high-risk sectors and sensitive waters. Enforcement follows suit, uneven and region-specific. Some sites face immediate compliance pressure; others can postpone upgrades for years. In this landscape, broad claims matter less than proof in the right place, at the right time.

Meanwhile, trade in certain critical raw materials remains heavily restricted. On paper, this elevates wastewater as a potential secondary source. In practice, the pull is weaker. Without common recovered-product standards and with buyer liability concerns lingering after contamination incidents, recovery remains strategically interesting but commercially cautious. For SMEs and startups, opportunity appears in pockets: where regulation bites, seasonal water stress tightens operations, and modular, site-specific solutions can be justified despite fragmented rules and hesitant incumbents.

3.1.4.2 Change drivers and dynamics

Targeted regulation, uneven enforcement. Micropollutant limits focus on specific sectors and waters. Inspections are frequent in some regions and rare in others. The result is a two-speed Europe: fast-moving Northern and Western jurisdictions where non-compliance is costly, and slower regions where delay remains cheaper than action.

Pilots without pathways. Funding flows easily into demonstrations. Subsidies support pilots, data collection, and proof-of-concept projects. But mechanisms beyond pilots are thin. Buyers ask a hard question: who carries the risk once the grant ends? Incumbents respond cautiously—allowing trials, but withholding sales channels and long-term commitments until mandates are unmistakable.

Recovery pulled, then restrained. Restricted trade strengthens the strategic logic for recovery, but missing standards and heightened liability cool enthusiasm. Buyers demand discounts, extra testing, and tight off-take controls. Country-by-country waste classification rules add friction, forcing redesign and re-permitting with every border crossed. Scaling becomes a sequence of local negotiations, not a single European rollout.

3.1.4.3 Day in the future

In 2030, two factories operate under the same EU flag, and live very different days.

In one region, inspections are frequent. Sector-specific limits apply. The plant manager reviews audit results weekly. "If we wait," she says, "we're betting against a fine we can't absorb". Quaternary upgrades are framed as insurance: auditable performance, predictable operation, reduced risk.

A few hundred kilometres away, another factory faces rare inspections. The same upgrade proposal lands on the CFO's desk, and stalls. "Why invest now," he asks, "when non-compliance is still cheaper?"

Summer brings a shared shock. Heatwaves trigger short abstraction bans. Even where annual water availability looks stable, peak stress forces plants to scramble. Some turn to modular add-ons that can be deployed quickly, especially as legacy high-water-use operations quietly shut down in sectors like textiles and plating.

Recovery enters the conversation, but cautiously. Restricted access to raw materials keeps interest alive, yet buyers hesitate. Standards are unharmonized. Liability feels personal after recent incidents. Technologies that perform consistently, despite uneven fouling resistance, despite site-specific challenges, are the ones that earn trust in the must-comply islands.

3.1.4.4 Trajectory

From the mid-2020s to 2030, momentum clusters where sector-based limits and active enforcement coincide. Procurement decisions are justified as compliance hedges, not revenue plays. Pilots multiply, supported by subsidies, but many stall without post-pilot funding or strong enforcement signals.

Incumbents reinforce selectivity. They permit trials, observe results, and wait. Only when regulation becomes unavoidable do they consider channelling solutions at scale. Operational reality sharpens this caution. Selective media reduce biofouling, but scaling and oils still cause downtime in some sites. Performance remains site-dependent, increasing the value of careful selection and local tailoring, especially under fragmented waste classification regimes that slow replication.

3.1.4.5 Outcomes

By 2030, Europe's quaternary market is ambitious but uneven. High-pressure regions and sectors advance decisively; lagging areas slow overall adoption. Demand concentrates in pockets where enforcement is real and water stress visible, leaving a long tail of unconvinced sites.

Recovery remains strategically relevant but commercially constrained. Without common standards and under persistent liability concerns, recovered products struggle to clear procurement hurdles except under tightly controlled conditions. Fragmented waste rules further complicate scaling, stretching timelines and compliance costs.

For SMEs and startups, success is possible, but filtered. Win where seasonal water limits force action, where sector-based rules create urgency, and where pilots can convert into operating spend without perfect funding continuity. Credibility is earned site by site. By 2030, success looks less like a single European breakthrough and more like a mosaic of locally optimized footholds.

3.1.4.6 Reflection

When regulation advances unevenly, does fragmentation slow innovation, or does it create the proving grounds where resilient solutions are forged one site at a time?

3.2 Company-specific scenarios: EdenTech

3.2.1 Scenario 1: The retrofit reflex: Compliance under constraint

3.2.1.1 Introduction

By 2030, quaternary micropollutant treatment in Europe has become less a story of technological breakthroughs and more a story of careful compliance under constraint. Regulators, utilities, and technology providers operate in a landscape shaped by pragmatic enforcement, tight capital, and cautious decision-making.

Regulatory pressure exists—but it arrives softly. France and Germany maintain regular inspections and moderate fines, while the UK applies regionally uneven enforcement with phased deadlines that reward credible upgrade plans over rapid transformation. At the EU level, the priority pollutant list expands incrementally, adding select high-risk PFAS and pharmaceuticals, yet thresholds remain broadly stable, and mixture rules are introduced slowly.

The result is a system that signals urgency without forcing radical change. Everyone knows more will eventually be required. Few feel compelled, or able, to move fast.

3.2.1.2 Change drivers and dynamics

Three interlocking uncertainties define this market.

First: capital scarcity. Utilities face severe CAPEX constraints. Budgets are frozen or tightly rationed, and Extended Producer Responsibility (EPR) reimbursements arrive late—often 18 to 36 months after expenditure. This delay quietly reshapes behaviour. Utilities gravitate toward the cheapest visible compliance step, prioritizing solutions that fit existing infrastructure and defer financial risk.

Second: technology trust and procurement friction. Low-cost adsorption media, now cheaper to operate and easy to slot into existing filters, becomes the default retrofit choice. In contrast, newer systems face a fragmented standards landscape: national protocols differ, certifications are not fully portable, and short site-specific pilots remain mandatory. Innovation is not rejected, but it must repeatedly prove itself, site by site.

Third: operational volatility. Energy markets remain unstable. Grid strain produces frequent peak-price events and time-of-use constraints. For energy-linked options such as photocatalysis, fluctuating LED module prices complicate lifecycle costing. What looks efficient on paper feels risky in procurement committees.

Together, these forces create a powerful reflex: minimize upfront cost, minimize perceived risk, and move only as far as compliance requires.

3.2.1.3 Day in the future

Elena, an engineering manager at a mid-sized European utility, stares at two spreadsheets on a late Tuesday afternoon.

One shows an adsorption retrofit: low CAPEX, familiar technology, minimal disruption. The other models a microfluidic photocatalytic unit, technically elegant, promising better removal for certain pharmaceuticals, but dependent on volatile component prices and a pilot that would delay sign-off by months.

Her budget is frozen. Her compliance deadline is phased but real. EPR reimbursement, if approved, will come long after installation.

“We can’t afford to be brave”, she says quietly to her team.

They choose the adsorption retrofit and schedule a short pilot, just enough to satisfy regulators. Across the region, similar meetings end the same way.

Meanwhile, a small photocatalytic technology vendor packs up a demo unit after a successful trial in a drought-stressed reuse site. The data is strong. The feedback is polite. The contract, once again, is deferred.

3.2.1.4 Trajectory

The market follows a steady but uneven path. Adsorption becomes the compliance default, reinforced by CAPEX freezes and reimbursement delays. Nationally fragmented standards keep pilots alive and slow universal adoption of newer systems.

Over time, insurers and utilities formalize third-party warranties and risk-pooling mechanisms. This reduces direct supplier liability but introduces new audit and premium costs, favouring firms with the balance sheets to absorb them.

Microfluidic photocatalytic systems do not disappear. Instead, they find narrow but durable footholds:

- Water-reuse sites in drought-affected southern France
- Select regions of Germany with reuse mandates
- Applications where adsorption performs poorly for specific pollutant mixtures

Innovation survives, but mostly at the margins.

3.2.1.5 Outcomes

By 2030, the market is constrained but navigable for emerging photocatalytic SMEs. Broad adoption remains elusive, not because the technology fails, but because the system rewards lowest-cost compliance over performance optimization.

Winning strategies look different than expected a decade earlier. Success comes from:

- Retrofit-friendly designs.
- Clear pilot evidence in real operating conditions.
- Insurer-underwritten performance guarantees.
- Transparent handling of residuals (sludge or brine) with manageable OPEX.

This is not a market of rapid roll-outs. It is a market of selective trust, incremental proof, and patient positioning.

3.2.1.6 Reflection

In a system optimized for cautious compliance rather than long-term performance, where should innovation push hardest—on technology, on financing models, or on changing how risk itself is shared?

3.2.2 Scenario 2: Innovation in the slow lane

3.2.2.1 Introduction

Europe in 2030 is not hostile to quaternary micropollutant treatment, but it is not welcoming either.

EdenTech, a small but technically ambitious SME, is pursuing microfluidic–photocatalytic solutions across France, the UK, and Germany. On paper, the need is clear. Regulators acknowledge the risks of persistent micropollutants, and policy frameworks exist. Yet the central question remains unresolved:

How does a novel technology scale in a system where urgency is always postponed? Two structural forces dominate. Enforcement remains light and penalties low, blunting immediate pressure on utilities. At the same time, municipal capital is being redirected toward visible climate resilience—flood defences, drought infrastructure, reuse capacity, leaving micropollutants as a secondary concern unless they can be bundled in. Together, these dynamics produce a market that prefers safe retrofits, limited liability, and gradual experimentation over decisive transformation.

3.2.2.2 Change drivers and dynamics

The regulatory environment sends mixed signals. EU frameworks exist, but priorities diverge regionally. Pollutant lists vary, forcing suppliers to adapt systems site by site. Certification remains fragmented, and no trusted EU-wide standard emerges. As a result, pilots become the proof mechanism, often running across multiple seasons before procurement decisions are made.

Financing reinforces caution. Extended Producer Responsibility schemes shift residuals risk away from utilities, which lowers one barrier to experimentation. However, reimbursements arrive slowly and cover only narrow cost categories. In practice, EPR funding supports pilots and phased rollouts, not full-scale deployment.

Procurement rules further dampen risk appetite. Utilities increasingly favour contracts that cap supplier exposure and share risk, protecting themselves but stretching SME balance sheets. Meanwhile, supply chains are reshaped by critical-materials regulation. Components are compliant and traceable, but more expensive, adding 15–25% to upfront costs.

Energy dynamics offer partial relief. Renewable penetration lowers electricity prices, improving operating costs for photocatalytic systems. Yet carbon costs shift into materials, keeping capital expenditure high. And as water stress intensifies, political attention tilts toward reuse infrastructure. Quaternary treatment is funded primarily when it rides along with reuse or resilience projects, not as a standalone priority.

3.2.2.3 Day in the future

In 2030, a regional utility in Germany reviews its investment plan.

The engineering team agrees that quaternary treatment will eventually be required. The regulator agrees too—just not yet. Inspections are infrequent, and fines remain symbolic. The capital budget, however, is already committed: flood defences upstream, reuse capacity downstream.

A pilot proposal sits on the table. EdenTech's system looks promising, and EPR rules mean the utility won't carry long-term residuals liability. Still, the procurement team hesitates. Certification will require a multi-season trial. Audits will increase. And no one wants to explain a costly technology choice when enforcement is still optional.

"Let's pilot", the director finally says. "But we're not locking anything in".

For EdenTech, this is progress, and another delay.

3.2.2.4 Trajectory

Between now and 2030, EdenTech's most viable pathway is incremental. Rather than chasing universal adoption, the company focuses on clusters of opportunity: sites where quaternary treatment can be bundled into reuse or resilience projects, and regions where pollutant profiles justify advanced solutions.

Proposals are tailored to regional standards, demonstrating adaptability across France, the UK, and Germany. Timelines and pricing assume long pilots and capped liability. Supply chains are documented meticulously to satisfy critical-materials rules, even as costs rise.

EPR reimbursements, predictable but slow, are used to finance staged deployments. Performance-linked contracts align with procurement norms, trading speed for credibility. The company accepts that weak enforcement will persist, and builds a business model designed to survive it.

3.2.2.5 Outcomes

By 2030, adoption remains selective rather than universal. EdenTech systems operate where reuse priorities, EPR funding, and bundled investments align. Sales cycles are long. Solutions are bespoke. Margins are tight.

Operating costs benefit from cheaper electricity, but higher material costs squeeze CAPEX-heavy deployments. Competitors make incremental gains, yet no single technology dominates. In this environment, success is not defined by rapid scale, but by endurance.

EdenTech survives, and grows modestly, by excelling at pilots, navigating regional complexity, and aligning innovation with a system that prefers caution over commitment.

3.2.2.6 Reflection

If regulation signals long-term intent but short-term patience, who should bear the cost of moving first, and what would finally make delay more expensive than action?

3.2.3 Scenario 3: No Place to Hide

3.2.3.1 Introduction

By 2030, the world EdenTech operates in has changed decisively. Enforcement is no longer sporadic or symbolic. It is proactive, producer-funded, and backed by courts willing to impose real penalties. Expanded inspection teams, financed through EPR, conduct regular checks, and non-compliance now carries reputational and financial consequences that utilities and suppliers can no longer ignore.

At the same time, regulatory attention has shifted abruptly. The focus is no longer limited to familiar pharmaceuticals or PFAS, but extends to industrial additives and their transformation products, harder to detect, harder to remove, and harder to explain to the public.

The effect is immediate: demand for quaternary treatment becomes urgent and regulator-driven. But the bar for success rises just as fast. It is no longer enough to perform well in periodic tests. Performance must be proven, continuously.

3.2.3.2 Change drivers and dynamics

The mechanics of the market are transformed. EPR funding no longer flows quietly through utilities; it is routed directly through procurement tenders. Vendors are paid, but access is gated. Winning business now means surviving competitive tenders where only a handful of suppliers are deemed credible enough to carry the risk.

Certification undergoes a similar shift. Intermittent lab reports are replaced by digital-only certification based on live monitoring. Sensors stream performance data continuously, feeding compliance dashboards that regulators can inspect at any time. Treatment performance becomes less like an audit, and more like a live broadcast.

Procurement expectations harden accordingly. Contracts impose supplier-heavy guarantees with strict penalties. If removal efficiency drops, if uptime falters, if residuals are mishandled, the liability sits squarely with the vendor. Insurance, partnerships, and balance-sheet strength become as important as chemistry.

3.2.3.3 Day in the future

In 2030, an EPR-backed tender opens for a regional water authority.

The requirements are unforgiving. Continuous compliance dashboards are mandatory. Limits on newly prioritized industrial additives are tight. Energy consumption and lifecycle residual handling must be disclosed upfront. Concentrated residual streams are now classified as hazardous, requiring licensed transport and documented disposal.

Inside EdenTech, the mood is tense.

Winning this tender would unlock scale across multiple regions. Losing it, or worse, winning and failing, could threaten the company's survival. Engineers double-check energy-per-kilogram calculations. Legal teams scrutinize penalty clauses. Management debates how much risk the company can realistically carry.

There is no room for optimism without proof. The system will be watching, every hour, every day.

3.2.3.4 Trajectory

EdenTech's path forward is shaped by two opposing forces: opportunity and exposure.

On one side, proactive enforcement and EPR-funded tenders accelerate deployment. CAPEX barriers fall as subsidies and direct vendor payments smooth roll-outs. On the other, upstream risks intensify. High-grade catalyst prices spike unpredictably, rising by 30–60%. Electricity prices remain elevated, with sharper carbon penalties amplifying cost volatility.

To survive, EdenTech must relentlessly optimize. Energy efficiency becomes a strategic weapon. Service intervals are extended. Digital monitoring is not just a compliance tool, but a sales argument, proof that guarantees can be met despite material and energy pressure.

Success depends on convincing buyers that supplier-heavy guarantees are survivable, even in a high-cost environment.

3.2.3.5 Outcomes

If EdenTech succeeds, meeting continuous monitoring standards, controlling energy use, and managing hazardous residuals cost-effectively, it can scale rapidly across EU-regulated markets. In this world, enforcement and tenderization reward technologies that are verifiably reliable, not merely promising.

If it fails, the consequences are swift. Inability to meet live performance thresholds, spiralling residuals costs, or weaker energy economics compared to alternatives such as electrochemical oxidation systems could confine EdenTech to niche pilots or force deep partnerships to share guarantee risk.

This is not a forgiving market. But it is a decisive one.

3.2.3.6 Reflection

When compliance is continuous and risk is pushed onto innovators, does this accelerate better solutions, or quietly favour only the biggest players who can afford to be watched all the time?

3.2.4 Scenario 4: Proof at Scale

3.2.4.1 Introduction

By 2030, EdenTech is operating in a European regulatory environment that leaves little room for ambiguity. Enforcement is strict, data-led, and backed by escalating fines. Continuous monitoring and public reporting are now standard, and mixture-based compliance, limits based on combined toxicity rather than single substances, has become central.

For utilities, this changes everything. What matters most is no longer whether a technology can remove a specific pollutant in isolation, but whether it can demonstrably and reliably reduce overall risk, across complex mixtures, under real operating conditions.

For EdenTech, an SME with compact, intensified microfluidic–photocatalytic systems, the opportunity is real, but so is the challenge. The question is no longer whether regulation will force action. It is whether a smaller player can turn regulatory clarity into repeatable, financed deployments before larger, more established technologies capture the ground.

3.2.4.2 Change drivers and dynamics

Several forces align to accelerate adoption, while sharpening competition.

An EU-wide certification scheme with fast-track permitting replaces fragmented national approvals. Once certified, technologies move quickly from approval to installation. At the same time, policymakers redirect EPR funds into upfront grants for certified systems, lowering capital barriers and turning regulatory compliance into a funded pathway rather than an unfunded mandate.

Procurement practices evolve in parallel. Utilities retain operational risk but structure contracts around verified performance bonuses. Suppliers are not asked to guarantee uptime, but they are expected to deliver certified results and enable rapid module replacement when performance dips.

Technical boundary conditions tighten. On-site residual destruction becomes mandatory for permits, forcing all solutions to confront waste handling directly. Meanwhile, membrane AOP hybrids gain favour at large plants, offering a robust, centralized option that sets a high-competitive benchmark.

Costs remain manageable but visible. Regional sourcing limits volatility, and carbon charges stabilize. Short, intense droughts introduce seasonal discharge limits, creating episodic demand for fast, modular upgrades funded through ring-fenced capital.

3.2.4.3 Day in the future

It is August 2030 at a coastal treatment plant in southern France.

River flows are low. Sensors tighten discharge limits automatically, and the compliance dashboard, publicly accessible, turns amber. The plant manager knows fines can escalate within days.

A certified quaternary module, installed at the start of summer, comes online. It was approved quickly under EU fast-track rules and partially funded through an EPR-backed grant. The contract is simple: a fixed service fee plus bonuses for verified mixture-toxicity reduction.

The operator remains responsible for uptime. But when performance begins to drift, the module is swapped within hours. Residuals are destroyed on site, logged automatically to satisfy permit conditions.

At a larger inland facility, the story is different. There, a membrane AOP hybrid handles baseline compliance. EdenTech’s modules appear only at the margins, used where space is tight, or where rapid seasonal response is critical.

3.2.4.4 Trajectory

Between now and 2030, the system settles into a new rhythm.

High-sensitivity monitoring becomes ubiquitous, enabling strict enforcement without constant inspections. Mixture toxicity caps are standardized, and certification becomes the main gateway to market access. EPR-backed grants turn certification into a commercial accelerant, shortening the path from pilot to paid deployment.

Utilities, operating under tight but manageable budgets, favour phased, modular retrofits. Procurement norms stabilize around utility-retained risk paired with verified performance incentives. Suppliers that can prove results, and swap modules quickly, gain trust.

At the same time, membrane AOP hybrids consolidate their position at scale, shaping a segmented market rather than a winner-takes-all outcome.

3.2.4.5 Outcomes

By 2030, a mixed market has emerged.

Certified, grant-backed microfluidic–photocatalytic modules achieve repeatable deployments in priority niches: seasonal permits, space-constrained sites, and situations demanding rapid response. Their value lies in low footprint, fast deployment, and the ability to show results, continuously.

Broader market dominance remains elusive. Many large plants default to membrane AOP hybrids, and strict residual-destruction rules favour solutions with proven, integrated waste handling. Operating costs are predictable, if slightly higher, and supply chains remain stable.

For EdenTech, success depends not on outcompeting every alternative, but on turning proof into pattern: leveraging EU certification, upfront EPR grants, and verified performance to build a modular, procurement-ready business that scales were speed and clarity matter most.

3.2.4.6 Reflection

When regulation rewards what can be proven, continuously and publicly, does this level the playing field for innovative SMEs, or quietly raise the bar in ways only established players can clear?

3.3 Company-specific scenarios: IRIS

3.3.1 Scenario 1: Reuse becomes routine infrastructure

3.3.1.1 Introduction

By 2030, the advanced wastewater market in Europe and Northern Africa expands through repeatable, financeable reuse projects rather than headline-grabbing breakthroughs. For an engineering SME like IRIS srl, growth is real but conditional: deliver on compressed schedules, prove performance continuously, and plan commissioning in phases when key equipment cannot arrive all at once.

The business case strengthens as many cities as possible price reclaimed water closes to drinking water. Municipalities and industrial sites treat reuse as standard infrastructure, with decisions driven by payback, guaranteed volumes, and the cost of interruptions.

Approvals for lower-risk uses such as irrigation and industrial cooling move faster thanks to clear templates, shifting the bottleneck from paperwork to proof. Compliance is increasingly demonstrated through certified online data, with infrequent lab checks and tighter expectations around data handling. In parallel, buyers pay upfront through multi-year framework tenders, but attach clear performance penalties and annual verification. Demand rises; confidence goes to suppliers that can make results visible, auditable, and stable under real operating conditions.

3.3.1.2 Change drivers and dynamics

Faster institutions and sharper accountability reshape the pipeline. Permits arrive in months for low-risk reuse, and tenders increasingly assume continuous digital evidence rather than sporadic sampling. Framework purchasing accelerates decisions, while penalties and annual verification harden the consequences of underperformance.

Climate pressure is steady, not catastrophic: hotter summers and recurring local drought keep reuse funded year after year. Gradually higher salinity and tighter reliability expectations push designs toward robust polishing and control.

Economics reinforce the shift. A high, predictable carbon price across EU industry turns energy performance into a primary evaluation filter. Energy costs still swing, but more sites add solar and batteries and prefer modular plants that can shift operations to cheaper hours.

Financing becomes less of a stop sign. Blended structures de-risk lending, unlocking mid-sized upgrades and pilots, segments where SMEs can win with modular delivery and execution. Cooperation between Europe and Northern Africa moves faster with shared rules and steady co-funded programs, while trade corridors stabilize import processes and shipping times.

Constraints stay practical. Surging demand lengthens lead times for membranes, UV, sensors, and controls, making phased installs and realistic schedules part of credible bids. Meanwhile, stricter industrial pre-treatment enforcement reduces shock loads, so value shifts from crisis response to consistent performance and continuous assurance. Competition remains selective: multinationals focus on the biggest tenders and key industrial clients, leaving room for SMEs in modular niches, though with moderate price pressure and higher service expectations.

3.3.1.3 Day in the future

In a drought-prone region, a utility board reviews its next upgrade. The question is no longer whether reuse is acceptable, but how to scale quickly without taking on compliance risk. Operators point to what has changed at the inlet: fewer sudden shocks than before, but a slow rise in salinity that makes “almost compliant” a dangerous place to live.

The finance case is straightforward. With reclaimed water priced close to drinking water, the project reads like a standard capacity investment. The debate tightens around execution: approvals, proof, and delivery timing.

Permitting for irrigation and industrial cooling now runs on templates and arrives in months. The tender is equally direct about evidence: certified online monitoring is the primary proof, lab checks are rare, and data must be audit-ready. Procurement comes through a multi-year framework, upfront payment paired with performance penalties and an annual verification report. Bidders understand they are selling operational certainty, not just equipment.

Energy is treated as a design constraint. The utility is expanding solar and batteries to buffer price swings and wants a modular plant that can run harder when power is cheaper. A predictable carbon price makes low-energy operation non-negotiable.

Across the Mediterranean, a similar tender looks familiar. Shared rules and co-funding shorten interpretation and contracting, and logistics are steadier, but hardware queues remain long. The bids that survive do not deny the waits; they phase commissioning while ensuring that monitoring, reporting, and compliance confidence are in place from day one.

3.3.1.4 Outcomes

By 2030, the market scales in a bankable pattern centered on mid-sized upgrades and decentralized deployments. Faster approvals for low-risk reuse expand volumes, while compliance built around certified online monitoring, with only occasional lab checks, makes continuous data a core deliverable. With reclaimed water priced near drinking water, projects clear internal hurdles more easily and move into routine infrastructure budgets.

Climate impacts keep pressure on without triggering panic. Local drought hotspots sustain demand, and rising salinity plus tighter reliability expectations raise the bar for stable operation. At the same time, a high and predictable carbon price across EU industry anchors tender decisions around energy efficiency over the full lifecycle.

Commercial terms become stricter as speed increases. Framework tenders pay upfront but enforce performance penalties and annual verification, rewarding suppliers that can commit to measurable outcomes and maintain a defensible evidence trail. Blended finance widens the project funnel and supports pilots that help SMEs build references.

Delivery discipline becomes a differentiator. Even with more predictable cross-border trade, demand surges create long lead times for key hardware, normalizing phased installations. Buyers increasingly prefer modular plants that can flex around energy price swings, often supported by on-site renewables and storage. Stricter industrial pre-treatment reduces shock loads, shifting operations toward steady polishing and continuous assurance. With shared rules and co-funded programs linking Europe and Northern Africa, an SME can compete effectively by executing quickly, proving performance continuously, and repeating deployments with minimal reinvention.

3.3.2 Scenario 2: Evidence driven and performance linked markets

3.3.2.1 Introduction

By 2030, across Europe and Northern Africa, advanced wastewater treatment has moved beyond scattered pilots into a repeatable market where credibility is earned in operation. Severe, multi-year droughts have turned reuse from a sustainability option into a continuity measure. Utilities and industrial sites treat reclaimed water as a strategic supply, while tougher conditions, especially rising salinity, push plants closer to their technical limits.

In this landscape, pricing, permits, and finance all behave like linked scoreboards. Reclaimed water is paid according to the quality actually delivered and the hours the system stays online. Permits come faster when operators open a digital file and keep regulators continuously informed with live quality data. Banks follow the same logic: plants that can prove outcomes with auditable monitoring secure better rates; those that cannot pay a premium for capital.

For engineering SMEs, the route to market is narrower but clearer. Standalone equipment sales rarely close deals. Winning work increasingly depends on pairing treatment hardware with monitoring, alarms, and long-term service, because performance must hold through spot audits, seasonal energy spikes, and the direct financial consequences that follow any slip in quality or uptime.

3.3.2.2 Change drivers and dynamics

Demand strengthens because drought is no longer episodic. Multi-year shortages in key basins raise the urgency—and the willingness to pay—for reuse, while forcing operators to manage more difficult water chemistry, including higher salinity. Buyers therefore prioritize systems that can stay stable under stress, not just meet targets on commissioning day.

Revenue signals reinforce that shift. Reclaimed water prices increasingly track delivered quality and uptime, turning reliability into a source of margin and a source of exposure at the same time. Downtime and performance drift translate directly into reduced income or penalties, pushing customers toward solutions that can be continuously verified.

Project pipelines are shaped by how smoothly suppliers and operators can work with authorities. Permitting speeds up when applications are submitted online and then backed by ongoing sensor data, which favours suppliers that bundle instrumentation, reporting, alarms, and responsive service. In parallel, compliance for micro-pollutants becomes more uniform: the same lab panels and methods are required, results are uploaded to a shared database, and spot audits become more common. The practical challenge becomes keeping day-to-day sensor signals aligned with what the lab confirms.

Commercial models and cost discipline evolve in parallel. Buyers increasingly opt for multi-year leases with uptime guarantees, shifting ownership and performance risk toward suppliers—often SMEs—who must fund equipment upfront and manage lifecycle outcomes. Lenders amplify this by offering better loan terms to projects that can prove results with monitoring data. Operating economics remain decisive: energy prices stay within a moderate band but spike in winter, making efficiency upgrades and fixed-price operations contracts more attractive.

Execution is further shaped by supply constraints and competitive dynamics. Multi-year component deals make lead times more predictable, but raise prices and limit the ability to switch suppliers midstream. At the same time, Europe's preference for local sourcing adds checks and fees to imports, increasing landed costs for non-local equipment. Competition expresses less through aggressive underbidding and more through evidence: providers with verified removal and uptime data can sustain premium pricing, while offers that cannot substantiate performance struggle to gain trust.

3.3.2.3 Day in the future

A typical 2030 project starts on a portal rather than in a binder. The permit application is filed digitally, and the decisive section is often the commitment to share live quality data once the plant is operating. Approvals move faster when regulators can see performance continuously instead of waiting for periodic reports.

After commissioning, the commercial rhythm is blunt. Payments depend on what the site actually delivers and how reliably it stays online. A quality drift or an unplanned outage is no longer an internal problem—it immediately shows up in revenue and contract consequences. Remote monitoring becomes the operational nerve center, and service is organized around shortening deviations and protecting uptime.

Compliance runs on two tracks. Sensors steer daily control, but standardized laboratory panels remain the official record. Results must be uploaded to a shared database, and more frequent spot audits raise the cost of any mismatch between sensor trends and lab confirmations. Plants that keep those signals consistent move through scrutiny with fewer interruptions; plants that cannot face questions, corrective actions, and avoidable downtime.

Upstream conditions improve where industrial contributors follow the new enforcement approach. Key industries are required to self-report online and provide sensor data, with audits aimed at anomalies rather than routine visits. Where reporting is consistent, municipal plants see fewer shocks and can run tighter processes; where it is not, risk concentrates around the outliers.

The deal structure ties everything together. Many buyers choose a 5–10 year lease with uptime guarantees, expecting the supplier to finance equipment and manage performance over time. Financing terms echo that expectation: better rates follow plants that can prove outcomes with auditable data. Delivery planning is shaped by long-term supplier commitments—lead times are steadier, but choice is narrower and switching is costly. With imports facing extra checks and fees, procurement increasingly favours supply chains that are locked in and EU-aligned, and that can localize sourcing or assembly where possible.

3.3.2.4 Outcomes

By 2030, market access is determined less by claims and more by evidence. Providers that can demonstrate reliable removal and high uptime with credible monitoring data can sustain premium pricing, while low bids without verification are treated cautiously and often fall apart during diligence. The same logic governs approvals and cash flow: faster permitting and stronger revenues accrue to plants that can both sustain performance and document it continuously.

Risk shifts toward suppliers, especially SMEs. Multi-year leases with uptime guarantees become a common route to sale, but require suppliers to retain ownership, fund equipment upfront, and carry long-term performance obligations. Preferential loan terms for data-proven projects can reduce financing pressure, but they do not remove the operational burden: performance must be maintained day after day.

Cost control remains central. With energy prices fluctuating moderately but spiking in winter, buyers favour efficiency-focused retrofits and fixed-price operations arrangements that convert volatile bills into predictable service economics. Suppliers who cannot manage energy intensity or price exposure find margins harder to protect under performance-linked revenues and service commitments.

Supply and cross-border delivery become more orderly, but less flexible. Multi-year component lock-ins stabilize schedules yet raise costs and limit pivots when preferences change. Europe's tilt toward local sourcing adds checks and fees to imported equipment, making EU-aligned procurement strategies more competitive.

In Northern Africa, project volumes expand where cooperation is driven by security priorities and backed by EU funding, but access comes with conditions: local assembly and training are expected, and delivery depends on meeting local participation requirements. Across both regions, trust converges on a single currency—verifiable performance—because regulators, lenders, and customers all rely on the same evidence to decide who gets approved, financed, and paid.

3.3.3 Scenario 3: Fragmented markets and uneven regulatory pace

3.3.3.1 Introduction

By 2030, the advanced wastewater treatment market across Europe and Northern Africa does not move as one market. What sells in one city is stalled in the next. Reclaimed-water tariffs are set locally and vary sharply, some places allow a workable return, others cap prices so tightly that projects struggle to pencil out. Permitting follows the same uneven pattern: a handful of authorities turn approvals around quickly, while others stretch decisions past a year and demand extra studies. The result is a lopsided map of momentum, with bankable demand clustering in a few jurisdictions and long pauses elsewhere, even where the technical case for reuse is well understood.

Operating and delivery conditions add another layer of uncertainty. Energy remains expensive, and supply contracts shift often enough that buyers treat electricity risk as a core design constraint, not an afterthought. At the same time, getting equipment to site is increasingly a test of planning discipline. Components sourced within the EU tend to arrive on time, but cross-border shipments into North Africa face longer customs queues and higher freight costs. Product compliance requirements also diverge across the region, adding paperwork and rework risk. For suppliers, especially engineering SMEs, this combination turns “selling a system” into managing a chain of approvals, operating-cost exposure, and logistics reliability.

3.3.3.2 Change drivers and dynamics

Demand is pulled forward less by a single, sweeping regulatory wave than by climate stress and the way compliance is measured. Weather patterns swing hard: dry spells tighten water availability, then intense storms flush unpredictable loads through sewers. That volatility changes what customers value. They look for treatment that stays stable through sudden shifts in influent quality, not just equipment that performs well under steady conditions.

At the same time, oversight of micro-pollutants becomes more about averaged reality than continuous dashboards. Monthly passive samplers, verified by labs, increasingly define whether a plant is “in compliance”. Real-time sensors still matter, but mainly to keep processes within safe operating ranges and trigger alarms, rather than to prove compliance minute by minute. This nudges technology choices toward solutions that deliver consistent average performance and can be defended with clear lab results.

On the buyer side, constraints in municipal balance sheets reshape procurement and, with it, who can realistically win. Tight budgets push tenders to slip, scope to shrink, and contracting to tilt toward pay-per-month structures rather than large upfront purchases. In parallel, build-own-operate players bundle multiple sites and buy treatment modules in bulk, concentrating purchasing power and standardizing what they are willing to accept. Large suppliers respond with low-priced modular packages tied to financing, intensifying competition and squeezing margins, especially in decentralized applications.

Meanwhile, influent management becomes a market in its own right. Municipal utilities increasingly protect their networks by imposing surcharges or restricting problematic industrial discharges at the connection point. Industries that want to avoid service limits and higher fees are pushed into installing their own pre-treatment, creating a steadier industrial demand stream that is less dependent on the timing and politics of municipal tenders.

Across all of this, carbon costs expand only unevenly and rise modestly. They show up in bid evaluations and internal business cases, but rarely become the single deciding factor, reinforcing a market where energy efficiency and operational predictability matter more than headline carbon signals.

3.3.3.3 Day in the future

In one European region, reclaimed-water tariffs are high enough to support investment, but the rules differ just across the border, caps and recovery formulas change with the municipality. A build-own-operate provider uses that local clarity to move ahead, issuing a bundled procurement for several small sites at once. The contract is written to match strained public budgets: instead of a large capital purchase, it asks for a pay-per-month treatment outcome with clear performance guarantees and a commissioning schedule that can't slip, because the next dry spell is already being planned for.

An engineering SME comes in hoping to offer a standardized module. The first obstacle is not technical, it is documentation. Approval requirements vary by jurisdiction, and even within the same country the operator asks for region-specific studies and design tweaks to satisfy local reviewers. The compliance discussion is also specific: the operator cares most about whether monthly lab-verified sampler results will hold under variable loads. Dense real-time monitoring is welcomed as a way to protect operations, but it does not carry the compliance weight it once promised.

Then the cost screen tightens. With energy remaining expensive and contracts changing frequently, the operator interrogates power consumption, redundancy choices, and expected operating ranges. Competing bids arrive with aggressively priced modular packages paired with financing, forcing the SME into uncomfortable trade-offs: either meet the price and accept thinner margins, or differentiate through service terms, uptime commitments, and operational risk-sharing.

A few months later, the same SME supports a North African pilot that exists because of a discrete co-funded cooperation effort. The technical scope is manageable; the delivery plan is not. EU-sourced components move quickly to the export gate, but customs procedures and compliance differences on arrival stretch timelines and raise costs. The team carries extra inventory and dual-sources where possible to avoid a single delayed part halting commissioning. The pilot reaches operation and produces credible results, yet it remains what it started as: a one-off project that proves a concept without automatically unlocking a broader, repeatable pipeline.

3.3.3.4 Outcomes

By 2030, the market grows, but in bursts rather than as a smooth curve. Local tariff choices and uneven approvals produce a geography of winners and waiters: some cities can move projects from concept to contract, while others linger in reviews or cannot make the economics work. This fragmentation limits cross-border scaling and rewards suppliers that can repeatedly repackage the same technical core into different regulatory and commercial formats.

Procurement becomes more concentrated and more service-like. Budget pressure reduces the size and frequency of traditional municipal tenders and increases pay-per-month contracting. Build-own-operate players bundle sites and purchase in bulk, tightening specifications and putting smaller vendors under negotiating pressure. In decentralized segments, aggressive pricing by large firms, often tied to financing, compresses margins and makes standalone equipment sales difficult unless paired with credible long-term service capability.

Technology and operating models are shaped by two practical realities: volatile influent conditions and energy risk. Systems that can maintain stable performance through swings in load and water quality command attention, particularly as dry periods are interrupted by intense storms. Yet energy costs remain high and unpredictable, penalizing energy-intensive solutions unless efficiency is demonstrated in operating data and reflected in credible cost forecasts.

Compliance practices reinforce this focus on consistency. Monthly passive sampling verified by labs becomes the accepted proof point for micro-pollutant performance, steering discussions toward average-load outcomes and defensible laboratory evidence, while on-site sensors retain a supporting role as process protection rather than the centrepiece of compliance reporting.

Cross-border delivery, especially into North Africa, stays possible but heavier to execute. Faster access to EU-made components helps within Europe, but customs delays, divergent compliance requirements, and persistent trade friction increase buffers, inventory needs, and project-management overhead. Cooperation between the EU and North Africa does produce a handful of co-funded pilots and some standards work, yet it remains selective and project-by-project, limiting the emergence of a predictable, scaled regional pipeline.

3.3.4 Scenario 4: Stop–start surges and rapid response permitting

3.3.4.1 Introduction

By 2030, the advanced wastewater treatment market across Europe and Northern Africa grows in spurts, quiet quarters punctuated by sudden, high-pressure bursts. Brief heatwaves repeatedly trigger strict water-use restrictions, and the moment taps tighten, reuse moves from “nice to have” to “procure now”. That urgency reshapes the permitting environment. Authorities respond with fast, temporary approvals to keep industry and cities running, then tighten requirements after the first season. As a result, buyers stop betting on single, fixed designs and instead look for modular upgrades they can install quickly, prove in operation, and expand when the regulatory bar rises.

Commercially, willingness to pay is uneven. Industrial users accept higher tariffs in exchange for reliable reclaimed water, while household prices are kept low to protect public support, good politics, but it leaves many urban schemes short of full cost recovery. Meanwhile, compliance starts to feel less like a one-off inspection and more like an ongoing, insurable obligation. Regulators increasingly accept insurance-backed compliance bonds, and those policies come with strings attached: audits, continuous sensor data, and documentation that can stand up to scrutiny. For engineering SMEs, opportunity exists, but it arrives on short notice and increasingly rewards those who can deliver fast, instrumented, and upgrade-ready systems.

3.3.4.2 Change drivers and dynamics

The first and loudest driver is the stop-start rhythm created by recurring heatwaves and restrictions. When water limits hit, projects are launched under urgent timelines, enabled by rapid, temporary permitting. But the same cycle produces a second wave of demand: once the immediate crisis passes, the rules tighten and early deployments need additional treatment steps, monitoring, or capacity. This favours suppliers that can start small, commission quickly, and build a credible path to expansion without tearing systems apart.

The second driver is the shifting economics of reuse and decarbonization under tight capital. Premium industrial tariffs can support investment decisions, especially where reliability is mission-critical, while discounted household pricing makes purely municipal business cases harder to close without careful structuring. Exporters and energy-heavy sites face added cost pressure from border-linked carbon rules, pushing them toward lower-carbon upgrade options even when carbon pricing remains uneven across sectors. Financing, however, is selective: banks prefer proven upgrades, interest remains elevated, and higher equity requirements slow projects led by smaller players and reduce the number of “experimental” technology bets.

A third set of drivers raises execution risk and reshapes the upgrade calendar. Supply chains improve compared with earlier years, yet periodic shortages, especially for membranes and control components—still create delays and cost spikes. Intermittent security flare-ups bring sudden border holds and higher insurance, which can pause projects and limit contractor travel at the worst moment. At the same time, industrial discharge enforcement follows a predictable arc: a compliance amnesty runs through 2028, then targeted crackdowns begin in 2029–2030 as limits tighten. That transition pulls demand toward pre-treatment and more robust designs, particularly where variable influent quality could jeopardize permits, insurance coverage, or production continuity.

3.3.4.3 Day in the future

In a typical heatwave year, the warning signs arrive early: reservoir levels drop, restrictions snap into place, and utilities and industrial sites scramble for alternatives. Regions issue temporary permits for reuse to stabilize supply, and procurement teams compress timelines that used to take months into weeks. The first question is no longer "What is the perfect end-state plant?" but "What can we install quickly that won't be obsolete after the first season?" Buyers lean toward solutions that can be commissioned fast, then expanded when requirements tighten and monitoring expectations rise.

Inside the plant, compliance has a different feel than it did a few years earlier. Operators secure insurance-backed bonds to cover fine exposure, and insurers insist on audits and continuous sensor data before they underwrite risk. That makes instrumentation, data integrity, and reporting routines central to operations, not optional add-ons. Commercially, most deals still look familiar: systems are purchased upfront, but they increasingly come bundled with multi-year monitoring and maintenance contracts that keep suppliers involved long after commissioning. A few outcome-based pilots appear, but they remain the exception rather than the rule.

Operating costs remain hard to pin down. Energy is usually inexpensive, which supports advanced treatment, but sudden policy-driven shocks still hit without much warning. Buyers respond by prioritizing flexible designs that can shift loads and manage peaks rather than locking themselves into a single energy profile. In Northern Africa, the market is more fragmented: regional cooperation slows, each country pursues its own plan, and standards vary, raising entry costs and increasing the need for local adaptation. Large multinationals deepen their presence through local partnerships and informal carve-ups, prices stay relatively steady, but access narrows for independent SMEs. Add periodic component shortages and occasional border holds, and it becomes clear why customers increasingly prefer vendors who can prove delivery reliability and sustain service continuity when conditions change mid-project.

3.3.4.4 Outcomes

By 2030, the market is larger, but it does not behave like a steady growth story. Reuse deployments accelerate in restriction-driven surges, and the emergency-permit pattern creates built-in follow-on work: systems installed quickly under relaxed conditions often require upgrades once the tighter post-season rules take effect. Where money is available, it concentrates in industrial reuse, because reliability commands higher tariffs. Urban schemes keep household prices low to maintain legitimacy, but that political choice limits full cost recovery and constrains how aggressively cities can scale without careful financing and phasing.

Late in the horizon, the enforcement timeline becomes a demand catalyst. After the amnesty period, targeted crackdowns in 2029–2030 drive spending on industrial pre-treatment and resilience upgrades, especially where influent variability can undermine performance and raise liability concerns. Across both municipal and industrial operators, compliance becomes more uniform and more data-dependent. Insurance-backed bonds normalize audited performance, pushing suppliers toward robust monitoring, documentation, and service capability. Procurement reinforces that shift: most buyers still pay upfront, but multi-year monitoring and maintenance contracts become standard expectations rather than differentiators.

Decarbonization pressures shape technology choices in parallel. Export-linked carbon costs and audits make energy-heavy sites more sensitive to the carbon profile of their wastewater systems, strengthening the case for low-carbon upgrade paths even though energy is generally cheap most years. The constraints remain tangible: selective green lending with higher equity needs slows project volume and filters out less-proven approaches; multinational partnership strategies keep pricing stable but reduce the number of open, standalone tenders SMEs can win without alliances or a sharp niche. Operationally, periodic shortages and sporadic border holds continue to add schedule and cost risk, reinforcing buyer preference for modular rollouts, staged commissioning, and suppliers that can navigate delivery disruptions without derailing compliance or production.

4.0 Opportunities and threats

This section translates the workshop outputs into a structured view of where value can be created and where exposure exists. Opportunities and threats were derived in two stages. First, each company developed scenario-specific implications using a facilitated, morphological scenario process supported by ScenAIrios for consistency analysis and rich narratives. Second, in strategy design workshops, participants mapped opportunities and threats per scenario, then isolated those that were valid across all scenarios. These became the targets for strategic measures in T 5.4: to exploit opportunity domains and avoid or mitigate threats that could derail uptake and scale.

The section is organised as follows. The first subsection synthesises the entire pool of opportunities across companies and scenarios into thematic clusters that reveal the big picture—systemic demand drivers, procurement logics, capability advantages, and value propositions that recur. It grounds every cluster in the underlying scenario logics to retain interpretability and to anonymise firm-specific data. The second subsection presents the scenario-independent opportunities (one by one), explaining why they hold across divergent futures and how they connect to the roadmap lanes we build next. Threats will be addressed after opportunities, in a separate step, mirroring this structure.

The value to iMERMAID and its beneficiaries is twofold. First, the synthesis provides actionable focal points for exploitation strategies, business models and partnering, coherent with evolving EU policy on water reuse, micropollutants and circular raw materials. Second, it anchors opportunity narratives in credible regulatory and market signals, from the EU Water Reuse Regulation (in force since June 2023) to PFAS and mixture-risk oversight, Extended Producer Responsibility for quaternary treatment, and the Critical Raw Materials Act, which shape investment, certification, and procurement.

4.0 Opportunities – Cross company, cross-scenario analysis

Across the futures examined in the workshops, one pattern repeats: buyers select what can be proven, installed quickly, and kept stable under stress. Opportunities therefore concentrate in three connected spaces: **modularity and speed of delivery, evidence and certification as currency, and selectivity and circularity where recovery reinforces resilience,** and they are animated by procurement behaviour, regulatory mechanics, and operating realities that recur across very different market moods.

The first space opens wherever time pressure and capital discipline meet. Heatwaves and seasonal restrictions compress procurement windows; surges in demand collide with long lead-times for membranes, UV units, sensors and controls; authorities issue temporary approvals that are followed by tighter rules in the next season. In such conditions, the opportunity lies in **modular, upgradeable designs** that can be commissioned within weeks, then expanded or re-configured as requirements rise. Futures that feature phased commissioning, framework tenders with explicit performance penalties, and standardised templates for lower-risk reuse create the same pull: **phased delivery** becomes a procurement filter. The opportunity is not just technical; it is also commercial. Providers that pair skids and cartridges with clear "phase-in" playbooks and outage-free tie-ins qualify for fast-track projects and, later, for the follow-on work when the bar tightens. National programmes that push reuse toward routine infrastructure, such as France's plan to lift reuse from roughly one percent today to a target ten percent this decade, reinforce this pattern by creating bursts of time-sensitive tenders that favour ready-to-deploy modules over bespoke one-offs.

The second space grows where compliance and cash flow depend on **evidence**. In several futures, approvals move faster when live quality data is committed upfront, financing terms improve when outcomes can be audited continuously, and revenues are linked to delivered quality and uptime rather than nameplate performance. Where monitoring intensity rises without a parallel explosion of new limits, risk-based systems still push operators to invest in what **reduces the chance of being flagged**. In both cases, **evidence becomes currency**. The opportunity here is to bundle core treatment with instrumentation, secure data handling, and **certification pathways** that reduce diligence friction for buyers and lenders. It is also to design processes that keep lab-confirmed results consistent with online signals, because the cost of divergence increases with standardised panels and public dashboards. These futures mirror real policy motion: the Drinking Water Directive's PFAS monitoring requirements became applicable in January 2026; the EU PFAS restriction process is on course for opinions in 2026; and the Commission's mixture-risk agenda is pushing effect-based monitoring and cumulative assessments into practice. Each element deepens the value of offers that can **prove** performance, align with certification, and make reporting auditable.

A third space emerges where **selective recovery** supports security-of-supply narratives and where **low-waste residuals** simplify compliance. In trade-volatile conditions, recovered outputs shift from "nice to have" to **hedges against supply risk**, provided quality can be certified and by-product status clarified. In quieter monitoring-centred settings, certification still decides whether concentrates become tradable products or legal headaches. The opportunity is to position recovery as disciplined risk management: **tailored selectivity, traceability, and off-take quality assurance,** linking it to the EU's Critical Raw Materials Act benchmarks for domestic recycling capacity and emerging standards for secondary-materials traceability. Where resource pressures and compliance overlap, recovery modules attached near the source turn into practical instruments for de-risking operations and for meeting circularity targets without over-complicating waste obligations.

Threaded through these spaces is an operating reality that creates further opportunity for those who master it: **performance under variability and mixtures**. When drought alternates with storm flushes, or when pre-treatment enforcement reduces shock loads but raises salinity, buyers value systems that hold average performance under unstable influent. Where oversight shifts toward mixture-based outcomes, advantage accrues to processes that **handle parents, transformation products, and complex matrices** while keeping energy and residuals under control. This favours hybrid trains, robust control strategies, and verification methods that connect online control to lab-defensible mixture proxies. The mixture-risk science and the Commission's programme on combined exposures provide the regulatory and methodological backdrop for this shift, and they help explain why offers that demonstrate **stable performance across variability** show up as opportunities in multiple futures.

Across all of the above, opportunities expand when providers move beyond equipment sales to **service-centric models** that align payment with proof. In futures where revenues track delivered quality and hours online, or where insurers and lenders require risk management plans and continuous evidence, performance-linked service becomes a competitive advantage. The market logics at play, fast permits contingent on live reporting, grants channelled to proven retrofits, multi-year leases with uptime expectations, favour suppliers who can **finance, monitor, and stand behind performance** while keeping lifecycle costs predictable.

In sum, the opportunity landscape that emerges from the explored futures is not a scatter of one-off ideas but an integrated design and delivery thesis: **modular systems commissioned fast, evidenced continuously, specialised where selectivity and circularity create value, and supported by service and financing models that turn proof into revenue**. The policy environment now in force or advancing, reuse regulation, PFAS and mixture-risk oversight, EPR financing of quaternary treatment, and critical-materials resilience, makes that thesis actionable and durable beyond any single company or country.

4.1 Opportunities – Scenario-independent set

4.1.1 Modular adaptive treatment

This opportunity holds across futures because delivery speed and future-proofing recur even when the triggers differ. Seasonal restrictions and temporary approvals create short commissioning windows; supply-chain stress makes phased installations credible; framework contracts with penalties reward upgrade-ready configurations. Modularity therefore reduces approval risk, delivery risk, and adaptation risk simultaneously, and it fits the burst-like procurement patterns seen where national reuse plans move from intent to projects.

4.1.2 Brownfield-first, retrofit-ready entry

Budgets remain tight in multiple futures, even when enforcement strengthens. Buyers favour technologies that “drop in” without civil overhauls, defer large capital, and meet phased deadlines. The EPR mechanism in the recast wastewater directive further channels funds toward quaternary upgrades at existing plants, making retrofit logic a structural feature rather than a transient preference.

4.1.3 Performance-based service revenue

Where approvals, finance and customer payments respond to verified outcomes, monitoring and service become integral to the offer. Revenues linked to delivered quality and uptime de-risk cash flow for buyers and create recurring revenue for suppliers, provided data integrity and lab alignment are maintained. This model generalises because the drivers, risk-based oversight, lender diligence, and reputational exposure, appear across futures with different intensity.

4.1.4 Low-OPEX, low-residuals processes

Operating cost discipline is a constant, whether energy spikes are episodic or persistent. At the same time, rules and practice around PFAS and hazardous residuals raise the cost of poorly managed concentrates. Solutions that keep energy intensity low and minimise problematic wastes gain advantage in both lenient and strict futures.

4.1.5 Broad-spectrum and mixture-robust treatment

As lists expand from familiar pharmaceuticals to industrial additives and as mixture-risk thinking advances, futures reward technologies that manage parents, transformation products and complex matrices. Even when numeric limits do not tighten, risk-based scoring and effect-based approaches push operators to value demonstrable mixture control. This makes spectrum breadth and mixture robustness a cross-scenario asset.

4.1.6 Partnership-driven market access

Channel power and diligence demand favour partnerships with utilities, OEM platforms and service integrators. In several futures, acquisition or framework alignment becomes the route from pilots to scale. This generalises because the underlying reasons—procurement risk, delivery reliability, certification readiness—do not disappear as policy shifts; they intensify.

Service-based deployment models (leases/PaaS).

When buyers face constrained CAPEX yet must demonstrate outcomes, leases and product-as-a-service models align incentives: payment follows proof, and lifecycle performance is managed with remote monitoring. The same logic appears in different guises across futures, which is why the model travels.

4.1.7 Tailored selectivity and recovery as resilience

Recovery is not a universal panacea, but futures that value risk reduction and circularity create stable niches for selective modules with quality-assured outputs. Certification and traceability convert recovery from a regulatory headache into a tradable product; the Critical Raw Materials Act strengthens the narrative by setting recycling benchmarks along the value chain.

4.1.8 Resource-security positioning

Supply-chain vigilance remains part of policy and buyer behaviour. Framing recovered outputs and selective capture as risk mitigation aligns with CRMA targets and Member State resilience strategies. This positioning holds whether trade is volatile or merely scrutinised, because the dependency question persists.

4.1.9 Expansion through reuse policy implementation

As reuse regulation moves from law to projects, Member States roll out national plans with decentralised deployments and fast-track approvals. This creates repeatable openings for compact, modular systems, especially where industrial users accept higher tariffs for reliability. The signal is strongest in countries pressing toward 2030 targets.

4.1.10 Low-waste residuals as a selling point

Where hazardous classification spreads and management plans become obligatory, minimising problematic residuals reduces both cost and risk. This resonates across futures because waste handling is a fixed constraint, not a trend.

4.1.11 Specialisation under influent variability

Volatile influent conditions are a shared feature: drought and storm cycles, salinity shifts, and pre-treatment enforcement all change operating envelopes. Technologies and control strategies that hold average performance under variability, and that keep online and lab signals aligned, therefore generalise across contexts.

4.2 Early opportunity-driven insights for roadmap development

These opportunities define the lanes and milestones the roadmap will prioritise: modular product architectures and phased delivery, embedded monitoring and certification pathways, service and finance models aligned to outcomes, and R&D that targets mixture robustness, energy efficiency, and low-waste residuals. They also point to geographies and segments where policy converts intention into procurement: reuse roll-outs under Regulation 2020/741 and national plans; quaternary treatment funded through EPR under the recast UWWTD; PFAS and mixture oversight that makes evidence decisive in approvals and revenues; and circular-materials frameworks that turn recovery into resilience rather than speculation.

4.3 Threats – Cross-company, cross-context synthesis

Across the futures examined in the workshops, the same forces that create market pull also generate strategic exposure. Three families of risk recur: **procurement and delivery dynamics that favour scale and speed; evidence-centred compliance that shifts liability and workload onto supplier**s; and **fragmentation, in regulation, markets, and standards, that traps value before it compounds**. These are not abstract categories; they are the practical forms that policy and market change can take when they touch real projects.

In one set of futures, reuse accelerates and compliance becomes routine, but the **industrialisation of procurement** quietly tilts the field. Framework tenders reward serial production, standardisation, and assured delivery; emergency approvals harden into tight commissioning windows; module serialization and multi-year component lock-ins privilege larger manufacturers with deeper inventories and working capital. Smaller providers risk being priced out or delayed at precisely the moment the market opens. Even where national programmes aim to decentralise and speed deployment, as in France's push to expand reuse rapidly this decade, the mechanics of delivering many small projects at once can reward those who already operate at scale. The threat here is not policy reversal but **policy success without inclusive channels**.

In a second set, the centre of gravity shifts from claims to **verification and liability**. Approvals move faster when operators commit to live sharing of water-quality data; lenders discount projects with continuous evidence; revenues track delivered quality and hours online. This creates a sturdier market, but it also **transfers performance risk** to suppliers. Contracts link penalties to uptime and delivered outcomes; insurers demand audits and data retention; any drift between lab confirmations and online signals becomes a commercial problem, not just an operational one. With PFAS monitoring obligations taking effect under the Drinking Water Directive and restriction processes advancing toward 2026 decisions, the burden of proof becomes heavier and more expensive to carry for those without mature monitoring, data integrity, and service capacity. The market rewards those who can stand behind outcomes, and exposes those who cannot.

A third family of threats appears where **policy tightens unevenly**. Local tariff formulas diverge; permitting and documentation requirements vary by jurisdiction; inspection frequency and penalty practices differ within a single economic area. In such conditions, companies face **bureaucratic overhead** and rework risk as designs are tailored city by city. At the same time, **pilot-rich but pathway-poor** funding environments multiply proof-of-concepts without clear conversion to stable operating spend; pilots become reputational necessities rather than launchpads. Where recovered materials are concerned, missing or inconsistent product standards and by-product classifications raise **liability concerns** and stall off-take. The Critical Raw

Materials agenda makes circularity a strategic goal, but value fails to compound if certification and recovered-product quality assurance lag.

Execution risks weave through all of this. Surges tied to heatwaves, drought or enforcement waves collide with periodic shortages of membranes, sensors and control components; border holds and permit checks interrupt logistics; manufacturing capacity and skilled labour become bottlenecks just when urgency peaks. In these moments, **supply-chain instability** and **insufficient production capacity** convert promising tenders into missed windows. When the procurement calendar is compressed, the penalty for delay is not just lost revenue, it is loss of place in the queue.

Finally, technical boundary conditions do not disappear simply because policy wants results. High organic loads and variable influent quality depress performance in some advanced processes; pre-treatment enforcement reduces shocks yet raises salinity, tightening control margins; mandatory pre-treatment or pre-filtration steps elevate capital costs. Where residual destruction or hazardous classification requirements toughen, solutions that transfer rather than degrade pollutants face rising operating complexity and cost. In short, **physics and chemistry still set the price of admission**, and markets that pay for proof punish optimistic sizing or incomplete residuals plans.

Read together, these threads define a coherent risk landscape: procurement logic that can outgrow small-scale delivery models; verification economies that move liability onto suppliers; fragmentation that absorbs capacity in paperwork rather than performance; supply-chain and production limits that turn surges into stress; and technical constraints that discipline ambition. Each of these elements is visible in the futures explored, and each is reinforced—not weakened—by policy trajectories currently in motion.

4.4 Threats – Scenario-independent set

4.4.1 Scale-driven displacement in serialised procurement

As reuse and quaternary upgrades become programmatic, tenders favour serial production, assured delivery and standardisation. These dynamic advantages large platforms and can marginalise smaller vendors even when technology fit is strong. Mitigation requires alliances, framework supply, and modular designs that truly interoperate.

4.4.2 Risk concentration under evidence-linked contracts

Outcome-based arrangements convert performance deviations into direct revenue loss or penalties, while insurance and lender requirements add audit and documentation burdens. Suppliers that cannot maintain lab–sensor alignment and service responsiveness face cumulative exposure. Robust monitoring, data governance, and insurer-aligned maintenance regimes become non-negotiable.

4.4.3 Administrative drag from regulatory fragmentation

Divergent approvals, documentation and inspection practices create rework and delay, especially across national borders. The cost is felt in engineering hours, legal reviews and slipped commissioning, not only in permit fees. Standardised application packs and local compliance partners reduce cycle time volatility.

4.4.4 Pilot-to-nowhere risk

Funding that favours demonstration without post-pilot conversion pathways traps capacity in trials and undermines business cases. Clear pilot-to-operation designs, off-take commitments, and performance-based step-ups are required to translate evidence into contracts.

4.4.5 Supply-chain instability and capacity bottlenecks

Crisis-driven surges and enforcement waves collide with periodic shortages of critical components and manufacturing slots, turning tenders into missed windows. Diversified suppliers, regionalised assembly, and capacity trigger plans are essential hedges.

4.4.6 Technical efficacy under variable influent and high organics

Fluctuating loads and elevated organic matter can suppress removal in advanced processes, increasing the risk of underperformance where continuous verification is expected. Hybrid trains, robust controls and conservative design margins are required where variability is structural.

4.4.7 Residuals complexity and hidden CAPEX

Mandatory pre-treatment or residual destruction requirements lift capital and operating costs; hazardous classification or unclear by-product status inflates handling and liability. Designing for low-waste outcomes and early engagement on classification reduces downstream surprises.

4.4.8 Financial exposure in performance-driven, fragmented markets

Uneven tariffs, variable approvals, performance-linked revenues and higher equity needs combine into unpredictable cash flows and longer paybacks. Financing structures that match evidence cadence and risk-sharing with buyers become a condition for growth.

4.4.9 Standards and certification gaps for recovered products

Where product standards and certification routes lag, recovered outputs remain hard to trade and raise liability concerns, stalling circular value. Aligning QA with emerging traceability standards and sector specifications is necessary to unlock off-take.

4.4.10 Competition from incumbent integration

Established providers can bolt selective steps onto existing trains and occupy the same compliance niche, compressing differentiation and price. Channel partnerships and undisputed performance evidence are required to sustain a place in the specification.

4.5 Early threat-driven insights for roadmap development

These threats point to concrete roadmap design choices: alliances and framework supply to counter scale bias; embedded monitoring, documentation and insurer-aligned service to manage evidence-linked liability; compliance toolkits and local partners to reduce administrative drag; regionalised assembly and capacity triggers to absorb surges; hybrid trains, conservative control envelopes and low-waste residuals strategies to meet technical boundary conditions; and financing models that stage risk and cash flows in line with verification rhythms and uneven tariff regimes. They also suggest where policy engagement has outsized value: harmonising certification and standards for recovered outputs and clarifying pilot-to-operation funding mechanics so that evidence reliably becomes deployment.

5.0 Roadmap

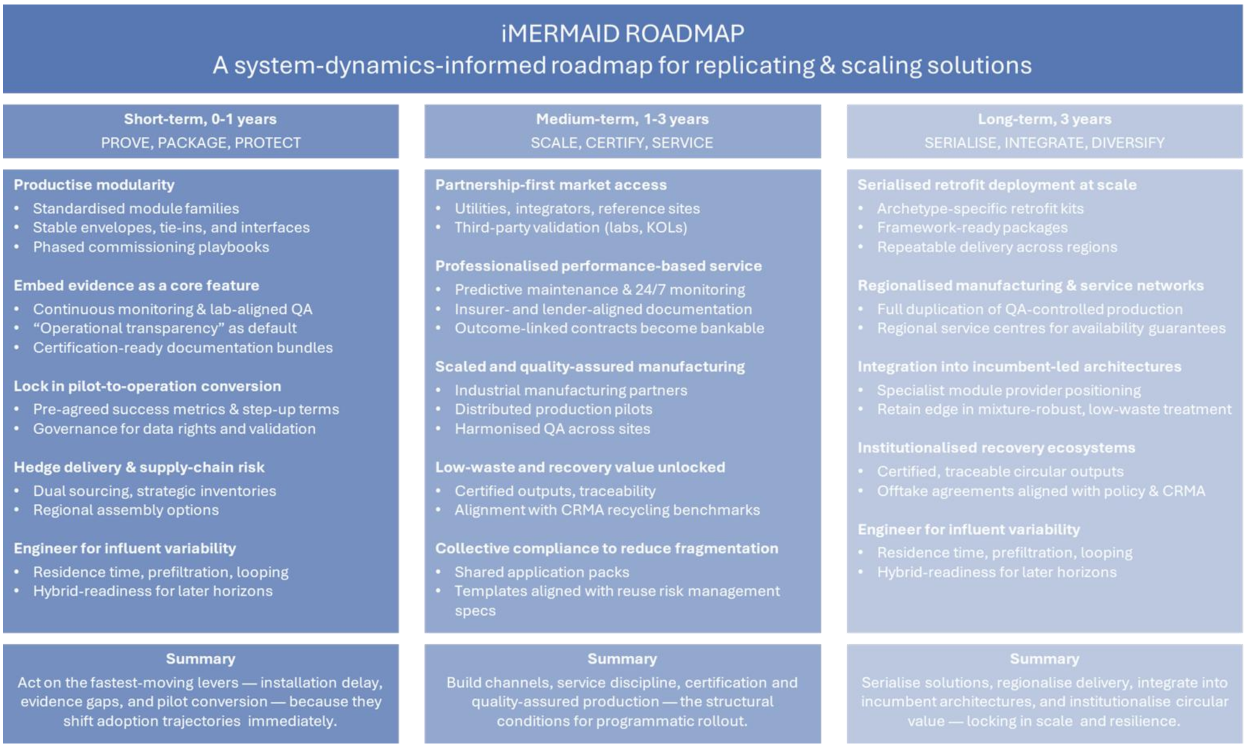

Figure 1: Roadmap for replicating and scaling iMERMAID solutions.

5.0 From opportunities and threats to targeted action

This roadmap translates the opportunity and threat landscape into a time-sequenced programme of action that a wider ecosystem can implement. It follows the synthesis method described earlier: company-specific outputs were compared, recurring patterns were grouped into consolidated categories, and differences in detail were resolved by focusing on the strategic intent of the measures rather than firm-specific wording. The result is a generalised roadmap aligned with iMERMAID's aims and grounded in the system-dynamics view of dependencies and leverage points established in Task 5.3. In particular, the modelling highlights how regulatory clarity and enforcement, evidence creation and diffusion, and reductions in installation delays act as high-leverage variables that shift adoption trajectories in wastewater markets. These insights are used throughout to prioritise measures that change stocks and flows which actually move the system.

The roadmap is presented in three horizons—0–1 years, 1–3 years, and 3+ years—because delivery rhythms, financing windows, and policy milestones differ meaningfully across these periods. Wherever relevant, the narrative situates actions in the EU policy context that already appears in the scenarios and in stakeholders' decision environments: the EU Water Reuse Regulation 2020/741 (applicable since June 2023) shapes risk management, permitting, and transparency for reuse schemes; the Drinking Water Directive recast introduces PFAS monitoring obligations from January 2026, making evidence central to oversight; work under REACH/POPs and the broad PFAS restriction proposal is advancing towards 2026 committee opinions, with compliance consequences for residuals and materials; the recast Urban Wastewater Treatment Directive (UWWTD) introduces Extended Producer Responsibility (EPR) to finance quaternary treatment, rebalancing costs and incentives; and the Critical Raw Materials Act (CRMA) establishes recycling benchmarks and resilience targets that strengthen the case for selective recovery and quality-assured secondary materials. These are not background notes; they define how procurement, certification, financing, and market access will actually work. [monacolife.net], [ecologie.gouv.fr], [europarl.europa.eu], [regartis.com], [iafcertsearch.org], [watereurope.eu], [waternewseurope.com].

The section integrates the measure set the workshops produced for each scenario-independent opportunity and threat, and interweaves those actions with the leverage points and feedback loops revealed by the system dynamics models for metal recovery, removal of contaminants of emerging concern (CoEC), and agricultural reuse. In doing so, it connects why a measure is needed (opportunity or threat), what it changes in the system (stock, flow, or loop), who must play which role for it to work (stakeholder responsibilities), and how it interacts with the broader regulatory environment.

5.0.1 Short-term (0–1 years) — prove, package, and protect

The immediate priority is to turn near-term opportunities into repeatable wins while neutralising the threats that most often stall or derail adoption: pilots that never convert, supply-chain shocks that break delivery promises during heat-wave or enforcement surges, and premature exposure to performance-linked liabilities before monitoring, service capacity, and documentation are ready. In this horizon, the most powerful moves are to industrialise modularity, to make evidence a product feature, to contract for pilot-to-operation before any site work, to hedge supply and production risks deliberately, and to engineer around boundary conditions like high organics and variable influent that otherwise degrade performance precisely when continuous verification arrives.

The centrepiece is packaged modularity for phased commissioning. Buyers increasingly prefer compact modules that can be installed fast, with tie-ins to existing trains and the option to phase commissioning when long-lead components arrive. The measures set points to the same place: define the technology as a modular product family for specific needs; design sub-modules that configure cleanly for different duty points; and make the modules manufacturable rapidly so that delivery promises hold under tight timelines. Crucially, this is not a call for ad-hoc flexibility but for a shift in mentality from bespoke projects to industrialised products with stable envelopes and interfaces. This shift increases the inflow into the "installed capacity" stock in the CoEC model, shortens decision–delivery delays, and dampens the price-demand oscillations that otherwise arise when supply cannot meet pilot-driven spikes. It also anticipates the serialisation bias that emerges as reuse becomes routine infrastructure later on.

In parallel, evidence must be bundled as part of the offer. As permits, financing, and even cash flows respond to proof of delivered quality and uptime, instrumentation, data integrity, and lab-alignment protocols become product features rather than optional extras. The measures articulate the posture succinctly: operate as an "operationally transparent" supplier that makes performance and quality data available, with secure pipelines and documentation built for certification. This reduces diligence friction for regulators and lenders and aligns with the Water Reuse Regulation's risk-management and transparency requirements and the Drinking Water Directive's PFAS monitoring obligations that took effect in January 2026. In system terms, credible evidence narrows the "customer demand and actions gap," increases the decision-to-invest rate, and stabilises the balancing loops where improving quality would otherwise reduce perceived pressure to invest.

A third, very practical move in the first year is to prevent pilot-to-nowhere dynamics. The workshops were explicit that pilots should not commence until the conversion path is contractually defined. That means agreeing in advance the success metrics, data rights and validation methods, and the commercial step-up to full scale or long-term service if the metrics are met. This measure changes the system's structure: it ensures that a portion of the "proof" stock is wired to flow directly into "installed infrastructure" rather than accumulating as non-converting demonstrations, which otherwise absorb scarce engineering capacity and create false signals of progress. In procurement terms, it also normalises conditional awards that are triggered by independently verified outcomes.

Because the most urgent windows coincide with supply-chain volatility, the roadmap calls for hedging delivery risk deliberately. Diversifying suppliers for critical components, maintaining small strategic inventories sized to seasonal water-stress patterns, and establishing regional assembly partnerships for final integration reduces exposure to border holds, customs delays, and sudden shortages of membranes, sensors, or control elements. When windows are compressed by temporary permits or enforcement waves, the penalty for delay is loss of place in the queue rather than a small schedule slip. In the models, smoothing these inflows reduces the amplitude of price and availability swings, important for maintaining decision momentum under the price-demand balancing loops.

Finally, engineering around technical boundary conditions avoids performance dips precisely when continuous verification and standardised audit panels expose weaknesses. The measures stress three near-term levers for streams with high organics or variable contaminant loads: increasing residence time, implementing prefiltration, and using internal looping to maintain removal efficiency. Where needed, design hybrids with membranes are planned into the next horizon when procurement and budgets allow. This keeps average treatment performance stable under variability, prevents compliance drift when mixture-risk thinking and effect-based approaches enter practice, and avoids unexpected cost or liability when residual classification tightens under PFAS and hazardous waste rules.

By the end of this first year, the goal is that modularity is productised, evidence is automated and auditable, pilots convert by design, supply and assembly capacity hold under stress, and process designs cope with variability. These gains are not only commercial; they are structural adjustments that change key feedbacks in the system the T 5.3 work described. They increase the credibility of investment decisions, bring installation and commissioning delays under control, and embed the evidentiary practices that lenders, regulators and insurers increasingly require.

5.0.2 Medium-term (1–3 years) — scale, certify, and service

With packaged modularity and credible evidence in place, attention turns to programmatic roll-out. Two lines of effort dominate this horizon. The first is to deepen market access through partnerships with system integrators and utilities who own the channels that matter for framework tenders and multi-site programmes. The second is to professionalise performance-based service so that revenue models linked to delivered quality and availability become reliable earnings rather than latent liabilities. Alongside these, manufacturing scales deliberately to meet surge windows without quality drift, and low-waste and recovery advantages are amplified and certified where possible to open circular value streams.

Partnerships come first because channel power concentrates in this period. Several measures are already seeded in the first horizon, collaboration at reference facilities; validation studies with academic and regulatory key opinion leaders; communication that leads with compliance value and traceability, and now become the default route to market. Channel partners shorten sales cycles, share responsibility for integration risks, and improve inclusion in framework tenders. The effect is visible in the models as increased perceived "ease of integration", which raises the decision-to-invest rate and offsets the balancing loops that otherwise slow adoption when first installations begin to reduce perceived pressure.

These partnerships also provide the venues where lab-confirmed panels and digital evidence streams are normalised with regulators, which matters for both permitting and project finance.

Performance-based service then moves from aspiration to operational discipline. Where payments, bonuses, or penalties link to delivered quality and uptime, suppliers need remote monitoring, predictive maintenance, and insurer-aligned documentation to manage risk, as well as trained after-sales teams who can shorten deviations when signals drift. The workshop corpus describes the posture as “operational transparency”, which fits both risk-based oversight and the Drinking Water Directive’s monitoring frame. In system terms, this capability reduces the discrepancy between desired and realised wastewater quality, thereby relaxing some of the regulatory pressure that pushes investments even as it increases trust, stabilises operations, and sustains recurring revenue without litigation.

Scaling manufacturing without losing quality is the third pillar in this horizon. To counter the threat of insufficient production capacity, the measures call for partnering with experienced industrial manufacturers for early-stage scale-up, recruiting seasoned production leaders, and piloting distributed manufacturing with harmonised quality assurance so that quality holds even when production is duplicated closer to surge regions. This is not over-engineering; it is a recognition that the market’s most valuable windows are uneven and often time-critical, and that quality drift in this period can destroy margin and reputation at once. In the price-demand loops, growing the installed capacity to produce dampens spikes and improves adoption by maintaining on-time delivery promises.

Because waste obligations tighten in parallel with programme roll-outs, mid-term actions must amplify the low-waste advantage and develop recovery niches where economics and certification allow. The measures suggest two tracks: making low-waste outputs an explicit selling point in markets sensitive to hazardous residuals, and building modular recovery cartridges where selective capture has a clear off-take route. This aligns with the CRMA’s recycling benchmarks and with circular procurement preferences that many public buyers adopt as they formalise long-run climate and resilience goals. Where recovered outputs enter commerce, aligning QA and chain-of-custody documentation with emerging traceability standards, for example the guidance in ISO 59014:2024 on sustainability and traceability of secondary materials, reduces liability and accelerates acceptance. In the T 5.3 metal-recovery model, such steps strengthen the commercial potential loop by adding a secondary value stream that justifies infrastructure investment beyond pure compliance.

Finally, administrative drag from regulatory fragmentation needs to be addressed collectively. The measures propose clustering with peers to create shared application and documentation packs that can be adapted across jurisdictions, beginning already in the first year and maturing in this horizon. This reduces time-to-permit, shares legal costs, and standardises arguments and evidence across agencies. It is directly reinforced by the Commission’s work on risk-management technical specifications for reuse under Regulation 2020/741, which gives applicants a stable template for analysis even when national practices differ. In the models, lowering bureaucratic variance reduces uncertainty in the decision to invest and keeps installation inflows steady rather than lumpy.

By the close of the 1–3-year period, the goal is to have partnership-first access as the standard, service discipline that makes outcome-linked contracts bankable rather than risky, manufacturing capacity that meets surges without drift, low-waste and recovery value propositions that are certified and recognised, and compliance toolkits that seriously reduce administrative overhead. These are the conditions required for the final horizon’s move to serialisation and regional duplication.

5.0.3 Long-term (3+ years) — serialise, integrate, and diversify

The long-term horizon is about moving from isolated roll-outs to programmatic deployment, while defending distinctiveness as large platforms push serialised offers through channel control. Three moves define the position here: serialise brownfield deployment without losing flexibility, regionalise manufacturing and service to match surge geographies, and integrate with incumbents to avoid channel exclusion while retaining a mixture-robust, low-waste edge. A fourth move, for solutions with credible circular value, is to institutionalise recovery and traceability so that resource-security positioning becomes part of standard procurement logic rather than a bespoke argument.

Serialising brownfield deployment means using compact, retrofit-friendly packages to win in the niches where they are structurally superior—sites with limited footprint, phased installation constraints, or distributed loads—and then standardising reference kits for each archetype. The measures envisage coupling these packages with service-based models that lower risk for buyers and make replacement and upgrade cycles predictable. In systems language, serialisation reduces installation delay across many sites at once, creating smoother goal-seeking trajectories toward target capacities in the agricultural-reuse model and reducing the probability that window-driven demand collapses into backlogs that exhaust working capital.

To sustain on-time delivery when waves of procurement arrive in different regions, manufacturing and service need to be regionalised, building on the distributed pilots of the mid-term. Technology transfer packages, harmonised quality control, and inter-laboratory calibration underpin even quality across sites, while service centres close to customers maintain availability guarantees under tighter, evidence-linked contracts. This reduces exposure to border holds and logistics friction and positions suppliers for both temporary-permit surges and long-run frameworks that schedule works by river basin or region.

Because incumbents can integrate selective steps into their trains, avoiding channel exclusion requires becoming a specialist module provider inside larger architectures rather than insisting on a standalone posture. The measures highlight collaborative integration where end customers typically require a combination of technologies and prefer to buy from system integrators; the path forward is to strengthen selected partnerships, define role boundaries, and retain a clear performance advantage where mixture robustness and low-waste residuals remain differentiators under continuous verification and stricter waste rules.

Where the economics and certification pathways support it, recovery and traceability should be institutionalised. Making the recovery ecosystem visible to the customer, through quality-assured outputs and traceability aligned with emerging standards, links circular value to procurement scoring and reduces reliance on virgin-material markets that the CRMA flags as strategic dependencies. Over time, offtake agreements for recovered materials can stabilise a portion of revenues and strengthen the resilience argument that many industrial buyers and public agencies now recognise in policy.

By the end of this horizon, the roadmap expects a market in which serialised, evidence-rich retrofit packages are normal, regional production and service underpin delivery reliability, integration with incumbents secures access to major frameworks, and circular outputs travel on certified, traceable pathways that align with industrial resilience agendas. These outcomes preserve distinctiveness even as the field scales.

5.1 Traceability from insights to actions

For transparency, each measure in this roadmap is anchored in either a clearly stated opportunity or threat and the corresponding workshop-derived action. For example, the choice to industrialise modularity and enable phased commissioning responds directly to the opportunity created by compressed delivery windows and long lead-times, and to the threat of losing tenders when commissioning cannot be staged. The call to bundle evidence responds to a cross-cutting opportunity, approvals, finance and even revenues increasingly respond to auditable outcomes, and to the threat that liability will otherwise sit with suppliers who lack monitoring depth and lab-aligned QA. The insistence on pilot-to-operation contracts before any site work answers the pilot-to-nowhere threat that the workshop participants repeatedly flagged. The measures that hedge supply-chain instability and capacity bottlenecks translate directly from the observed seasonal surges and crisis cycles. The engineering choices around residence time, prefiltration, looping, and planned hybrids are responses to performance risks under high organics and variable influent which the scenarios and the models both anticipate. The mid-term emphasis on partnerships, on the professionalisation of performance-based service, on manufacturing scale-up without drift, on low-waste and recovery value propositions, and on collective compliance are likewise explicit within the measures corpus and map exactly onto the opportunities and threats articulated in the earlier sections. The long-term moves; serialise brownfield, regionalise manufacturing and service, integrate with incumbents, and institutionalise recovery; reflect both the trajectory of market consolidation suggested by the scenarios and the leverage-point analysis in Task 5.3.

5.2 Roles

Governance of implementation follows the same rhythm. In the first year, technology developers and integrators lead the productisation of modularity and the embedding of evidence, while utilities co-design conversion-ready pilots and host references. Regulators and public labs co-develop validation panels and documentation, and insurers and lenders publish evidence checklists that align underwriting and financing with operational realities. In the mid-term, channels consolidate through partnerships; service organisations professionalise availability guarantees; contract manufacturers commence regional assembly; and standards bodies and policymakers refine risk-management specifications for reuse and traceability for secondary materials, enabling credible offtakes. In the long run, framework tenders embed serialised retrofit packages; regional factories and service centres meet surge windows; recovery ecosystems rely on certified quality classes and CRMA-aligned targets; and industry alliances maintain living application packs that steadily reduce fragmentation drag. Each step mirrors the leverage-point emphasis of Task 5.3: concentrate effort where small, structural changes alter many decisions at once.

5.3 How the system-dynamics perspective shapes priorities

The priorities in this roadmap are not a matter of taste; they arise from the structure and behaviour of the wastewater system as revealed by the Task 5.3 system-dynamics models. Three sectoral models, selective recovery of metals, removal of contaminants of emerging concern (CoEC), and agricultural reuse, were built around stocks, flows and feedbacks that explain why adoption accelerates in some conditions and stalls in others. The models identify the high-leverage places to act, the delays that must be shortened, and the balancing loops that must be neutralised if innovation is to scale. The roadmap therefore concentrates effort where relatively small changes shift the entire system.

In the metal-recovery model, adoption is driven by relative attractiveness versus virgin materials, the installed recycling capacity, and the profitability that governs both utilisation and investment inflows. Several balancing loops keep the system in check. When utilisation rises, supply widens, the supply–demand gap increases and prices soften, reducing profit and dampening further investment; conversely, tight supply lifts prices and invites new capacity until the market cools again. The same logic operates on the virgin-materials side: high prices trigger extraction and capacity expansion, which closes the gap and brings prices down; at the extreme, elevated virgin-material supply can undercut recycled outputs despite their circular benefits. The model surfaces nine leverage points, from regulatory standards and economic incentives to investment discouragements for virgin extraction and awareness campaigns that shift demand toward circular inputs. It is not surprising, then, that this roadmap places early emphasis on quality-assured recovery modules, on traceability that converts environmental value into an accepted product, and on policy alignment with circular-materials frameworks: such measures change the variables that determine the relative attractiveness of recovered materials and they harden demand in the face of commodity cycles. They also explain the insistence on evidence as product and on partnerships with off-takers and certifiers in the middle horizon, because certification and market access reduce the amplitude of those balancing loops and allow capacity to accumulate without being wiped out by a single price swing.

In the CoEC-removal model, the key stock is the installed quaternary-treatment infrastructure, and the decisive flows are the rate at which new systems are installed and the rate at which existing systems become obsolete because of performance drift or regulatory upgrades. Here, five loops structure outcomes. On the demand side, increases in market price, often a symptom of supply lagging demand, curb the very demand that created the tightness, while improvements in observed wastewater quality reduce the perceived need for further investment, at least until new limits arrive. On the regulation side, as more sites meet the desired standard, pressure eases and the system, left to itself, tends to idle until a new impetus is introduced. Against this backdrop, the model highlights a handful of levers that move the most: regulatory requirements and fines that raise the desired quality; technological advances that improve cost or performance; knowledge of CoEC harmfulness that raises expectations among regulators and buyers; and subsidies or grants that close early economics. The roadmap's short-term focus on productised modularity, pilot-to-operation contracts, and bundled evidence therefore maps directly onto the variables that the model shows to be decisive. Modularity shortens installation delays and raises the installation inflow into the installed-infrastructure stock; conversion-ready pilots rewire the proof stock so that it becomes an on-ramp rather than a cul-de-sac; and auditable evidence tightens the link between delivered quality and regulatory or financing responses, increasing the decision-to-invest rate while reducing the variance that otherwise pushes decisions back and forth. The medium-term emphasis on service discipline further stabilises the system by reducing the discrepancy between desired and realised quality, which the model identifies as a core source of oscillation.

In the agricultural-reuse model, scarcity is both a problem and a driver. A first balancing loop ties local agriculture to water intake, natural resources and scarcity: when scarcity rises, the system seeks new sources or reduces demand; when scarcity falls, appetite for investment slackens unless other drivers are present. A second loop traces a path from scarcity to demand for regenerated water, from the capacity gap to commercial potential, and from private R&D to higher technology readiness, larger installed capacity and, ultimately, reduced dependence on natural sources. A reinforcing loop operates through political will, public R&D, advances in detection, and the production of knowledge that raises both expectations and the willingness to regulate. Once again, the roadmap's choices are simply the shortest path across these loops. Rapid, modular deployment changes the shape of the capacity-gap trajectory; evidence as product aligns political will, risk management and finance with observed performance; and collective compliance packs shorten bureaucratic lags that can otherwise eat an entire drought season. As water-reuse plans translate into hundreds of projects, the long-term push to serialise brownfield retrofits and regionalise manufacturing shortens the time constants that govern the system's response to scarcity signals, which the model indicates is the surest way to avoid oscillatory, crisis-driven adoption.

Across all three models, delays matter as much as parameter values. Procurement delays inflate the market price of solutions and trigger balancing responses; commissioning delays erode trust and shift decision-makers back toward familiar options; certification delays raise financing costs and reduce the attractiveness of service-based business models; and evidence delays allow perceived quality improvements to suppress investment prematurely. The early-horizon measures that standardise module envelopes and tie-ins, pre-agree conversion terms for pilots and instrument installations with auditable data are therefore not convenience features: they remove delay from places where the model shows delay to be decisive, thereby flattening the oscillations and allowing capacity to climb.

The models also reveal threshold effects and path dependence. In metal recovery, once traceability and quality classes are in place, offtake can scale because buyers accept recovered products without renegotiating legal status on every shipment. Below that threshold, value is trapped by liability concerns and the recovery loop fails to reinforce. In CoEC removal, once monitoring and laboratory confirmation are aligned and trusted, revenue and permitting follow delivered outcomes with fewer disputes, which allows service-based models to reinforce rather than punish suppliers. Below that threshold, the same service contracts concentrate liability and slow adoption. In agricultural reuse, crossing the technology-readiness and social-acceptance thresholds changes the political-will loop from episodic to sustained, because results are seen, understood and accepted. The roadmap's sequencing, evidence and conversion first, partnerships and service second, serialisation and regionalisation third, is designed to take the system across those thresholds in the order that the models suggest will stick.

Finally, the system view clarifies why policy alignment is a practical necessity rather than a narrative garnish. Implementation of the Water Reuse Regulation with its risk-management specifications creates a common template that reduces variance in permitting; alignment with the Drinking Water Directive's PFAS monitoring milestone changes both lender behaviour and operational routines; the UWWTD's EPR provisions transform how quaternary upgrades are financed and scheduled; and the CRMA's recycling benchmarks give recovered outputs a policy-anchored role in resilience strategies. These instruments act as exogenous shocks in the models, resetting desired quality, shifting relative attractiveness and throttling investment psychology. A roadmap that leans into those currents—by building documentation to the reuse specifications, by making PFAS and mixture evidence auditable from the start, by designing retrofit-first packages that fit EPR-funded upgrades, and by hard-wiring traceability into recovery—uses policy to push on the very levers the models identify.

Taken together, the system-dynamics perspective explains why the roadmap cares so much about modularity, evidence, conversion, service discipline, manufacturing capacity, compliance tooling, and traceable circularity. Each priority corresponds to a stock, flow or loop that determines whether adoption compounds or stalls. The sequence across horizons is the minimum-effort path through those structures, from removing delays in the first year to reinforcing stable channels in the medium term and then eliminating residual capacity bottlenecks and channel risks in the long term.

5.4 Roles

Successful implementation in this domain depends as much on who acts when as on what is technically feasible. The roadmap therefore assigns responsibilities over time to a set of actors who together control the levers identified by the models: technology developers and integrators, utilities and plant operators, regulators and public laboratories, financiers and insurers, contract manufacturers and supply-chain partners, standards bodies and off-takers, and, in the reuse case, agricultural users and public-communication actors. The sequence follows the three horizons precisely because different actors are pivotal at different stages.

In the first year, technology developers and integrators carry the main load. They must convert flexible engineering into industrialised module families with stable envelopes and interfaces, publish acceptance-test procedures that travel between sites, and design tie-ins that minimise outages in brownfield settings. They must also embed instrumentation and secure data handling so that evidence is not an afterthought but an integral feature, and they must propose pilot-to-operation contracts that define conversion metrics, data rights and commercial step-ups before any equipment is moved. Utilities and plant operators, for their part, need to convene reference sites, co-design conditional awards, and adopt operational routines that maintain sensor quality and laboratory alignment so that data can be trusted. Regulators and public laboratories should publish panel expectations and recognition pathways that allow applicants to build once and reuse many times, leaning on the Water Reuse Regulation's risk-management framework and its emerging technical specifications to create predictable documentation requirements for exposure assessment and mitigation. In parallel, financiers and insurers should publish underwriting and financing checklists that tie lending terms or performance bonds to the presence of auditable monitoring and conversion-ready pilot structures, since the Drinking Water Directive's PFAS monitoring milestone and the tightening PFAS restriction process will otherwise raise perceived risk and slow approvals. Contract manufacturers and critical suppliers must be qualified early, with small strategic inventories and regional assembly options prepared so that temporary-permit or enforcement surges do not exceed capacity. Standards bodies and off-takers have a role even now: early alignment on traceability and product-quality documentation for recovered outputs will prevent circular value from stalling on legal classification and liability. In governance terms, the first horizon is about shortening delays and de-risking conversion, and that requires decision-rights and deliverables to be explicit: technology developers own productisation and data integrity; utilities own conversion gates; regulators own template clarity; financiers and insurers own evidence-conditional terms; suppliers own readiness to deliver.

In the second horizon, responsibility shifts toward channel formation, service discipline and manufacturing scale-up. Integrators and utilities should formalise framework arrangements that allow programmatic roll-out rather than single-site tenders, with documentation packs mapped to the reuse risk-management specifications and to standardised laboratory panels for mixture-risk assessment. Technology developers must stand up remote monitoring and predictive-maintenance capability, not only to meet performance-linked contracts but to qualify for insurance underwriting that prices risk against documented response times and corrective-action histories. Insurers and lenders can reinforce this by codifying service expectations into coverage terms and financing covenants, which in turn reduces cash-flow volatility in performance-based models. Contract manufacturers should begin distributed manufacturing pilots with harmonised quality assurance to prepare for longer-term regional duplication, while supply-chain partners arrange dual-sourcing and component forecasts tied to water-stress and enforcement calendars. Regulators and public laboratories can use this period to consolidate third-party validation practices and to host template repositories for common documentation, thereby reducing administrative drag; the Commission's delegated act on risk-management technical specifications for reuse provides a useful anchor for this harmonisation process. Industry associations and consortia should lead collective compliance packs that translate successful approvals into re-usable artefacts, cutting the time spent re-arguing settled ground in new jurisdictions. Standards bodies and off-takers need to advance traceability and quality-class definitions for recovered materials so that offtake is not negotiated plant by plant, and so that circular procurement policies have something concrete to point to when they score bids. In the background, Member States will be implementing the recast UWWTD, with EPR obligations due to apply by December 31, 2028; utilities and producers can prepare by clarifying data pipelines and cost-allocation mechanisms so that quaternary upgrades are not delayed by accounting disputes.

In the long-term horizon, leadership becomes organisational rather than project-based. Integrators and utilities should embed serialised retrofit packages into framework tenders by archetype so that brownfield deployments proceed with minimal re-engineering and with installation delays reduced to logistics and tie-ins. Technology developers and contract manufacturers should convert distributed pilots into regional manufacturing and service footprints, with technology-transfer packages, inter-laboratory calibration for quality control, and digital twins that maintain process fidelity as factories are replicated. Regulators can shift their effort from one-off approvals to audits of serialised designs and monitoring systems, freeing capacity while raising the floor on performance. Standards bodies and off-takers should move recovery into normal commerce through certified traceability and quality classes aligned with the CRMA's recycling benchmarks, so that resilience considerations are priced into procurement rather than appended as narratives. Financiers and insurers will be in a position to price risk against observable, multi-year performance and to offer products that reward sustained availability and mixture-risk control, which further reinforces service-based models. Throughout, administrative harmonisation should continue: as more Member States apply the reuse regulation and the UWWTD recast, the expectation should be that documentation packs compiled earlier in the decade transfer across borders with only marginal adaptations.

Two further governance considerations complete the picture. First, data stewardship must be settled early and renewed at each horizon. Evidence is the backbone of approvals, finance and service contracts, so data rights, retention, access and cyber-security need named owners, with clear pathways for regulators and lenders to receive and trust the data they depend on. Second, decision gates should mirror the roadmap's thresholds. There should be a gate that prevents pilots from starting without conversion terms; a gate that prevents serialisation without two or more validated reference kits per archetype; and a gate that prevents regional duplication without QC and calibration proven across at least two manufacturing sites. These gates institutionalise the very thresholds that the models suggest are needed for adoption to compound rather than oscillate.

The result of this allocation of responsibilities is a living implementation architecture. In the first year, the centre of action sits with those who can shorten delays and de-risk conversion; in the middle years, it shifts to those who can build channels, stand behind performance, and scale production without quality drift; and in the long run, it rests with those who can serialise, certify and regionalise so that supply keeps pace with policy-driven demand without periodic breakdowns. At each stage, the actors are acting not in the abstract but on the variables that the system-dynamics work shows to be decisive.

5.5 Alignment with policy and external enablers

The actions in every horizon are designed to match the trajectory of EU policy rather than fight it. The reuse measures and documentation pack concepts mirror the Water Reuse Regulation's requirements for risk management and transparency and the Commission's ongoing technical specifications work. The emphasis on continuous, auditable evidence matches the Drinking Water Directive's PFAS monitoring milestone and the way lenders and insurers already price risk. The consolidated treatment and residuals strategies anticipate REACH/POPs and a 2026 horizon for committee opinions on the broad PFAS restriction, reducing surprises in projects with long service lifetimes. The retrofit-first, brownfield stance aligns with the UWWTD EPR cost-allocation model by making quaternary upgrades deliverable within existing plants, while the circular themes connect directly to the CRMA and to emerging traceability guidance for secondary materials that off-takers and auditors can accept. Policy is therefore not an externality in this roadmap; it is the structure within which the roadmap becomes feasible and financeable.

5.6 What changes between horizons

In the narrative above, the horizons are not arbitrary. The first year concentrates on proving, packaging, and protecting because those steps alter the system's feedbacks fastest: modularity reduces delivery delays; evidence converts uncertainty into momentum; pilot-to-operation contracts turn demonstrations into assets; supply hedges reduce oscillations; and boundary-condition engineering prevents performance dips when monitoring tightens. The second horizon focuses on scaling, certifying, and servicing because channels and contracts dominate how rapidly installations accumulate; this is when partnerships, documentation, and service discipline move projects from intermittent wins to predictable programmes. The third horizon serialises and regionalises because platforms and frameworks now shape access; success requires duplication with quality and integration with incumbents, while circular value is locked in with traceability and offtakes rather than ad-hoc claims. At each step, the measures are the smallest changes that move the largest stocks and shift the most consequential loops.

5.7 Closing reflection

The roadmap's centre of gravity is pragmatic. Build what can be deployed quickly, prove it continuously, and structure contracts so that proof converts to revenue. Do this while hedging delivery risks and preparing for scale through partnerships and distributed capacity. In that configuration, mixture-robust performance, low-waste residuals, and credible recovery are not optional flourishes; they are the distinctive edge that remains when the market is mature and the field is crowded. Each action presented here is traceable to the workshop corpus, is positioned within the EU policy architecture that governs procurement and compliance, and is defended by the system-dynamics logic that explains why it will work in practice.

6.0 Iterative follow-up and update mechanism

The roadmap developed in this project should not be treated as a static plan but as a living strategic instrument that is continually refined to reflect new evidence, market signals, policy developments and technological progress. The academic literature on roadmapping consistently emphasises that robust roadmaps function as dynamic coordination tools, updated through cycles of sensing, sense-making and strategic adjustment rather than being executed linearly from beginning to end (Phaal & Muller, 2009). This iterative orientation is particularly important in fields such as emerging wastewater treatment and chemical-pollution mitigation, where regulatory timelines, technology readiness levels and adoption rates shift due to external forces that could not have been fully anticipated during the initial roadmap creation (Kostoff & Schaller, 2001).

6.0 Roadmap as an iterative process

Treating the roadmap as an iterative process ensures that organisations across the wastewater value chain—technology developers, utilities, regulators and integrators—can adapt their strategic priorities as new risks emerge or policy levers become activated. The literature stresses that effective roadmaps are refreshed on a regular cadence, often annually or semi-annually, with interim "micro-updates" triggered when early warning signals indicate that core assumptions are shifting (Willyard & McClees, 1987). In this project, iteration is critical because several of the roadmap's priority domains, such as regulatory tightening on PFAS and CoECs, infrastructure cycles for quaternary treatment, and the scaling conditions for modular systems, are governed by non-linear and path-dependent dynamics, which necessitate revisiting not only actions but also the causal assumptions behind them (Rohrbeck & Schwarz, 2013).

An iterative roadmap also supports organisational learning, a central function highlighted by scholars of strategic foresight. Each update cycle becomes an opportunity to evaluate which measures have produced measurable progress, which barriers proved more persistent than expected, and where unexpected accelerators have emerged, thereby improving strategic agility.

6.1 Monitoring signals and decision points

To operationalise iteration, the project ecosystem must establish a structured approach to monitoring signals and defining decision points that trigger roadmap updates. Foresight research identifies three categories of signals that typically require close monitoring: "hard signals" such as regulatory enactments, technology-performance data, or procurement-volume changes; "semi-soft signals" such as shifts in investment flows, early market-adoption patterns or coalition-building among incumbent integrators; and "soft signals" such as media narratives, emerging scientific concerns or stakeholder expectations (Ansoff, 1975).

For iMERMAID's context, monitoring should focus on signals that directly influence the leverage points identified in the system-dynamics analysis. These include changes in regulatory pressure on CoECs, incentives and penalties associated with treated effluent quality, funding availability for agricultural reuse infrastructure, price movements in virgin metals affecting the economics of selective recovery, and performance data from modular installations operating under variable conditions. System-dynamics literature emphasises that decision points should be tied to threshold effects, moments where incremental change in a variable produces disproportionate system-level impacts, making it essential not only to detect signals but to understand where each sit within the system's feedback architecture (Sterman, 2000).

In practice, the roadmap should therefore define a set of trigger conditions, such as regulatory tightening, breakthrough improvements in monitoring or treatment efficiency, supply-chain instability exceeding a tolerance range, or accelerating adoption patterns that justify serialising manufacturing or expanding regional footprints. These triggers provide objectivity to update decisions and ensure that roadmap revision cycles are not driven purely by intuition or internal opinion but by observed system behaviour (Heger & Rohrbeck, 2012).

6.2 Governance for updates

A clear governance structure is required to ensure that the iterative follow-up mechanism functions consistently and that roadmap updates are based on shared evidence, not fragmented local assessments. Roadmapping research emphasises the importance of a multi-stakeholder governance model where roles are distributed across organisations responsible for technology development, regulatory oversight, operations, and value-chain integration (Phaal, Farrukh & Probert, 2004). This ensures that updates reflect the full set of technological, economic, regulatory and social factors that shape adoption dynamics.

In the context of this project, governance should be anchored in a joint steering function, for example, a cross-partner Roadmap Coordination Group, that convenes periodically to review monitored signals, assess deviations from planned trajectories, and determine whether adjustments are required. Academic studies note that update governance is most effective when responsibility is distributed between:

- an evidence owner, responsible for gathering and validating performance, policy, and system-dynamics data;
- a strategic integrator, responsible for interpreting signals and aligning changes with programme-level objectives; and
- a decision authority, responsible for approving roadmap modifications and ensuring alignment with long-term innovation aims (Kappel, 2001).

Such governance also benefits from embedding feedback loops, whereby insights from each roadmap cycle not only inform planning but also refine the monitoring process itself, mirroring the double-loop learning principles widely recognised in strategic-management literature (Argyris & Schön, 1978). Over time, this governance model strengthens the resilience of the roadmap by making it responsive to system complexity, policy timing, and technological maturation patterns central to iMERMAID's mission.

6.2.1 Roles within the iMERMAID project and innovation cluster

Within the iMERMAID cluster, technology developers, solution providers, and research partners serve as the primary evidence contributors to the update process. Their responsibility is to provide validated insights on technology readiness, operational performance, robustness under variable conditions, and scalability constraints emerging from pilots, demonstrations, and early deployments. These inputs ensure that roadmap updates remain grounded in empirical evidence rather than expectations or advocacy positions.

Utilities, operators, and end-user stakeholders linked to the cluster play a complementary role as practice validators. Through feedback on procurement processes, permitting timelines, operational reliability, cost structures, and organisational capacity, they help test whether roadmap priorities remain realistic under real-world constraints. Their experience is especially important for identifying emerging bottlenecks related to retrofit feasibility, service models, uptime obligations, or administrative burden that may not be visible from a purely technological or policy perspective.

Research organisations and system-dynamics experts retain a distinct function as system interpreters. Building on the modelling work conducted in T 5.3, these actors support the governance process by assessing whether new signals represent incremental noise or structural changes in system behaviour, for example, by altering feedback loops related to regulatory pressure, investment incentives, or adoption timing. This function is critical to ensuring that roadmap updates address root causes rather than symptoms.

6.2.2 Role of EU institutions and policy actors

EU institutions play a structural and ongoing role in shaping the relevance, timing, and feasibility of roadmap priorities. Their influence operates both indirectly, through regulation and funding, and directly, through signalling and guidance that inform decision-making across the ecosystem.

The European Commission, particularly through Directorates-General responsible for environment, water policy, chemicals, and research and innovation, acts as a central policy signal generator. Regulatory proposals, delegated acts, guidance documents, and funding instruments constitute key inputs to the roadmap update process. Changes in regulatory scope, monitoring obligations, risk-management requirements, or financing logic may therefore trigger targeted roadmap revisions, especially where they affect compliance thresholds, cost allocation, or technology qualification.

EU regulatory agencies and advisory bodies contribute as standard-setting and risk-assessment references. Their evolving positions on topics such as contaminant mixtures, PFAS management, water reuse risk assessment, or the classification of recovered materials as products or waste directly influence the credibility and bankability of roadmap measures. While these bodies do not participate directly in roadmap governance, their outputs provide authoritative reference points that must be considered in any substantive update.

The EU research and funding ecosystem plays a further role as a selective accelerator. Changes in programme priorities, eligibility criteria, evaluation emphasis, or co-funding expectations can materially affect which roadmap actions are feasible or attractive. Governance for updates must therefore assess not only funding availability, but also whether the logic of funding aligns with the strategic intent of the roadmap.

6.2.3 Decision authority, cadence, and transparency

Within this multi-actor structure, the Roadmap Coordination Group holds the decision authority to formalise updates. Decisions should be based on structured assessments that link monitored signals to the roadmap's core assumptions, leverage points, and sequencing logic. Minor adjustments, such as re-sequencing actions within a time window or refining emphasis, can be addressed through lightweight updates, while substantial changes, such as accelerating or deprioritising entire strategic paths, should trigger a documented revision cycle.

The governance framework benefits from a dual update cadence. Periodic reviews aligned with reporting or programme milestones provide continuity and learning over time, while ad-hoc reviews are triggered by clearly defined external events, such as regulatory enactments, major technological breakthroughs, or abrupt shifts in market adoption or financing conditions. Importantly, each update cycle should also refine the monitoring framework itself, improving the quality and relevance of signals used in subsequent cycles.

Overall, this governance model ensures that responsibility for roadmap evolution is distributed yet coherent. It anchors updates in evidence, system understanding, and policy reality, while preserving the roadmap's role as a shared strategic reference rather than a prescriptive implementation plan. In doing so, it supports iMERMAID's long-term objective: enabling the sustained, system-level uptake of innovative monitoring and remediation solutions under evolving European regulatory and market conditions.

7.0 Conclusions and recommendations

7.0 General conclusions

7.0.1 Market pull is real, but it is uneven and evidence-centred.

Across the full scenario set and the consolidated workshop outputs, the most durable opportunities arise where buyers can deploy modular upgrades quickly and where approvals, finance and revenues respond to continuous, auditable performance rather than claims. This pattern recurs in municipal reuse, CoEC removal, and selective recovery, and it is reinforced by EU-level instruments that make monitoring, risk management and transparency binding features of compliance and procurement. In practical terms, the centre of gravity has shifted from technology promise to verified outcomes, which is why instrumented, certifiable modules and conversion-ready pilots repeatedly outperformed bespoke propositions in the stakeholder material.

7.0.2 System behaviour, not point decisions, determines scalability

The Task 5.3 system-dynamics work shows adoption is constrained or enabled by a small set of stocks, flows and feedbacks: installed capacity inflows (and the delays that govern them), discrepancy between desired and realised quality (and the evidence that resolves it), and path-dependent thresholds in certification, residuals handling and offtake trust. Roadmap elements that reduce installation delay, wire pilots to installation, and institutionalise evidence creation change the shape of adoption trajectories more than isolated, one-off investments.

7.0.3 Retrofit-first routes dominate the next wave; serialisation follows

In the operating conditions described by partners, tight CAPEX, programme-based tenders, seasonal and crisis-driven surges, the fastest path to volume is brownfield and modular. That path leads naturally to serialised, archetype-specific kits when projects scale. The workshop corpus and scenarios converge on this sequence, and the EU policy environment (reuse risk management, UWWTD EPR implementation, and PFAS monitoring) increases the value of solutions that can be dropped into existing trains with minimal tie-ins and maximum evidentiary clarity.

7.0.4 Circular value is credible where traceability and quality classes exist; otherwise, it stalls.

Selective recovery creates resilience and revenue only when recovered outputs travel on quality-assured, certifiable pathways. The CRMA's recycling benchmarks and emerging traceability guidance give a policy anchor, but buyers still need assurance on product status and liability. Without that, recovery remains an engineering success that fails to compound commercially.

7.0.5 Fragmentation and administrative drag remain material risks, but are tractable

Divergent documentation and inspection practices create rework and delay, especially across borders. Yet the Water Reuse Regulation's technical specifications for risk management, together with shared templates and collective compliance tooling, can reduce the variance that undermines schedules and cash flow.

7.1 Strategic recommendations

7.1.1 Make "evidence-as-a-product" the default, not an add-on

Build and ship treatment modules with embedded monitoring, lab-alignment protocols and secure data pipelines so that approvals, finance and service models can price against verified outcomes. This directly targets the "discrepancy between desired and realised quality" loop in the CoEC model and aligns with the Drinking Water Directive's PFAS monitoring obligations now in force. It also reduces diligence friction under reuse risk-management and accelerates lender and insurer acceptance.

7.1.2 Convert pilots into assets by contracting the step-up before any deployment

Adopt a standard pilot-to-operation template that defines success metrics, data rights, third-party validation and commercial step-ups contingent on evidence. This rewires "proof" from a sink that consumes resources to a feeder flow into "installed capacity", which the system-dynamics work identifies as the biggest compounding lever in early markets. The approach is consistent with programme-based public procurement and the framework-tender logic observed in the scenario material.

7.1.3 Industrialise modularity and phase-in designs to shorten installation delays

Stabilise envelopes and tie-ins, define sub-modules for different duty points, and publish acceptance-test sequences that travel with each archetype. In the models, reducing installation delay has an outsized effect on the installed-capacity trajectory and dampens the price–demand oscillations that otherwise slow adoption. In policy terms, this is the practical route to meet temporary-permit windows and EPR-funded upgrade schedules.

7.1.4 Professionalise performance-based service before scaling commercial exposure

Outcome-linked contracts must sit on 24/7 monitoring, predictive maintenance and insurer-aligned documentation; otherwise, they concentrate liability and erode margin. As continuous verification becomes normal practice, service discipline becomes both a competitive advantage and a precondition for bankability.

7.1.5 Regionalise manufacturing and service footprints to meet surge windows reliably.

Pilot distributed manufacturing with harmonised QA in the medium term and commit to regional duplication where surge demand and logistics risks justify it. This mitigates the supply-chain instability threats repeatedly surfaced in the workshops and seen in the stop–start procurement patterns of reuse and enforcement cycles.

7.1.6 Institutionalise recovery with traceability; engage early on product status.

Where recovery economics pencil out, couple selective modules with chain-of-custody documentation and quality classes aligned to CRMA targets and ISO guidance. Work with off-takers and legal counsel to secure stable by-product or product status and to clarify liability. This moves circular value from a marketing claim to a bankable component of the offer.

7.1.7 Reduce administrative drag through collective compliance and template reuse.

Form cross-partner "compliance cells" that maintain living application packs mapped to reuse risk-management specifications and standard laboratory panels for mixture risk. Regulators and public labs can amplify the effect by publishing reference templates and validation expectations. This reduces variance in approval timelines across Member States and speeds inflows into installed capacity. [eur-lex.europa.eu]

7.1.8 Align sequencing with leverage points: prove and package first; scale and certify second; serialise and regionalise third

Follow the roadmap's horizon logic because it mirrors the order in which thresholds are crossed in the models. Evidence and conversion change decisions immediately; partnerships, service and manufacturing scale create programme capacity; serialisation and regional duplication remove residual bottlenecks and defend distinctiveness as incumbents integrate selective steps.

7.2 Implications for iMERMAID and its stakeholders

For technology developers and integrators, the immediate task is to finish productising modularity and to hard-wire evidence creation so that pilots convert and service models are bankable. This is the most efficient way to move the system's adoption curve with the resources available, and it is fully compatible with the reuse, PFAS and UWWTD policy architecture.

For utilities and operators, the highest leverage comes from adopting conditional awards tied to independently validated outcomes, enforcing lab–sensor alignment routines and prioritising retrofit-friendly packages that minimise disruption and commissioning delay. These changes shorten cycles and increase the reliability of both compliance and budgets.

For regulators and public laboratories, publishing risk-management templates and validation panels, and recognising serialised archetypes for reuse and quaternary upgrades, will reduce rework without relaxing standards. This is the fastest path to more predictable schedules in multi-site programmes.

For financiers and insurers, aligning underwriting and financing covenants to the presence of auditable monitoring, conversion-ready pilots and insurer-grade service documentation will reward the behaviours that the models show to be decisive levers for scalable adoption, while reducing avoidable exposure in outcome-linked contracts.

For circular-economy stakeholders and off-takers, early participation in traceability, quality-class definition and contract templates will unlock recovery niches that otherwise stall on classification and liability, allowing circular value to contribute to resilience objectives anchored in the CRMA framework.

As a closing statement, evidence gathered through scenarios, workshops and system-dynamics modelling points to a clear, practicable thesis: scalable success in advanced wastewater treatment and circular recovery will come from solutions that can be deployed fast, proven continuously and scaled programmatically through partnerships, service discipline and serialised delivery. The policy environment—reuse risk management, PFAS monitoring and restriction, UWWTD EPR, and CRMA-aligned circularity—does not merely permit this thesis; it favours it. The recommendations above prioritise the smallest, most controllable interventions that move the largest stocks and reshape the most consequential feedbacks in the system, translating iMERMAID's insights into an implementable agenda for the next three to five years.

8.0 References

Ansoff, H. (1975). Managing strategic surprise by response to weak signals. California Management Review, 18(2), 21-33.

Argyris, C., & Schön, D. (1978). Organizational Learning: A Theory of Action Perspective. Addison-Wesley.

Cheng, M.N., Wong, J.W.K., Cheung, C.F. & Leung, K.H. (2016). A scenario-based roadmapping method for strategic planning and forecasting: A case study in a testing, inspection and certification company. Technological Forecasting and Social Change, 111, 44-62.

Heger, T., & Rohrbeck, R. (2012). Strategic foresight for collaborative exploration of new business fields. Technological Forecasting & Social Change, 79(5), 819-831.

Hussain, M., Tapinos, E., & Knight, L. (2017). Scenario-driven roadmapping for technology foresight, Technological Forecasting & Social Change, 124, 160–177.

Ilevbare, I.M., Probert, D. & Phaal, R. (2014). Towards risk-aware roadmapping: Influencing factors and practical measures, Technovation, 34(8), 399-409.

Kappel, T. (2001). Perspectives on roadmaps: How organizations talk about the future. Journal of Product Innovation Management, 18(1), 39-50.

Kostoff, R., & Schaller, R. (2001). Science and technology roadmaps. IEEE Transactions on Engineering Management, 48(2), 132 - 143

Lindgren, M. & Bandhold, H. (2009). Scenario Planning - Revised and Updated. The Link Between Future and Strategy, Ed. 1, Palgrave Macmillan London.

OECD (2025), Strategic Foresight Toolkit for Resilient Public Policy: A Comprehensive Foresight Methodology to Support Sustainable and Future-Ready Public Policy, OECD Publishing, Paris.

Phaal, R., Farrukh, C. J. P., & Probert, D. R. (2004). Technology roadmapping—A planning framework for evolution and revolution, Technological Forecasting & Social Change, 71(1), 5–26.

Phaal, R., & Müller, G. (2009). An architectural framework for roadmapping: towards visual strategy. Technological Forecasting & Social Change, 76(1), 39-49.

Probert, D.R., Farrukh, C.J.P. & Phaal R. (2003). Technology roadmapping—developing a practical approach for linking resources to strategic goals. Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture, 217(9),1183-1195.

Rohrbeck, R., & Schwarz, J. (2013). The value contribution of strategic foresight: Insights from an empirical study of European firms. Technological Forecasting & Social Change, 80(8), 1593-1606.

Samson-Onuorah, Caroline & Bakare, Felix. (2025). Dynamic Strategic Foresight Using Predictive Business Analytics: Modeling Competitive Advantage Under Volatile Market and Technological Conditions. International Journal of Research Publication and Reviews. 6. 1044-1060.

Saritas, O., & Aylen, J. (2010). Using scenarios for roadmapping: The case of clean production, Technological Forecasting & Social Change, 77(7), 1061–1075.

Siebelink, R., Halman, J.I.M. & Hofman, E. (2016): Scenario-Driven Roadmapping to cope with uncertainty: Its application in the construction industry. Technological Forecasting & Social Change, 110, 226–238.

Sterman, J. (2000). Business Dynamics: Systems Thinking and Modeling for a Complex World. McGraw-Hill.

Willyard, C., & McClees, C. (1987). Motorola's technology roadmap process. Research Management, 30, 13-19.

Zumbrunn, L. (2023). Resilience through foresight: implications for the public sector. Yearbook of Swiss Administrative Sciences, 14(1).

The Mediterranean Sea and its surrounding regions support a diverse variety of essential socioeconomic activities. It is one of the highly exploited water ways and the influence of anthropogenic activities on its marine habitats and ecosystems has grown significantly since the industrial revolution. Because of this, the Mediterranean Sea basin is very vulnerable to chemical contamination and build-up. To safeguard the Mediterranean Sea basin from contaminants for emerging concerns (CoEC), iMERMAID will integrate, coordinate, and synergize innovative preventive, monitoring, and remediation solutions. iMERMAID will build an evidence-based multidimensional framework that will guide policymaking and transform societal perceptions to reduce CoEC usage, emissions, and pollution. Furthermore, next generation sensor and remediation solutions will be developed within iMERMAID to monitor and remove prioritized chemicals from its source while reducing upstream pollution. iMERMAID builds an ideal interdisciplinary team by bringing together prominent SMEs, researchers, regulators, and innovation professionals who have been essential in improving the knowledge and awareness of CoEC. Beyond state-of-the-art techniques, iMERMAID will strive to strengthen regulations against CoEC, expand economic possibilities and competitiveness, improve the standard of living for EU residents, while preventing the accumulation of chemical pollution in the Mediterranean Sea basin. iMERMAID will empower the efforts to create a zero pollution, contaminant free waters by enabling the Chemical Strategy's goals to become a practical reality.

iMERMAID

Funded by
the European Union