Robert Atkinson

THE HYDROGEN AGE

Costs, Constraints, and the Real Shape of the Energy Transition

Foreword by Silke Frank
President, German Hydrogen Association (DWV)



Carl Grossmann
Verlag

Praise for *The Hydrogen Age*

"This exceptionally well-structured book provides a comprehensive and accessible overview of the hydrogen sector, offering a pragmatic, balanced, and candid perspective on its opportunities and challenges. It successfully bridges the gap between the technical realities of hydrogen science and technologies and the broader political, industrial, economic, and societal dimensions of the energy transition. Importantly, it goes beyond technology to highlight the people, institutions, policies, and strategic decisions that will ultimately shape hydrogen's role in the global energy landscape of the future."

— Prof. Dr. Bruno G. Pollet
Executive Vice-President of the International Association for Hydrogen Energy (IAHE)
President of the IAHE Clean Hydrogen Division

"The Hydrogen Age offers a thorough presentation of the relevant factors for hydrogen in 2026. Robert Atkinson brings to any reader, new to the industry or seasoned participant, a strong understanding and realistic view of the benefits, opportunities, and impediments for inserting hydrogen and hydrogen technologies in the broader global energy mix."

— Frank Wolak
President and CEO
Fuel Cell and Hydrogen Energy Association, USA

"'The Hydrogen Age' by Robert Atkinson is a comprehensive compilation of the diverse, competing issues – ranging from science and technology to economics – that shape hydrogen applications across the entire value chain. The book inspires to further deep dive into the various topics to enhance the understanding of them in greater detail. Essential reading for anyone in energy policy or investment."

— Roland Schulze
Managerial Adviser / Conseiller de Direction, Low Carbon Energy Technology
Energy Department, European Investment Bank (EIB)

"Hydrogen has never suffered from a lack of promises. What has been missing is an unsentimental reality-check of where it really makes sense. Robert Atkinson provides one. The Hydrogen Age is technically literate, economically hard-headed and above all refreshingly free of evangelism. Its strength lies in separating what is possible from what is useful — and in showing that hydrogen will matter most when we no longer need to keep explaining why. This book was needed."

— Jürgen Pfeiffer
Hydrogen commentator and moderator
Host of 'The Hydrogen Elevator' Podcast

"Timely, honest and forward looking, this book gives Hydrogen the reckoning it needs, balancing the fundamentals with a clear eye view of the future. Highly recommended whatever your background"

— Kerry-Ann Adamson
Head of Hydrogen consulting, Capgemini

"An expansive and impressively detailed exploration of the hydrogen sector, taking a genuinely whole-system view from production technology and use-case to policy, finance and wider societal considerations. I learned a great deal and the depth is matched by style, flair (and humour!). Highly recommended for early-career and established professionals alike, particularly anyone looking to broaden their understanding beyond their own specialism."

— Dr Paula McNamee

"The hydrogen ecosystem involves a wide range of technologies, markets, applications, and challenges, oftentimes making it difficult to see the bigger picture. Working in the hydrogen industry myself, I truly appreciate the way 'The Hydrogen Age' provides a well-structured overview to such a complex and rapidly evolving field.

The book provides a comprehensive picture on everything hydrogen in an easy-to-understand, engaging, and entertaining way. This book does an excellent job of connecting the dots and making the world of hydrogen more accessible to both newcomers and experienced professionals.

Thank you, Rob, for the time, effort, and passion you have invested in making this complex industry easier to understand. This is a valuable contribution to helping more people navigate, understand, and ultimately contribute to the development of the hydrogen economy."

— Joachim Ladra
Head of Sales, Marketing & Communication
cellcentric

"The Hydrogen Age offers a timely, clear-eyed and highly relevant perspective on hydrogen's role in the energy transition. Robert Atkinson succeeds in combining technological insight with an honest assessment of costs, constraints and system-level challenges. This balanced approach makes the book a valuable contribution for decision-makers, industry professionals, researchers, students and anyone seeking to understand where hydrogen can deliver real impact—and where other solutions may be more appropriate."

— Univ.-Prof. Dr.-Ing. Thomas von Unwerth
Professor of Advanced Powertrains, Chemnitz University of Technology
Chairman, HZwo e. V.

Robert Atkinson

The Hydrogen Age

The Hydrogen Age

Costs, Constraints, and the Real Shape of the Energy Transition

Insights from the Front Lines of the Hydrogen Economy

by

Robert Atkinson

With a foreword by Silke Frank,

President of the German Hydrogen Association (DWV)

First edition 2026

Carl Grossmann Verlag

Author:

Robert Atkinson

Bibliographic information published by the Deutsche Nationalbibliothek (National Library, Germany): The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available on the Internet at http://dnb.d-nb.de.

© 2026 Robert Atkinson

Licence: CC BY-SA 4.0



This work is published open access under a Creative Commons licence. It may be shared and adapted, provided that the author is credited and any adapted work is distributed under the same licence as the original.

Further details are available at:
https://creativecommons.org/licenses/by-sa/4.0/



DOI: 10.24921/2026.94115984

This work has been prepared with great care. Nevertheless, the author and the publisher accept no liability for the accuracy of information or advice, or for any printing errors. The views expressed in this book are solely those of the author and do not represent the views of any employer, past or present, or of any organisation with which the author is affiliated.

The typeface Lora is licensed under the SIL Open Font License, Version 1.1.

Printed in Germany and Poland on acid-free paper.
Production of the publisher's edition (print and digital):

Carl Grossmann Verlag, Berlin
www.carlgrossmann.com

ISBN (print edition, hardcover): 978-3-941159-83-9
ISBN (e-book, open access): 978-3-941159-84-6

Suggested citation:
Atkinson, Robert: *The Hydrogen Age. Costs, Constraints, and the Real Shape of the Energy Transition*. Berlin: Carl Grossmann Verlag, 2026.
DOI: 10.24921/2026.94115984

Foreword

Silke Frank, President of the German Hydrogen Association

For many years, Robert Atkinson and I have been connected by far more than our professional discussions about hydrogen. Through countless conversations, we have developed a relationship founded on mutual respect and appreciation.

A foreword should give readers an additional perspective on a book and explain why it is worth reading. It should honour both the work and its author. Meeting this expectation was the real challenge for me when Robert Atkinson — Rob to me — asked me to write this foreword.

Our conversations have always combined technical depth with a shared conviction: hydrogen plays an essential role in energy supply, European competitiveness, and global resilience - wherever it creates tangible value within an integrated energy system.

This is precisely the perspective that shapes The Hydrogen Age.

Robert writes from an especially valuable standpoint: practical experience. His career of more than two decades has given him a deep understanding of complex energy systems, their technical possibilities, and the commercial conditions and human factors required to turn technology into a viable business model.

This book invites you to understand the wider context, challenge assumptions, and bring different perspectives into conversation. Rob combines his distinctive optimism with realism, and his enthusiasm with a willingness to address uncomfortable constraints openly.

The subtitle captures this approach perfectly: 'Costs, Constraints, and the Real Shape of the Energy Transition' are inseparably linked. A systemic perspective is essential to understanding the specific contribution hydrogen can make and how it can develop into a strong and prosperous hydrogen economy.

I was particularly moved by the very personal starting point of this book. A conversation between Robert and his son, together with his son's suggestion that he

should capture these many complex interconnections in a book, set in motion the process that became The Hydrogen Age.

This motivation also connects us, as we are both parents. For us, responsibility towards the generations that follow is self-evident. Every day, we see clearly who we are making decisions for and who will live with their consequences far longer than we will.

What kind of world will we leave to our children? At the end of our lives, will we be able to tell them with credibility that we used our knowledge, our opportunities, and our influence to the best of our ability?

My longstanding commitment to hydrogen and to the continuous development of our energy system is therefore, for me personally, far more than a professional responsibility or a matter of technology policy. At its core lies a deeply held ethical responsibility: to protect creation and preserve the foundations of life that we pass on to the generations that follow us.

We are guests on this earth for a limited time. This understanding represents an ethical principle for me, but above all a universal responsibility that applies to all of us: we must treat this planet in a way that allows the generations after us to live well on it.

One of my intellectual mentors, Dr. Johannes Töpler, captures this responsibility in his lectures with a simple yet fundamental insight: the laws of nature and the rules created by human beings are not automatically aligned. Nature does not follow our economic expectations. It is our responsibility to shape our actions within its boundaries.

This principle is deeply rooted in the way I think and act. When intelligently connected, ethical responsibility and economic interests reinforce each other. The tension between them can become a powerful force for better technologies, viable business models, and responsible social structures.

Earth Overshoot Day shows just how urgent this task is. It marks the estimated date on which humanity's demand for ecological resources and services exceeds what earth can regenerate in that same year. In 2026, this day was especially early, falling on July 30th. In 2000, it was still in mid-September. The long-term trend is clear: for more than five decades, humanity has consistently consumed ecological resources faster than our planet can regenerate them.

The German Hydrogen Association (DWV) also pursues a clear social and non-profit mission. Its statutes place the protection of nature and the environment, the conservation of resources, and the preservation of our natural foundations of life at the heart of the association's work. At the same time, they connect these goals with economic performance, competitiveness, and security of supply.

The association's mission goes far beyond promoting hydrogen as an isolated technology. With its unique focus on the entire hydrogen value chain, the DWV helps shape the development of a sustainable hydrogen industry from production and storage through transport infrastructure to applications. It contributes to shaping legislative and regulatory frameworks and promotes interdisciplinary exchange.

Hydrogen creates options for energy storage and transport, connects renewable energy generation with industrial applications, and provides solutions where direct electrification reaches its limits. In this way, it makes a decisive contribution to a resilient energy supply that also respects our planet's natural boundaries.

In this book, Robert Atkinson examines efficiency losses and costs alongside infrastructure requirements, financing, regulation, and geopolitical dependencies. This realistic assessment is what makes The Hydrogen Age so valuable. The book moves confidently between electrons and molecules, technical ambition and market reality, investment requirements and commercial viability.

The author explicitly invites his readers to disagree with him, develop his ideas further, and contribute their own perspectives. This openness runs throughout the book and reflects the Robert I have known and valued for many years.

Rob, with this book, you provide a strong foundation for an honest and necessary discussion.

I wish this book a wide and attentive readership, the intensive discussions it deserves, and above all tangible impact. Read it with curiosity! Challenge its arguments! Bring your own experience into the debate.

The decisive question is where hydrogen creates clearly defined and systemically indispensable value.

I thank Robert Atkinson for a book that reveals the real conditions required for progress and enriches the discussion on building a strong and competitive hydrogen economy.

Dear readers, I wish you an insightful read, fresh perspectives, and the courage to turn knowledge into action.

Silke Frank
President of the German Hydrogen Association (DWV)

For Fynn and Christina

& in memory of
Regis Fenollar
A true teammate and friend
1978-2025

Note on this Edition

This is the first edition of The Hydrogen Age.

The book draws upon the author's professional experience across the energy and industrial sectors, together with publicly available information and research available at the time of writing. It aims to bridge technical, commercial, and policy perspectives to explore the role of hydrogen within the wider energy transition.

As the hydrogen economy continues to evolve rapidly, some of the policies, projects, and market conditions described in these pages will inevitably change over time. The underlying principles, however, remain centred on understanding hydrogen as part of a broader, integrated energy system.

Constructive feedback from readers is always appreciated and will help inform future editions of this work.

Contents

Preface: Why This Book, Why Now 1
1. Understanding Energy 7
2. Hydrogen — The System Adhesive 19
3. Hydrogen Use Cases 29
4. Technology Pathways and Innovation 51
5. Operating Reality and Industrial Readiness 81
6. The Role of Nuclear (and also Oil & Gas) 95
7. Integrating the System — Hydrogen and the Wider Energy Mix 119
8. Understanding LCOH (Levelized Cost of Hydrogen) 133
9. Financing the Hydrogen Economy 159
10. The Current Challenges in Delivering This Future 179
11. The Social and Environmental Dimensions 189
12. The Geopolitics within the Energy Transition 211
13. Energy Today vs. Energy Tomorrow – Changing Infrastructure 263
14. What Business Can Do 311
15. Digitalisation and AI — Accelerating the Hydrogen Transition 321
16. Action Today 339
17. Hydrogen in 2050: Three Possible Futures 353
18. Conclusions 365
Afterword 391
Thanks 395
About the Author 397
Illustration Sources and Rights Information 399

Preface: Why This Book, Why Now

There are advantages and disadvantages to having teenagers roaming around the house. For me, one of those moments became one of the reasons for this book.

My son, now at the stage where career discussions have started at school, asked what I actually do — and whether it might be something for him to consider in the future. My answer was an enthusiastic yes. The energy transition is one of the defining challenges of our time, and it will take decades of commitment and collaboration to achieve.

After a long conversation full of explanations and a few sketches on scrap paper, he said, "There's too much for me to remember in one go. You'll have to write a book if you expect me to keep up."

That offhand remark stayed with me. After more than twenty years in industry — mostly in energy — perhaps I did have a broad enough perspective to explain this story to others, not just to my son. Additionally, there was perhaps value in putting it all down — not as a technical manual or policy paper, but as something readable and grounded.

A Life in Energy

My career began with a degree in electronics and communication engineering, followed by work commissioning large variable-frequency drives and automation control systems — the hidden machinery that keeps pipelines, drilling rigs, and LNG plants running. Those early years gave me an appreciation for the scale and complexity of the systems that power modern life.

From there, I moved through several corners of the energy world. I spent time in wind power before returning to oil and gas, developing service strategies for my company's deliveries to major LNG projects, such as Gorgon and Ras Laffan—specifically around medium-voltage variable frequency drive (MV VFD) systems. Later, I supported the integration of an acquired compression business into the wider organisation and spent approximately six years supporting my company's

Oil & Gas market development activities. Working as part of the team responsible for translating strategic direction into practical market execution provided a broader perspective on how technology, business, policy, and industrial priorities intersect within the energy sector. About seven years ago, my focus pivoted again back to automation, and this time toward hydrogen and the wider energy transition.

Today, I help shape my company's strategy in this field and for the last three years have served on the steering committee driving our hydrogen business forward. That vantage point — where traditional energy and emerging technology meet — has given me a front-row seat to how industries evolve under pressure.

I've always wanted to understand how things work. That curiosity led me to engineering, but over time my focus widened — from how systems operate to how entire industries function. From working on boil-off gas compressors aboard LNG tankers off South Korea to gas storage facilities in southern England, and standing beside pipelines in India early in my career, I began to wonder: how does all this connect? How do these millions of moving parts fit into a single, global picture?

Why I Chose Hydrogen

Over the years, I've watched hydrogen drift in and out of fashion. At times it has been overpromised, at others, dismissed outright. I've seen it treated as a silver bullet, a distraction, a failure, and a fantasy — often by people who have never had to make an energy system function day after day.

When the opportunity came to pivot my own work more directly toward hydrogen, I was genuinely excited by it. After years spent moving between oil and gas, renewables, and infrastructure projects, this felt like a chance to step closer to the centre of the transition — to work on something that was no longer just adjacent to change but shaped by it. It wasn't a rejection of my earlier career, but a continuation of it in a different direction.

Hydrogen kept coming back into the picture for a different reason. As renewable electricity began to scale, the same questions kept resurfacing: what happens when supply and demand no longer line up? What do we do with excess energy when the grid is full? How do we keep systems stable not just over minutes, but over weeks, seasons, and continents?

Hydrogen appeared repeatedly at the edges of those questions. Not as a solution on its own, but as something that changed how the rest of the system behaved.

Something that gave engineers more options, planners more flexibility, and energy systems more room to breathe.

That is why hydrogen matters to me. Not because it is elegant, efficient, or inevitable — but because it exposes the uncomfortable parts of the transition. It forces us to confront trade-offs, losses, and compromises we would rather avoid. It makes clear that a clean energy system is not only about generating electricity, but about deciding how energy moves, waits, and survives when conditions change.

Hydrogen sits at the point where many of the practical challenges of the energy transition come together. Understanding it has helped me understand the transition more clearly. My hope is that it will do the same for the reader.

The Role of Hydrogen

In recent years, the energy transition has become the defining framework for answering those questions. To work in energy today is to be part of a transformation reshaping technology, economics, and geopolitics all at once.

Yet for all the attention renewable electricity has rightly received, one crucial question remains: how do we build an energy system that works reliably when nature doesn't cooperate? The sun sets; the wind stills — but modern life does not.

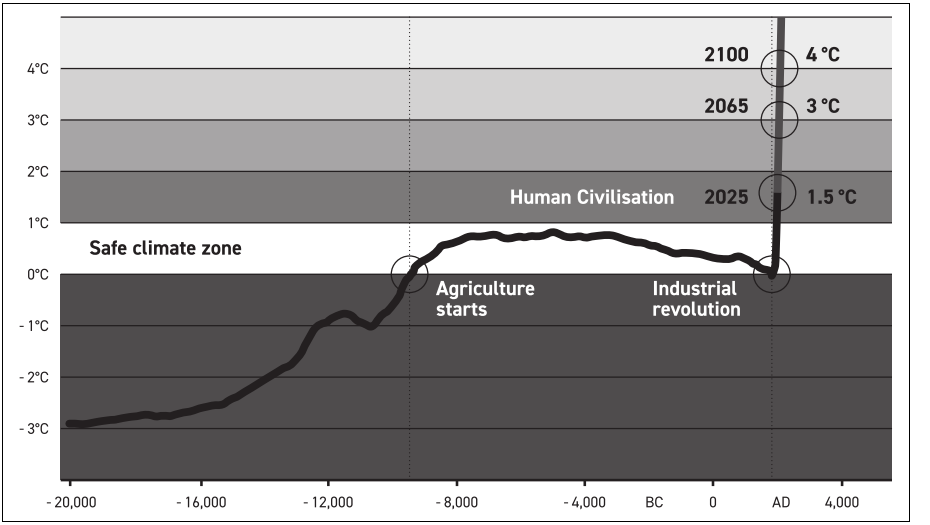
Hydrogen helps fill that gap. It isn't a replacement for renewables — it works alongside them. Electricity can do extraordinary things, but molecules still matter. Heavy industry, shipping, aviation, heating, and chemical manufacturing all depend on energy that can be stored and moved, and that's where hydrogen steps in.

Hydrogen sits precisely at this intersection — between electricity and molecules, innovation and infrastructure, ambition and reality. It's the quiet enabler of an energy system that can finally be both clean and continuous.

A Changing World

To be unaware of global warming today would require the isolation of a castaway. For anyone in energy, the science and urgency are undeniable. The question is not whether we must act, but how we can act fast enough. The transition, however, is not driven by climate alone. Energy security — the ability of nations to control and rely on their own energy systems — has always shaped geopolitics,

industry, and trade. Hydrogen therefore sits not only within sustainability debates, but within the strategic architecture of future energy systems.



Adapted from a climate visualisation originally attributed to Robert Corell, Heinz Center; illustration redesigned and updated for this book.

Over the past twenty-five years, I've watched renewables evolve from expensive curiosities into the backbone of modern energy systems. In the early days, critics dismissed wind and solar as unreliable and uneconomic. The wind doesn't always blow, they said, and the sun doesn't always shine. The grids of the 1980s and 1990s weren't built to handle such variability.

Through innovation, persistence, and policy support, however, those doubts faded. Today, renewables are central, not peripheral. The technology has matured, the costs have fallen, and the global mindset has shifted. Yet one fundamental challenge remains: we can't always use energy the moment we generate it.

Why Write This Book

The aim of this book is to explain, as clearly and simply as possible — simple enough, I hope, for my fifteen-year-old to read — how energy works today and how hydrogen can help shape the future.

It's for anyone who wants to understand what's happening beneath the headlines — how the pieces fit together, and where the real constraints lie.

The energy transition faces many challenges: cost, infrastructure, uncertain demand, regulation, financing, technology maturity, geography, coordination, and public perception. I'll explore each of these in turn, but also the progress already underway — and the role that innovation and collaboration are beginning to play.

My goal is to make a complex topic understandable without oversimplifying it. Anyone willing to think critically can grasp how it works and where it must go.

Looking Ahead

The next decade will be shaped not only by turbines, pipelines, or electrolysers, but by intelligence. Digitalisation and artificial intelligence are already reshaping how we model energy systems, forecast demand, optimise production, and maintain complex infrastructure. In hydrogen especially, where every molecule counts, data will be as valuable as the energy itself.

These tools can accelerate what once seemed slow, making the energy transition not just technological but digital — powered by algorithms as much as electrons.

In Closing

Hydrogen will not solve the energy transition on its own, but it helps to explain why the transition is harder — and more interesting — than it first appears. My hope is that this book helps readers see both the complexity and the promise of what lies ahead — to understand not just the science or the economics, but the story of how we got here and where we might go next.

If my son's generation is to inherit a world powered by cleaner energy, we owe them clarity, commitment, and courage. This book is my small contribution toward that goal — a guide to understanding the challenges, the potential, and the practical steps that can bring tomorrow's energy closer to today.

This is not a book about saving the world. It is a book about not breaking it while trying to change it.

1. Understanding Energy

'In the beginning the Universe was created. This had made many people very angry and has been widely regarded as a bad move.' The Hitchhikers Guide to the Galaxy, by Douglas Adams

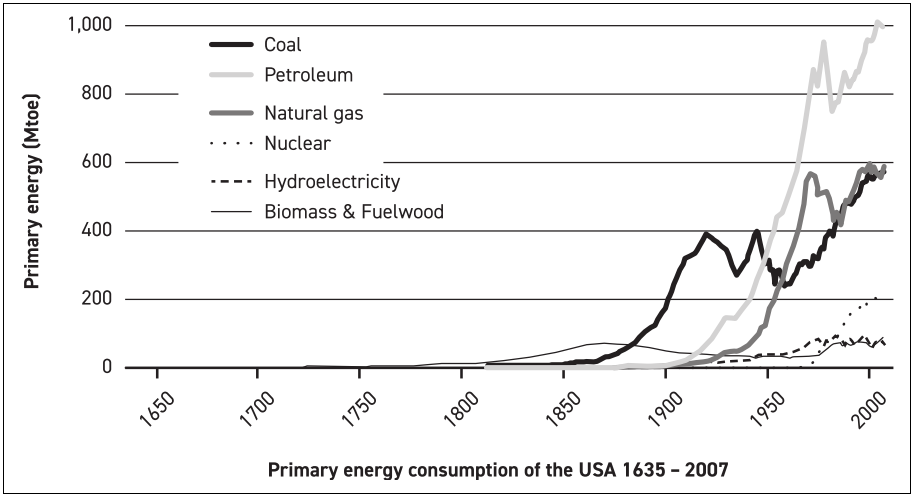
A New Kind of Transition

The world has never transitioned away from an energy source before.

It's often said that 'we've done this before' and that it shouldn't be difficult. We haven't.

Looking back through history, what we call energy 'transitions' have not been transitions at all, but additions; at no point have we fully replaced one dominant energy source with another.

At no point in the history of the planet - other than recently with a small down-turn of coal usage - have we completely replaced an existing energy source with a new one.



Adapted from Mikael Höök, Historical Energy Use of the United States from 1635 to 2000; illustration redesigned for this book.

Historical energy-use data from the United States illustrates this pattern clearly: at no point do we see the complete elimination of one energy source as another emerges.

Even biomass and fuelwood — despite their low energy density — still retain a place in the energy mix.

To understand the energy transition, you need to first understand energy as we use it today and where it comes from.

If Energy is Everywhere, Where Does It Actually Come From?

Nearly all the energy used on Earth can be traced back to the Sun. The Sun is the primary driver of almost every energy source, whether renewable or non-renewable. Its radiation fuels natural processes, sustains life, and underpins almost all of the planet's energy systems.

Solar energy itself is the most direct example. Technologies like solar panels and solar thermal systems capture sunlight and convert it into electricity or heat for human use. However, even energy sources that seem unrelated to the Sun, such as wind, biomass, and fossil fuels, are also solar in origin. Wind energy is generated by the uneven heating of the Earth's surface by the Sun, which creates air movement. Hydropower depends on the water cycle, which is driven by solar

energy that evaporates water, forms clouds, and precipitates rain that flows into rivers and dams.

Biomass energy, including wood and other plant-based fuels, is another indirect form of solar energy. Through photosynthesis, plants use sunlight to convert carbon dioxide and water into chemical energy stored in their tissues. When we burn biomass, we're simply releasing the solar energy that plants captured during their growth. Even fossil fuels—coal, oil, and natural gas—are ancient forms of biomass. They originate from the remains of plants and microorganisms that stored solar energy millions of years ago, which was then trapped and concentrated underground through geological processes.

In essence, whether we are burning coal, harnessing the wind, or capturing sunlight with solar panels, all of these energy forms are manifestations of solar energy accumulated and transformed over time. The only major exceptions are geothermal energy, which comes from the Earth's internal heat, and tidal energy, which is influenced mainly by the gravitational pull of the Moon. But for nearly everything else, the Sun is the ultimate source — the original engine powering our planet's energy systems.

This relationship can be summarised as follows:

Table: The Sun's Role in Global Energy Sources

Energy Source	Connection to the Sun	Type of Relationship	Explanation
Solar energy	Direct	Direct solar origin	Energy is captured directly from sunlight using photovoltaic panels or solar thermal systems.
Wind energy	Indirect	Driven by solar heating	Uneven heating of the Earth's surface by the Sun creates air pressure differences that generate wind.
Hydropower	Indirect	Driven by the water cycle	Solar radiation causes evaporation and precipitation, moving water through rivers and reservoirs.
Biomass energy	Indirect	Stored solar energy	Through photosynthesis, plants convert sunlight into chemical energy that can later be released as fuel.
Fossil fuels (coal, oil, natural gas)	Indirect	Ancient solar energy	Formed from ancient organic matter that originally stored solar energy millions of years ago.
Ocean and tidal energy	Minimal	Influenced by lunar gravity	Generated mainly by the gravitational pull of the Moon, with only minor influence from solar heating of oceans.
Geothermal energy	None	Internal planetary heat	Originates from heat within the Earth's mantle and core, independent of solar radiation.
Nuclear energy	None	Atomic energy from stellar processes	Released from atomic nuclei; uranium and other elements formed in earlier generations of stars, not by the Sun.

From Sunlight to Power

While the Sun provides the foundation for almost all energy on Earth, it is our ability to capture and convert this natural energy that defines how we use it today. The various energy sources—solar, wind, hydro, biomass, and fossil fuels—represent the primary energy drivers of our modern world. Each serves as a reservoir of natural energy that, through human technology, can be transformed into more usable forms, such as electricity, heat, or motion.

Electricity is the common thread linking these diverse sources. Whether it is generated by burning coal, channelling water through turbines, or absorbing sunlight with photovoltaic cells, the underlying process is the same: converting a natural energy flow into a controlled and transferable form. This conversion allows energy to be distributed across vast networks, powering homes, industries, and transportation systems with remarkable efficiency and reliability.

Our mastery of these transformations is what turns natural phenomena into practical energy. The movement of air becomes wind power; the flow of water becomes hydropower; and the stored energy in organic matter or ancient biomass becomes the fuels that drive engines and generate heat. Even sources beyond the Sun's influence, such as geothermal and nuclear energy, depend on the same fundamental principle—human control over natural energy processes.

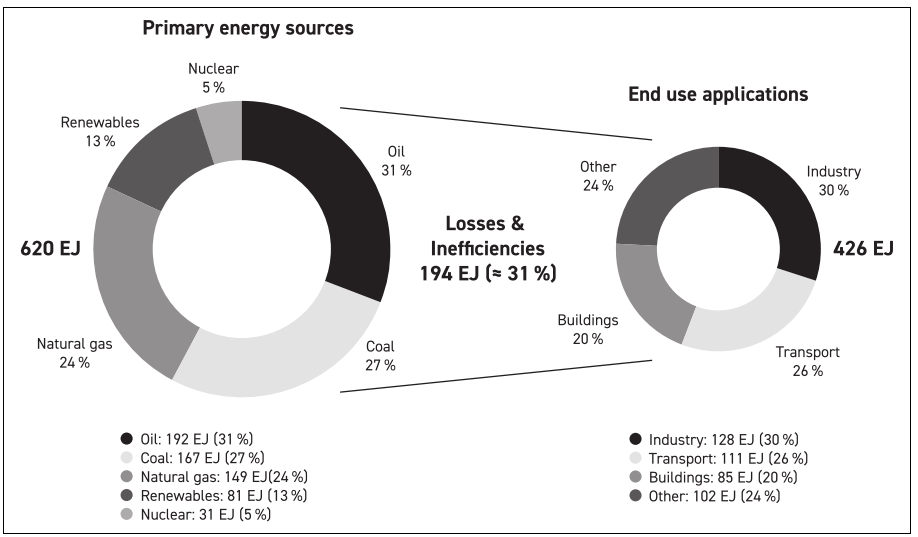
Ultimately, the energy that sustains our modern lives is not new energy created by human hands, but rather ancient energy repurposed. It is the result of our capacity to harness, channel, and refine the immense natural forces that have always shaped our planet.

The underneath table shows roughly how much of each of the carbon intensive energy sources the planet uses each year.

Material	Golden Age	Energy use in Golden Age EJ[1] / Year	Energy use today EJ[1] / Year
Wood + Biomass	Pre Industrial era	~ 20	~ 13
Coal	Late 19th – Early 20'th	~ 180[2]	~ 160[2]
Oil	Mid 20th to today	~ 150	~ 190
Natural Gas	Late 20th to today	~ 100	~ 145

1 EJ = Exajoule, **1 EJ** is roughly equal to the **annual energy consumption of a small-to-medium-sized country** (e.g., Norway or New Zealand)
2 was up to 192 EJ in 2022 and has only rolled back in the recent years

In total, the planet uses ~620 Exajoules (EJ) of Energy in a year, which is a considerable number. However, not all of the energy we produce is usable. One of the biggest issues within our current energy system is the amount of energy lost through inefficiency.



As you can see from the above graphic, energy conversion processes are almost never 100% efficient. During these processes, a portion of the energy is inevitably lost in various forms, reducing the overall efficiency of the system. These losses, often referred to collectively as energy losses, can arise from different causes depending on the type of conversion taking place.

Mechanical losses often result from friction or air resistance. Thermal losses occur when excess heat is generated and released into the environment instead of being utilised. Electrical losses arise due to resistance within circuits, while chemical losses stem from incomplete or unintended reactions.

Each type of loss can significantly influence the performance of an energy conversion system. For example, in a power plant, mechanical and thermal losses reduce the amount of electricity that can be produced from a given amount of fuel. By contrast, in batteries, electrical and chemical losses limit how much energy can be effectively stored and retrieved.

These losses explain why the type of energy we use — and how concentrated it is — matters so much. The more energy-dense a source, the fewer conversion steps we need, and the fewer losses occur.

This challenge does not disappear in the energy transition — if anything, it becomes more visible.

So, with the foundation established on where our energy comes from and how we use it, the next step is to understand energy density. Energy density refers to how much energy can be stored in a given amount of a material—essentially, how much usable energy we get per unit of weight or volume. The higher the energy density, the more power we can generate from a smaller quantity of fuel.

This concept helps to explain why our global energy systems have evolved the way they have. In the early stages of human civilization, energy needs were modest. People relied on wood for heating and cooking, and later on coal to power early industrial machines and steam engines. But as societies industrialised and energy demands soared—driven by transportation, manufacturing, and electricity generation—the limitations of low-energy-density fuels became clear.

Fossil fuels—coal, oil, and natural gas—offered a solution. Among them, oil and natural gas in particular possess far higher energy densities compared to traditional biomass like wood or even coal. This means they can release much more energy per kilogram, making them incredibly efficient sources of power. In practical terms, it takes far less oil to produce the same amount of energy you'd get from a much larger mass of wood or coal.

But energy density alone does not explain their dominance. Oil and natural gas are also relatively easy to transport and store, whether through pipelines, tankers, or other infrastructure, which makes them well suited to large-scale and global distribution. They are versatile, capable of powering everything from transport and industry to electricity generation, while also serving as feedstocks for a wide range of products.

Over time, this was reinforced by the development of extensive infrastructure and technology built around their extraction, processing, and use. Entire economic systems—and in many cases geopolitical relationships—became tied to them, creating a degree of inertia that is difficult to unwind. At the same time, fossil fuels offer a level of reliability that many alternatives struggle to match, providing consistent and controllable energy output regardless of weather or time of day.

In essence, the dominance of fossil fuels is not just a historical accident—it's a direct outcome of their high energy density, versatility, and reliability, combined with decades of investment in infrastructure and technology. These factors made fossil fuels the backbone of modern industrial society.

As the energy scholar Vaclav Smil has argued throughout much of his work, modern civilisation is fundamentally a story of energy and materials. The prosperity enjoyed by much of the world today rests upon vast flows of fuel, electricity, steel, cement, fertilisers, and countless other energy-intensive products that underpin modern life. Understanding these physical realities is essential when discussing any energy transition, because they define both what is possible and the scale of the challenge ahead.

The Price of Power

As our understanding of environmental issues has grown, the costs of fossil fuel dependence have become harder to ignore. The same fuels that powered industrial growth also release large amounts of carbon dioxide and other greenhouse gases, driving climate change, air pollution, and ecological damage. Rising temperatures, more frequent extreme weather, and stressed ecosystems make it clear that the current energy model is under increasing strain. Moving towards cleaner, renewable sources of energy is no longer just desirable — it is becoming necessary.

Beyond environmental concerns, the energy transition is also a matter of economic and geopolitical importance. Fossil fuels are finite resources concentrated in specific regions of the world, which creates supply vulnerabilities and exposes

nations to market volatility and political tensions. Renewable energy sources, such as solar, wind, and hydroelectric power, by contrast, are widely available and can be harnessed domestically. Transitioning to these sources enhances energy security, reduces dependence on imported fuels, and can help stabilise long-term energy costs.

The shift is also supported by technological and industrial progress. Advancements in renewable energy technologies, battery storage, and smart grids have made clean energy more accessible and affordable than before. What was once seen as an idealistic goal is increasingly a practical and economically competitive pathway. At the same time, the transition is creating new areas of industrial activity, from electric vehicles to hydrogen and sustainable construction.

Finally, the transition also has an ethical and generational dimension. Continuing to rely on high-carbon energy sources means passing on the consequences—environmental degradation, resource scarcity, and health risks—to future generations. Moving towards a cleaner, more sustainable energy system is, in part, about taking responsibility for those longer-term impacts. The transition becomes increasingly difficult to avoid if we are to maintain a stable and secure energy system.

Knowing this, the available options are relatively limited. For electricity generation, the primary sources are wind, solar, hydro, and nuclear. This provides a clearer view of where energy can come from, but each comes with its own set of challenges.

From Scepticism to Success: How Wind and Solar Overcame Early Doubts

When wind and solar power first appeared on the global energy scene, they were met with a wave of doubt — not just from entrenched fossil fuel interests, but from engineers, economists, and policymakers who saw them as noble ideas that simply wouldn't work in practice.

Wind, they said, was too unpredictable. The wind doesn't always blow, and when it does, it's not necessarily when electricity is needed. Solar faced the same criticism — the sun sets every night, clouds roll in, and energy demand doesn't politely align with daylight hours. Critics argued that such intermittent sources couldn't be integrated into electricity grids designed around steady, controllable power plants. Grid operators of 30 years ago worried that variable renewables would destabilise systems built on predictable supply.

Cost was another major obstacle. Early wind turbines were expensive to build and maintain, while solar panels were famously pricey — often seen as futuristic novelties rather than serious power sources. Many claimed these technologies would never compete with fossil fuels without heavy subsidies.

Wind turbines, some said, spoiled the view, made noise, and threatened birds. Solar farms were criticised for taking up space that could be used for agriculture or nature.

Yet, over the past few decades, almost every one of these objections has fallen away. Technological innovation, policy support, and economies of scale have transformed both wind and solar energy from expensive curiosities into the backbone of the clean energy transition.

Modern grids have evolved too. With advanced forecasting, flexible storage technologies, and smart digital controls, integrating variable renewables is now routine. What was once seen as "intermittency" is now managed with precision and foresight.

Wind turbines are quieter, more efficient, and less harmful to wildlife than ever before. Solar panels are cheaper than most other ways of generating power, and their footprints are shrinking as rooftops, parking canopies, and even building facades become part of the solar network.

Today, wind and solar aren't experimental — they're essential. They've gone from the margins of energy systems to their very centre, proving that the obstacles of the past were not insurmountable truths, but temporary challenges on the path to progress.

But even as these clean technologies now generate vast amounts of energy, they remind us of one lingering truth: we can't always use energy the moment we make it. The wind doesn't blow on command, and sunshine can't be bottled.

Electricity's Balancing Act

Electricity is one of the most essential yet fickle forms of energy we rely on. I once had a boss in the switchgear world who used to say that when it comes to electricity, expectations are simple: 100% availability—blackouts are not acceptable.

The reality behind that expectation is far more complex. Unlike fuels that can be stored for later use, electricity must be generated and consumed almost instantaneously. This means that at every moment, the amount of power produced must match the amount being used. Even a small imbalance can cause fluctuations in

voltage or frequency, leading to inefficiencies or, in extreme cases, widespread blackouts. Maintaining this delicate equilibrium requires constant coordination across the entire electrical grid.

The grid itself operates as a continuous balancing act between supply and demand. When consumption rises—such as during a heatwave when millions of air conditioners are switched on—power generators must quickly increase output. When demand falls, generation must be reduced to prevent overloading the system. This fine-tuned control is managed through a process known as frequency regulation, where grid operators monitor the system frequency (typically 50 or 60 hertz, depending on the region) and make rapid adjustments. If generation exceeds demand, the frequency rises; if demand exceeds generation, it drops. Maintaining that steady frequency is essential for keeping the entire system stable.

What makes electricity management even more complex today is the growing integration of renewable energy sources. As renewables like wind and solar now supply significant portions of electricity, operators must manage variability across the grid. Advanced forecasting, storage, and grid management techniques help address these challenges in real time. To counteract these fluctuations, operators now rely increasingly on energy storage technologies, such as batteries and pumped hydro, as well as smarter grid systems that use predictive analytics to anticipate changes in both production and consumption.

In essence, the stability of our electricity supply depends on a constant, dynamic effort to balance what is needed with what can be delivered. It is one of the most complex engineering challenges in modern society—keeping an invisible, instantaneous flow of energy steady, reliable, and responsive in the face of ever-changing human and natural demands.

Energy Storage: Keeping the Lights On

As more renewable energy is added to the grid, these challenges will only become more pronounced. Some of the most impressive work over recent decades has been the ability of grid operators to integrate these sources into daily life. There will, however, still need to be available capacity to cover periods of low generation, and this is where energy storage comes into play.

Energy storage lies at the heart of solving the challenges posed by a renewable-heavy grid. As more solar and wind power come online, the need to store excess energy during periods of high generation—and release it when production drops—becomes increasingly vital. Effective storage provides the flexibility that

the grid needs to maintain stability and reliability, ensuring that electricity remains available even when the sun isn't shining or the wind isn't blowing.

Currently, the most common forms of energy storage include pumped hydro, lithium-ion batteries, and thermal storage systems. Pumped hydro, the oldest and still the largest-scale technology, works by moving water between two reservoirs at different elevations, generating electricity when water flows downward through turbines. Lithium-ion batteries, on the other hand, have rapidly advanced in recent years, offering high efficiency and fast response times. They are particularly well suited for short-duration storage and frequency regulation, helping to smooth out fluctuations in power generation and demand.

To balance production and demand over hours, days, or seasons, storage solutions are essential as renewable penetration grows.

The topics introduced in this chapter only scratch the surface of the complex relationship between energy, materials, industry, and modern society. Readers interested in exploring these subjects in greater depth may wish to read Vaclav Smil's *How the World Really Works*, which provides an excellent examination of the physical realities and constraints that underpin the modern world. Many of the challenges discussed throughout the remainder of this book ultimately stem from those same realities.

2. Hydrogen — The System Adhesive

"Our house, in the middle of our street,

Our house, it has a crowd,

There's always something happening, and it's usually quite loud."

'Our House' by Madness

ATCO — The Quiet Normality of Change

Unlike the house in the song, this one is not "quite loud" — it is, in fact, entirely unremarkable. And that is precisely the point.

On a quiet suburban street in Sherwood Park near Edmonton, Alberta, a new house sits comfortably among rows of familiar bungalows and two-storey homes. The siding is light blue, the driveway freshly paved, and from the outside it looks like any other home in Canada. There is nothing about it that signals innovation, let alone revolution.

But step inside, and something subtle has changed.

The furnace hums softly in the basement, the water heater glows with quiet efficiency, and the oven preheats for dinner. Everything feels normal — and yet nothing here burns natural gas. This house, the first of its kind in Canada, is fuelled entirely by hydrogen.

Developed by ATCO and Qualico, the Sherwood Park "hydrogen home" is not a futuristic prototype so much as a deliberate exercise in restraint. The engineers who built it did not reinvent the concept of home; they changed what flows through its pipes. In this single building, the invisible molecules heating the air,

boiling the water, and cooking the food carry no carbon — only hydrogen, producing nothing but water vapour as they burn.

The scene is almost disarmingly ordinary. There is no futuristic hum of electronics or gleaming laboratory equipment. No dramatic architecture announcing a break with the past. Just a comfortable family home. That quiet normalcy is the most striking part of the project.

Because when energy systems truly change at scale, they rarely announce themselves. They integrate. They recede into the background. They become reliable, invisible, and ultimately taken for granted.

This ordinariness matters. People do not adopt new energy systems because they are technically elegant; they adopt them because they fit into daily routines without friction. Transitions that feel disruptive tend to stall. Those that feel familiar tend to scale. In Sherwood Park, hydrogen is not presented as something exotic or experimental. It behaves as energy always has — dependable, unobtrusive, and assumed.

For utilities, that familiarity is equally important. Instead of replacing every gas pipeline and every household appliance, hydrogen offers the possibility of adapting what already exists. Infrastructure built over decades does not need to be torn out and replaced wholesale; it can, in principle, be modified incrementally and pragmatically. For communities, this continuity matters just as much. No new habits to learn. No cultural overhaul. Just warmth, hot water, and everyday comfort — delivered differently beneath the surface.

It is a small project in physical terms, but a revealing one conceptually. It shows that decarbonisation does not have to arrive as a visible rupture. Sometimes the most consequential change is the one that leaves daily life untouched.

Hydrogen is often described as an energy carrier. In system terms, it may be more useful to think of it as something else entirely — a way of connecting parts of the energy system that do not naturally align.

But before exploring what hydrogen can connect, it is worth understanding what it actually is.

The Lightest, Fastest, Most Restless Element

If hydrogen were a character, it would be restless and hard to pin down. It's the lightest element known, so light that individual atoms can slip through metal, leak from tanks, and even escape Earth's gravity over time. That featherweight nature

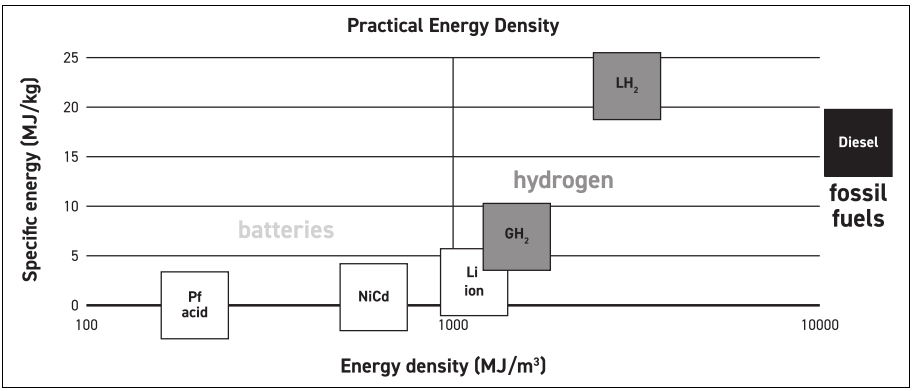
gives hydrogen a superpower — it carries more energy per kilogram than any other fuel. Roughly three times as much as gasoline, in fact.

But here's the catch: that energy is spread thin. By volume, hydrogen is incredibly diffuse — meaning you must squeeze it to 700 times atmospheric pressure, or cool it to –253°C, just to get enough of it into a tank to be useful. Those extremes make it expensive to store and transport.

Then there's its personality. Hydrogen burns invisibly and explosively — not a problem if you handle it right (rockets do, after all), but it demands respect. The upside is that hydrogen itself contains no carbon, so when it burns it produces no CO_2 from the fuel and no soot — primarily water vapour, although combustion in air can still produce NOx.

How Hydrogen Stacks Up

Energy Carrier	Energy per kg (MJ/kg)	Energy per Liter (MJ/L)	What It Means
Hydrogen (compressed)	~120	~5.6	Lightest by far — great for rockets, tricky for cars.
Gasoline	~46	~34	Denser, easy to store — the fossil benchmark.
Ammonia	~18.6	~11.5	Easier to handle than hydrogen, used as a carrier.
Lithium-ion battery	~0.9	~2.5	Low by weight, but rechargeable and efficient.



Hydrogen wins the race by weight but loses by volume. That trade-off — light yet bulky — is why engineers are still figuring out where hydrogen really fits. It shows more immediate promise in shipping, where space and refuelling constraints are less severe, while aviation remains technically challenging — and less so for your daily commute.

Hydrogen is the smallest thing in the universe, and yet it might help power its future. One proton, one electron — that's it. You can't get simpler than that. But simplicity doesn't mean insignificance: hydrogen makes up about three-quarters of all normal matter. The stars burn it, water contains it, and nearly every organic molecule — from oil to oatmeal — has a bit of it inside.

On Earth, though, hydrogen is a bit of a recluse. You'll never find it just floating around in the air, because it's too reactive — it immediately latches onto something else, like oxygen (to form water) or carbon (to form methane or other hydrocarbons). That means if we want pure hydrogen, we must make it by splitting it loose from whatever it's bound to. And that takes energy.

Not a Source — a Carrier

Here's the trick: hydrogen isn't a source of energy. It's an energy carrier — a middleman. Like a battery, it stores energy that must come from somewhere else first. What matters, then, is how that hydrogen is made. That's where the "colours" come in — not literal hues, but a kind of shorthand for the production process. These labels are informal and not all are equally important; they mainly describe carbon intensity and maturity:

🌈 **Hydrogen Colour Codes — Complete Reference (2026)**

Colour / Name	Feedstock / Energy Source	Main Process / Technology	Carbon Footprint	Notes & Examples
🩶 **Grey Hydrogen**	Natural gas, coal, or oil	Steam Methane Reforming (SMR), Coal Gasification, Partial Oxidation	High CO_2 emissions	Most common today; no carbon capture.
💙 **Blue Hydrogen**	Fossil fuels + Carbon Capture	SMR, ATR, or POX with CCS	Low CO_2 (60–90% captured)	Transitional low-carbon method; infrastructure compatible.
💚 **Green Hydrogen**	Renewable electricity (solar, wind, hydro)	Water Electrolysis (PEM, Alkaline, SOEC)	Near-zero CO_2	True zero-emission hydrogen; cost falling rapidly.
💛 **Yellow Hydrogen**	Grid electricity mix (renewable + fossil)	Water Electrolysis	Variable CO_2	Depends on grid mix; transitional method.

Colour / Name	Feedstock / Energy Source	Main Process / Technology	Carbon Footprint	Notes & Examples
💗 **Pink Hydrogen**	Nuclear electricity	Electrolysis powered by nuclear	Very low CO_2	Also called Red Hydrogen in some regions.
🔴 **Red Hydrogen**	Nuclear energy (heat and electricity)	High-temperature electrolysis using nuclear heat	Very low CO_2	Similar to pink but emphasizes nuclear heat utilization.
🟣 **Purple Hydrogen**	Nuclear + electrolysis + heat integration	Thermochemical water splitting using nuclear reactors	Very low CO_2	Variant of pink/red; used in nuclear R&D contexts.
🩵 **Turquoise Hydrogen**	Natural gas, biogas, or flare gas	Methane Pyrolysis — thermal or plasma-assisted	Low CO_2 (solid carbon byproduct)	Emerging tech: plasma process avoids CO_2 entirely.
⚫ **Black Hydrogen**	Bituminous coal or lignite	Coal gasification (no CCS)	Very high CO_2	Similar to brown; term varies by coal type.
🟫 **Brown Hydrogen**	Coal (sub-bituminous)	Coal Gasification (no CCS)	Very high CO_2	Oldest, dirtiest form; large in China and India.
🟩 **Biohydrogen**	Biomass, waste, or organic feedstock	Biomass Gasification, Dark Fermentation, Biophotolysis	Carbon-neutral	CO_2 recycled by plants; can achieve net-zero with CCS.
🌿 **Gold Hydrogen**	Biogas or biomass + Carbon Capture	Reforming of renewable methane with CCS	Net-zero or negative CO_2	Hybrid of bio and blue hydrogen pathways.
🧊 **White Hydrogen**	Naturally occurring underground deposits	Geological extraction	Zero production emissions	Found in natural seeps (France, Mali, USA); exploration phase.

For the rest of the book, the focus is on the pathways that matter at scale: grey (status quo), blue (transitional), green (target).

As the range of production pathways suggests, hydrogen itself is neither clean nor dirty. What matters is how it is made.

Today, most of the hydrogen used worldwide is produced using a method known as steam methane reforming — commonly referred to as grey hydrogen. It has been the backbone of industrial hydrogen production for decades, supplying refineries, fertiliser plants, and chemical manufacturers with a steady, inexpensive source of hydrogen.

Grey hydrogen dominates for simple reasons: it is technically mature, highly efficient, and relatively cheap. But those advantages come at a cost. The process

relies on natural gas and releases large quantities of carbon dioxide into the atmosphere. For every kilogram of hydrogen produced, roughly ten kilograms of CO_2 are emitted.

This creates a paradox at the heart of the hydrogen debate. Hydrogen is often promoted as a clean energy carrier because, when used, it produces no carbon emissions. Yet most hydrogen in use today carries a substantial carbon footprint embedded in its production. The fuel itself may be clean at the point of use, but the system that produces it often is not.

That contradiction explains much of the scepticism surrounding hydrogen. It is not that hydrogen cannot support decarbonisation — it is that, so far, it mostly has not. Any serious discussion of hydrogen's role in a low-carbon energy system must therefore begin by acknowledging this baseline.

Why Hydrogen Matters — Beyond Electrification

The question is now how this comes together with our energy needs. Well firstly not everything can be electrified and there are many reasons for that. Electrification is often presented as the cornerstone of a clean energy future. Solar panels, wind turbines, and batteries are central to reducing emissions, and for many applications, electricity can indeed replace fossil fuels. Yet, electricity alone cannot decarbonise everything.

Certain sectors—heavy industry, long-haul transport, and seasonal energy storage—pose challenges that direct electrification struggles to overcome. Steel and cement production require extremely high temperatures and chemical reactions that are difficult to achieve with electricity alone. Ships, airplanes, and heavy trucks demand energy-dense fuels that current battery technology cannot economically supply. Even storing renewable electricity over long periods, such as from summer to winter, remains prohibitively expensive with today's batteries.

These challenges reveal the inherent limits of direct electrification. Batteries, while effective for short-term storage and light vehicles, are heavy and bulky. The electric grid may lack the capacity to supply certain industrial processes at the scale and reliability needed. Moreover, renewable electricity is intermittent: the sun does not always shine, and the wind does not always blow. Without a flexible way to store or transport energy, electricity cannot reliably meet all energy demands.

This is where hydrogen enters the picture as a versatile "bridge molecule." Hydrogen can convert excess renewable electricity into a storable chemical form,

creating energy that can be transported over long distances and used when and where it is needed. It provides molecular energy in sectors that electrons cannot easily reach, from high-temperature industrial processes to aviation fuel. In essence, hydrogen extends the reach of renewable electricity beyond the limits of wires and batteries.

The role of hydrogen can be visualised as a circular flow of energy. Renewable electricity is used to produce hydrogen through electrolysis. Hydrogen can then be stored, transported, or converted into heat, fuel, feedstock, or even electricity again when needed. This flexibility allows hydrogen to bridge both temporal gaps—storing energy for months or seasons—and spatial gaps—moving energy across regions or even internationally. By acting as a carrier of clean energy, hydrogen enables decarbonisation in places where electrification alone falls short.

One of the most useful ways I have heard hydrogen explained in recent years did not come from a paper, a model, or a policy document, but from an analogy.

It is often compared to jam.

Think of renewable energy like fresh strawberries. They are abundant, full of vitality, and unquestionably good for you. But strawberries are also perishable. You can enjoy them at their best only for a short time, and after that, they spoil. Abundance alone is not enough, timing matters.

To make strawberries last, we turn them into jam — a form that can be stored, transported, and enjoyed long after the harvest. Something is lost in the process, but something crucial is gained: longevity, flexibility, and usefulness beyond the moment of production.

Hydrogen plays a similar role in the energy transition. When we use renewable electricity to produce hydrogen, we are effectively turning a perishable energy resource into something that can endure. Excess electricity generated on a windy afternoon or a bright summer day can be converted into hydrogen and stored, preserved for later use when conditions change and demand rises.

In doing so, we accept a trade-off. Converting electricity into hydrogen involves losses; not all of the original energy survives the process. But what remains is no longer bound to the instant of generation. It can be stored for days, weeks, or even seasons. It can be transported to where electricity cannot easily go. It can be converted back into electricity, or used in other forms, when the system needs support.

In that sense, hydrogen is not a rival to renewable electricity, but its complement. It does not compete with wind or solar; it makes them usable beyond the moment they are produced. Hydrogen is the jam that keeps the strawberries of the energy transition from going to waste — not perfect, not free of loss, but necessary if abundance is to become reliability.

This distinction matters. The challenge of the energy transition is no longer simply generating clean energy but ensuring that it is available when and where it is needed. Hydrogen addresses that challenge not by replacing electricity, but by allowing energy to persist beyond the fleeting conditions under which it was created.

Hydrogen as the "System Connector"

Hydrogen's value in the energy transition lies less in what it replaces than in what it connects. Its defining role is to link parts of the energy system that do not naturally fit together — across energy forms, sectors, and geography.

At a physical level, hydrogen allows energy to move between states that are otherwise difficult to reconcile. Electricity can be converted into hydrogen through electrolysis, storing energy in chemical form rather than forcing it to be used immediately. Surplus electricity from Danish wind farms, for example, can flow into electrolysers that produce hydrogen, storing energy until it is needed. That same hydrogen can later be converted back into electricity, such as in fuel cells supplying power into the German grid, allowing generation and consumption to be linked across time.

Hydrogen also connects energy use across sectors that have traditionally operated in isolation. Renewable electricity can be carried into industrial furnaces, chemical processes, heavy transport, or heating systems where direct electrification is difficult or impractical. This is not about a single pathway, but about optionality: electricity, transport, and industry can support one another rather than compete for limited grid capacity.

Geography introduces a further constraint. Renewable resources are unevenly distributed, while demand is often concentrated elsewhere. Hydrogen reduces this mismatch by allowing energy to move between regions. In Australia, renewable electricity can be converted into hydrogen and further into ammonia, enabling large-scale export. That energy can then be shipped to Japan and used for power generation or industrial applications, arriving in a form compatible with existing infrastructure and market structures.

This system-level connectivity is only possible because of hydrogen's underlying chemistry. Electrolysis enables electricity-to-hydrogen conversion, while fuel cells reverse the same reaction to release electricity again. Hydrogen can also be bound into carriers, such as ammonia or methanol, increasing energy density and simplifying transport. For example, hydrogen produced from North Sea offshore wind can be partially converted into methanol and transported to German industrial hubs, where it can be used as fuel or feedstock.

Reversibility adds another layer of flexibility. The same chemical processes that store energy can, when required, release it again. This allows hydrogen to function not just as a carrier, but as a buffer — absorbing surplus energy when supply exceeds demand and returning it when conditions change.

Hydrogen's compatibility with existing infrastructure further reinforces its connective role. It can be blended into gas pipelines, stored in underground caverns, or integrated into industrial systems with relatively modest changes. Pilot projects in Germany and as we have seen above in Canada, already inject hydrogen into regional gas networks, while industrial plants blend hydrogen or ammonia into operations to reduce emissions without complete system redesign.

Taken together, these properties explain why hydrogen functions less like a standalone technology and more like a connective layer within the energy system. It does not optimise any single component, and the losses involved in conversion are real. But by allowing variable generation, continuous demand, distant resources, and diverse end uses to function together, hydrogen helps stabilise the system as a whole.

In that sense, hydrogen does not replace existing energy systems — it joins them. Like any connector, it works best where interfaces are compatible, a limitation that will matter throughout this book.

Hydrogen: The System's Memory

The challenge that remains after electrification is not generation, but time.

Electricity must be produced and consumed almost instantaneously. Renewable energy intensifies this constraint: supply rises and falls with weather and daylight, while demand follows human schedules, seasons, and economic activity. Short-term balancing can be managed over seconds, minutes, and hours. What becomes difficult is spanning longer gaps — days, weeks, or entire seasons.

This is the function hydrogen performs within the energy system. By converting electricity into a chemical form, energy can be preserved beyond the moment it is generated. Hydrogen allows surplus electricity to be carried forward in time, stored when supply exceeds demand and released later when conditions reverse. In that sense, hydrogen gives the energy system memory.

That memory is not free. Converting electricity into hydrogen and back again involves losses, additional infrastructure, and cost. Hydrogen does not optimise efficiency; it trades efficiency for flexibility. Its value lies not in doing everything well, but in doing something that electricity alone cannot — sustaining energy availability across timeframes that wires and batteries struggle to cover.

Understanding hydrogen in this way clarifies both its promise and its limits. It is neither a universal solution nor an optional add-on. It is a specific response to a specific constraint: the mismatch between when renewable energy is generated and when energy is needed.

The question that follows is therefore not whether hydrogen belongs in the energy system, but where it actually makes sense to use it.

That question — of fit rather than faith — is the subject of the next chapter.

3. Hydrogen Use Cases

‘Hydrogen doesn’t have a chicken and the egg problem, the problem is purely egg, egg, egg, with egg being the demand!’ Statement from the *Plenum Hydrogen World Technology Expo Hamburg 2025*

In the opening chapters, I have focused primarily on why hydrogen is needed within the energy transition. A reasonable next question is therefore not whether hydrogen matters in theory, but where it enters the real world, and under what conditions it begins to make sense in practice.

That is the purpose of this chapter.

Hydrogen use cases emerge not because today’s energy systems fail, but because the goals they were built to serve are no longer sufficient. This chapter examines where that shift in objectives is already shaping practical decisions.

More importantly, this chapter approaches hydrogen through the way energy systems really behave. Hydrogen becomes relevant where electricity alone struggles to meet multiple requirements at once — where energy must be stored for long periods, moved across sectors, or delivered with high reliability rather than maximum efficiency. In these roles, hydrogen is less a fuel than a means of translating energy from one part of the system to another.

It is not intended to prove that hydrogen will dominate the energy system, nor to suggest that every proposed application is sensible or scalable. Neither will it attempt to resolve the economics of each use case. Questions of cost, competitiveness, and financing — including the levelised cost of hydrogen — are addressed later in this book, where they can be treated with the depth they require.

It is also important to acknowledge that there is no single, universally accepted view — even among experts — of how hydrogen will ultimately be deployed. Engineers, economists, policymakers, and system operators often reach different

conclusions, not because one group is careless or uninformed, but because they optimise for different constraints, time horizons, and definitions of risk. In a transition of this scale, such divergence is not a weakness of the debate; it is evidence that the system is still in motion.

Many objections to hydrogen are grounded in today's realities: today's costs, today's infrastructure, today's inefficiencies. Those objections are often technically correct. But an energy transition, by definition, is a process unfolding over time. What does not yet make sense everywhere must first make sense somewhere — otherwise the transition never progresses beyond theory.

A similar pattern appears in debates around re-electrification — for example, the use of hydrogen and fuel cells for backup or emergency power. These applications are often dismissed on the basis of round-trip efficiency, comparing electricity converted to hydrogen and back again. Yet this framing overlooks the alternative that currently dominates; vast fleets of diesel generators that exist precisely because reliability, autonomy, and resilience are valued more highly than efficiency alone. If decarbonisation is a central objective of the transition, then leaving these systems unchanged simply because a cleaner alternative is not perfectly efficient, represents a category error. The relevant comparison is not against an idealised electric system, but against the fossil-based infrastructure that already exists.

A related assumption appears more broadly across the transition: that entire system functions will ultimately be handed over to a single dominant solution, most commonly batteries. This is unlikely. Energy systems have never operated as monocultures, and the transition does not change that reality. Different use cases impose different requirements — on energy density, response time, autonomy, refuelling or charging patterns, spatial constraints, and risk tolerance. No single technology optimises all of these simultaneously. The outcome is therefore not technological replacement, but technological layering. This is particularly evident in mobility, where policy frameworks increasingly assume a mix of solutions rather than a single winner, reflecting the fact that diverse transport tasks impose fundamentally different operational demands.

In the previous chapter, we encountered the ATCO house in Edmonton — a project that, at first glance, appears easy to dismiss. I was sceptical myself until the system logic behind it was explained. Once it was, the use case became coherent, even if it remained limited in scope. That experience has shaped how I approach hydrogen applications more broadly.

Dismissing use cases simply because they conflict with prior assumptions misses the point. In practice, projects that fail to meet economic or operational requirements do not get built. Those that do are signals — not of inevitability, but of pressure within the system. Where consensus is lacking, practical deployment becomes one of the few ways to test assumptions against reality.

These use cases should therefore be understood as stepping stones rather than end states — mechanisms through which demand, infrastructure, and capability are built incrementally over time.

This chapter treats hydrogen use cases as indicators, not promises. Some will scale, whilst others will not. What matters is understanding why they appear, and what that reveals about the energy system that we are trying to transform.

Industrial Feedstock & Process Heat

Of all the potential entry points for clean hydrogen, the most natural place to begin is where hydrogen is already in use today. In these sectors, the challenge is not introducing a new molecule but replacing carbon-intensive inputs with cleaner alternatives.

Among all proposed hydrogen applications, its role as an industrial feedstock and source of high-temperature process heat is the least speculative. Hydrogen is already deeply embedded in sectors such as chemicals, refining, and steelmaking — not as a fuel of convenience, but as a functional input. In these contexts, hydrogen is valued for what it does chemically, not simply for the energy it carries.

Hydrogen as Feedstock, Not Fuel

In much of the public debate, hydrogen is framed primarily as an alternative energy carrier, competing with electricity or other fuels. In several industrial sectors, however, this framing misses the point. Here, hydrogen is not valued for the energy it releases when burned, but for the role it plays in enabling specific chemical reactions. It is a material input, not merely a source of heat.

This distinction is clearest in industries such as ammonia and methanol production, petroleum refining, and large parts of the chemical sector. Hydrogen is already used at scale in these processes, often in quantities that dwarf potential early uses in transport or power generation. In ammonia production, hydrogen is the foundational building block. In refining, it is essential for desulphurisation and upgrading heavier fractions. In these contexts, the presence of hydrogen is not optional; it is intrinsic to how the process functions.

Decarbonisation therefore does not begin with introducing hydrogen into these sectors, but with changing how that hydrogen is produced. Historically, it has been generated from fossil fuels, with associated carbon emissions treated as an unavoidable by-product. The transition challenge is to replace that carbon-intensive supply with low-emissions alternatives, while preserving the underlying industrial processes that depend on the molecule itself.

Electrification alone cannot solve this problem. While electricity can provide energy, it cannot replace a missing reactant. Where hydrogen participates directly in chemical transformations, there is no equivalent electrical substitute. Process redesign may reduce hydrogen demand at the margins, but it cannot eliminate it without fundamentally altering the product being made.

In these industries, hydrogen is therefore not competing with electrons; it is competing with carbon. The question is not whether hydrogen belongs in these processes, but whether its production can be decoupled from fossil fuels at the scale and reliability that heavy industry requires.

Steel and High-Temperature Process Heat

Steel production sits at the centre of the industrial decarbonisation challenge. It is emissions-intensive, globally traded, and deeply embedded in infrastructure, construction, and manufacturing supply chains. Unlike many downstream applications, steelmaking is not easily relocated or redesigned without significant economic and political consequences. As a result, any credible pathway to industrial decarbonisation must engage directly with how steel is produced.

The interest in hydrogen within this sector is not driven by fashion or policy signalling, but by chemistry. In conventional blast furnace routes, carbon serves a dual role: it provides energy and acts as the reducing agent that removes oxygen from iron ore. Direct reduced iron (DRI) processes using hydrogen replace that chemical function without altering the fundamental requirement of reduction. The appeal of hydrogen here is therefore structural. It enables an established metallurgical pathway to operate with significantly lower direct emissions.

High-temperature process heat more broadly follows a similar logic. Many industrial processes require sustained temperatures well above those easily delivered by conventional electric heating, often with tight tolerances and continuous operation. While electrification can and will play a role in some cases, it is not a universal substitute. In environments where heat demand is intense, continuous, and integrated into complex process chains, hydrogen offers a familiar combustion-based route that aligns with existing industrial practices.

What constrains deployment is not conceptual viability, but practical reality. Steel plants are capital-intensive assets with lifetimes measured in decades. Operators do not change primary production routes lightly, nor do they make decisions based on short-term policy signals alone. Movement occurs when multiple pressures align; regulatory requirements, customer demand for low-carbon materials, and access to finance willing to support large-scale conversion.

Hydrogen in steel and high-temperature industry is therefore best understood not as a disruptive technology, but as an enabling one. It offers a pathway to reduce emissions within processes that are otherwise difficult to alter, while remaining bounded by the economic, logistical, and temporal constraints of heavy industry.

Why Industrial Substitution Is Constrained

If the technical rationale for using hydrogen in certain industrial processes is increasingly clear, the pace at which substitution occurs remains slow and uneven. The difficulty lies in execution rather than principle. Heavy industry is shaped by physical assets, supply chains, and regulatory environments that are designed for stability, not rapid change.

The first constraint is existing capital stock. Industrial facilities represent multi-billion-euro investments, optimised around specific feedstocks, temperatures, and operating regimes, and are often deeply integrated into surrounding infrastructure. Retrofitting or replacing them is costly, disruptive, and rarely undertaken unless driven by sustained external pressure or clear long-term signals. Incremental improvements are common; fundamental process changes are not.

Supply chains add a further layer of inertia. Industrial processes do not operate in isolation, but as part of tightly coupled networks of upstream suppliers, logistics providers, and downstream customers. Introducing hydrogen at scale requires coordinated changes well beyond the point of use, including production, transport, storage, and safety systems. Each additional interface introduces cost, complexity, and risk, making coordinated deployment difficult even where individual actors are willing.

Operational and regulatory considerations further constrain substitution. Handling hydrogen at industrial scale necessitates new infrastructure, revised operating procedures, and regulatory approval. Permitting frameworks are cautious by design, reflecting legitimate concerns around safety, environmental impact, and land use. Even where technology is mature and political support is

strong, approval timelines can extend over many years, creating uncertainty for investors and industrial operators alike.

Scale and reliability of supply are equally critical — and remain among the least resolved challenges. Heavy industry consumes hydrogen in volumes far beyond early pilot projects or demonstration plants. Steelmaking, chemicals, and refining require continuous, predictable supply to maintain safe and efficient operation. Meeting even a fraction of this demand with low-emissions hydrogen implies production, transport, and storage capacity at a scale that does not yet exist in most regions.

A small number of projects are beginning to test what such supply chains might look like in practice. In Denmark, Everfuel is developing integrated green hydrogen production and distribution aimed at supplying industrial off-takers, moving beyond demonstration toward early commercial volumes. In Sweden, Stegra is attempting to align large-scale hydrogen production directly with industrial demand, illustrating both the ambition of these pathways and the degree of coordination required across energy, infrastructure, and industry.

These initiatives are significant not because they solve the problem, but because they reveal its magnitude. Even the most advanced projects today represent early steps toward supply levels that heavy industry would require at scale.

Underlying all of these constraints is cost. Low-emissions hydrogen is currently more expensive than fossil-derived alternatives when viewed through short-term operating economics. For industries exposed to international competition, these cost differentials matter, even where decarbonisation objectives are well understood. The ability to absorb, offset, or pass on higher costs depends on policy design, customer demand, and access to capital — issues that will be addressed later in this book.

Taken together, these constraints do not negate hydrogen's industrial role. They explain why progress is gradual, regionally concentrated, and uneven across sectors. The challenge is not whether alternatives exist in principle, but whether they can be deployed at industrial scale within the combined constraints of time, cost, infrastructure, and regulation. Understanding these limits is essential to interpreting how industrial hydrogen use evolves in practice, rather than how it might appear in theory.

Having established where hydrogen is structurally embedded in industry, the next question is where it fits in transport — and just as importantly, where it does not.

Mobility — Where Hydrogen Works

If there is one area where discussions around hydrogen tend to polarise, it is mobility. Cars, trains, ships, and aircraft are frequently discussed as a single category, despite having very different operating conditions, performance requirements, and constraints.

This polarisation is partly a consequence of visibility. Transport is personal, highly visible, and closely tied to everyday experience. It is also driven by the fact that, unlike many industrial processes, several mobility applications already have viable low-carbon alternatives. As a result, hydrogen's role in transport is necessarily more selective.

For that reason, this section does not assess hydrogen by vehicle type, but by operating constraint. The question is not whether hydrogen can move vehicles, but where its characteristics align with requirements that are difficult to meet through electrification alone — particularly where range, payload, utilisation, and refuelling time dominate system design.

Let us start with the biggest elephant in the room:

Battery Cars and Hydrogen Trucks.

One particularly insightful discussion on this topic comes from Rory Sutherland, Vice Chairman of the marketing think tank Ogilvy, in a recent podcast interview, where he offers a behavioural reframing of how technologies are adopted in practice. He often points out that focusing only on engineering data can lead to poor decisions. People adopt technologies as much for emotional, psychological, and experiential reasons as for purely technical ones.

Sutherland's famous thought experiment is to invert the question: what if electric vehicles had come first? Suddenly, petrol engines — with their flammable liquids, mini-explosions, hundreds of moving parts, and constant maintenance — would seem bizarre. The point is simple: our perceptions shape the debate just as much as physics does.

If we look at the facts through this lens, battery-electric vehicles make a great deal of sense. For the average driver travelling perhaps 20–60 km per day, overnight charging is convenient. While it would technically be possible in the future to install a micro-electrolyser at home to produce hydrogen — and I have, over the years, seen a number of small-scale hydrogen systems designed with this idea in mind, though far fewer today — few people would want to operate what is, in effect, a mini-refinery in their garage. Likewise, a battery-electric car typic-

ally uses 70–85% of the original renewable electricity that went into charging it, whereas hydrogen delivers significantly less to the wheels due to conversion losses incurred when electricity is converted to hydrogen and back again. Finally, echoing Sutherland's point about simplicity, EV drivetrains are mechanically elegant and low maintenance from an engineering standpoint. To be fair, hydrogen-electric drivetrains do not add much in this context beyond the storage characteristic.

Hence, for the average driver, batteries feel natural, cheap, quiet, and low stress.

This is not a judgement call; it is a reflection of where the use case aligns cleanly with the technology.

That does not imply a single solution for all transport — only that different constraints favour different outcomes.

Hydrogen enters the equation when the use case changes, and that change usually comes with distance and load. Payload matters. In logistics, there is a clear trade-off when it comes to weight: batteries are heavy, and the heavier the truck, the larger the battery required, which in turn reduces the payload that can be transported. Long-haul trucks often drive 600–1,000+ km per day and are frequently "on the clock". In this context, time becomes a critical constraint. The amount of time required to charge very large battery systems does not align well with high-utilisation operations. The infrastructure implications are also significant, particularly along major motorway routes.

There are additional considerations beyond distance and weight, particularly operational patterns where fleet operations run within fixed logistical structures.

A practical example of hydrogen working in high-utilisation operations comes from STILL in Hamburg, which has converted its intralogistics forklifts to hydrogen fuel-cell power. By producing its own fuel-cell systems in-house, STILL has designed forklifts that can refuel in minutes and operate continuously over multi-shift periods, without the downtime associated with battery swapping. The hydrogen refuelling infrastructure is concentrated at depots and logistics hubs, making distribution predictable and efficient. This real-world case demonstrates the principle that hydrogen thrives where operational patterns are regular, utilisation is high, and downtime is costly. It also illustrates a pathway for hydrogen beyond heavy trucks: the same logic applies to high-mileage or fleet vehicles, suggesting that while battery-electric cars dominate typical daily driving, hydrogen vehicles could occupy well-defined niches where rapid refuelling, predictable routes, or extreme utilisation are key considerations.

At the end of the day, the argument is not that hydrogen trucks are better. It is that hydrogen trucks fit the physics and economics of certain heavy-duty, high-utilisation segments.

However, it would be wrong to assume that hydrogen has no role at all in personal vehicles. Even though battery-electric cars dominate the mainstream, hydrogen cars may still have a market where the use case matches hydrogen's strengths.

Battery-electric vehicles make strong sense in regions of high population density, but as density thins out, infrastructure challenges re-emerge. If one were to draw a blanket over Europe from Berlin to Barcelona and from London to Milan, BEVs and their associated infrastructure make clear sense. Move further north toward Trondheim, or to the geographic fringes of Europe in places such as Portugal or Romania, and infrastructure density — and available electrical capacity — becomes more challenging.

The situation becomes more pronounced when looking at regions such as Australia, large parts of the Americas, Central Asia, and western China. In these contexts, refuelling infrastructure that resembles today's liquid-fuel model may remain attractive simply because it aligns with geography and distance.

Other considerations also play a role: high-mileage taxis with limited downtime, corporate fleets operating from central depots, regions with significant green-hydrogen production near ports or industrial clusters, cold climates where battery performance degrades, or consumers who value predictability and refuelling familiarity over maximum efficiency. Some users will accept lower efficiency in exchange for convenience, simplicity, and reduced cognitive load.

Shipping — Fuels Before Platforms

Compared with road transport, the debate around hydrogen in shipping is generally less polarised. This is partly because the constraints are clearer, and partly because the sector has long operated with a diverse set of fuels and operating models. The question in shipping is less about whether change is possible, and more about which pathways align best with regulation, energy density, and global operations.

Regulatory pressure is already shaping decisions. In parts of the world, particularly Europe, emissions rules for maritime transport are tightening, with requirements focused on both carbon intensity and absolute emissions. These measures are not abstract future signals; they are beginning to influence vessel design, fuel choices, and investment decisions today. For shipping companies operating glob-

ally, regulatory fragmentation also matters, favouring solutions that can scale across routes rather than remain confined to local niches.

Batteries play a role where the use case allows it. Short-distance ferries operating on fixed routes, particularly in Norway and other Nordic countries, have successfully adopted battery-electric propulsion. In these applications, predictable schedules, frequent docking, and relatively modest energy requirements make electrification practical. However, these conditions do not extend to long-haul shipping, where range, payload, and continuous operation dominate.

For deep-sea vessels, energy density becomes the defining constraint. Hydrogen in its pure form is difficult to store onboard at the volumes required for long voyages, pushing attention toward hydrogen-derived fuels instead. Methanol and ammonia, both produced using green hydrogen as a feedstock, offer higher volumetric energy density and greater compatibility with existing maritime infrastructure. As a result, the shipping sector is increasingly focused on fuels rather than propulsion platforms.

This is reflected in the strategies of major operators such as Maersk, which has committed to deploying vessels capable of operating on low-emissions fuels, including green methanol. These moves are not framed as technological bets, but as risk management responses to regulation, fuel availability, and customer expectations.

In this sense, shipping illustrates a recurring pattern in the hydrogen transition. Hydrogen rarely appears directly at the point of use. Instead, it underpins a family of molecules that better meet operational constraints. The role of hydrogen here is foundational rather than visible — embedded upstream in fuel production, while vessels themselves adapt to carriers that balance energy density, safety, and global scalability.

The Morning the Rails Went Quiet

Dawn arrives gently over the Bavarian countryside, washing the fields around Augsburg in pale gold. A thin band of mist clings to the meadows east of the city, softening the silhouettes of barns and hedgerows. It is the kind of morning suspended between seasons: cold enough for breath to cloud the air, warm enough that the frost on the grass seems unsure how long it will remain.

At the small station where the first regional trains pass through, everything is still. A handful of commuters gather along the platform, footsteps muted on the concrete, scarves pulled tight against the chill.

They expect the familiar soundscape of a diesel morning. For decades, this line has been served by sturdy units that announce themselves long before they arrive—engines rumbling through the valley, a faint haze of exhaust drifting above rooftops. The community has learned to live with it: the lingering scent of fuel, the low vibration underfoot, the instinct to raise your voice when trains idle nearby.

But today, as the clock edges toward the hour, something unfamiliar glides into view.

There is no roar. No buildup. Just a train emerging from the mist as quietly as a thought. Its silver-and-blue outline catches the early light, softening the steel beneath. The Siemens Mireo Plus H coasts into the station with a grace that makes several people look up, uncertain whether they simply hadn't heard it—or whether something has changed.

When the train stops, the doors open with a gentle hiss. Warm air spills into the cold. Inside, the carriage glows softly. Everything looks ordinary. And yet nothing feels quite the same.

A man in his sixties rests his hand on the table, as he always does, waiting for the familiar tremor of an idling engine. The surface stays still. He frowns.

Across the aisle, a student removes her earbuds. He leans toward her.

"Did they switch to electric trains?"

She smiles. "Hydrogen."

He pauses. "Fuel cells?"

She nods. "It makes electricity and water. That's it."

He runs his hand across the table again, feeling the absence of vibration, then gives a small nod—as if conceding a point in a debate he hadn't realised he was having.

Up in the cab, the driver settles into his seat. He trained on diesel units, and on electric ones where infrastructure allowed. This feels familiar—same controls, same responses. The difference lies behind the panel, where fuel cells convert hydrogen into electricity, supported by a battery that smooths power delivery and captures braking energy. Even the sound is different: a low, steady hum.

The signal turns green. The train pulls away—not with the usual diesel grunt, but with an effortless glide. It slips past the houses at the edge of the station and out into open countryside, fields stretching toward the foothills of the Alps.

Inside, a boy presses his face to the window.

“It doesn’t sound like anything,” he says.

His friend tilts his head, listening.

The clatter of wheels remains, but everything else—the vibration, the engine noise, the smell—has disappeared. It is the kind of quiet that makes the world outside feel sharper: frost melting on grass, shadows stretching behind trees, the faint outline of the Alps forming in the distance.

For years, the question of how to decarbonise regional rail hung unresolved over Bavaria. Electrification offered a solution in theory, but in practice it was slow, expensive, and difficult to justify on long rural stretches. Diesel persisted—not out of preference, but inertia.

Hydrogen changes the equation by changing very little. Same tracks. Same schedules. Lightly adapted depots. A different outcome.

At the next stop, passengers board without hesitation. A mother with a pram. A cyclist guiding her bike. A nurse settling into her seat. No one discusses the technology. They move as though nothing has changed. That is the point.

The train continues south, climbing towards the Allgäu. Sunlight burns through the mist. Farmhouses appear along the route, their steep roofs catching the light. The landscape feels unusually vivid, as though the absence of noise has brought it into sharper focus.

Miles away, an operator monitors the train’s telemetry. Fuel cell temperatures stable. Battery cycling within range. Hydrogen consumption exactly as expected. The system behaves not like an experiment, but like part of the fleet. He takes a sip of coffee and looks away.

At another station, an elderly couple step onto the platform. The woman glances back at the train.

“It’s strange,” she says. “You don’t notice anything.”

Her husband nods. “That’s how it should be.”

The train pulls away. The platform remains quiet—no fumes, no haze, no lingering trace of fuel. Just birds in the distance and the soft crunch of gravel.

Further along the line, the train rounds a wide curve, revealing the Alps rising at the horizon. Inside, passengers look up one by one, drawn to the view. It is a sight that has greeted travellers for generations.

But something has changed. The world outside is the same. The way they move through it is not.

Quietly. Cleanly. Almost invisibly.

And perhaps that is how the energy transition succeeds—not through spectacle, but through mornings like this. Where nothing in the rhythm of life appears altered, yet everything beneath it has shifted. Where the rails fall silent not because the world has slowed, but because it has begun, almost imperceptibly, to breathe differently.

Rail — Electricity Where It Works, Hydrogen Where It Doesn't

Rail transport offers one of the clearest examples of how use cases should drive technology choice. Wherever a nation has a policy to electrify its network, electricity remains the preferred solution by a wide margin. Electrified railways are among the most energy-efficient modes of transport available and, when powered by low-carbon electricity, are a very effective way to reduce emissions from both passenger and freight movement.

In the European Union, just over half of the railway network is electrified — nearly 57% of total lines — and these sections carry most of the rail traffic today. Electrification rates vary significantly from country to country, with some like Luxembourg and Belgium well above the EU average, while others in the Baltic States and parts of eastern Europe remain much less electrified.

Where overhead lines are present, electric rolling stock delivers strong performance, low operating emissions, and improved reliability.

But the economics and logistics of electrification are not uniform everywhere. Extending overhead power lines across remote or low-traffic routes can be expensive and slow, especially where population density is low, freight volumes are limited, or challenging terrain increases construction and maintenance costs. In such places, the traditional logic of electrification — build wires, use electric locomotives — becomes harder to justify if measured only by traffic density or short-term cost.

Just as with passenger cars, where batteries make sense for urban and regional trips but lose advantage in high-distance and high-time-pressure logistics, the same type of constraint applies on rail. Non-electrified lines remain common outside dense networks, and diesel traction still serves them today because alternatives are not yet sufficiently viable at scale.

In those contexts, hydrogen emerges as a potential bridging solution. Fuel-cell locomotives and hydrogen-powered railcars do not require continuous overhead power, yet they can replace diesel traction with a low-emissions alternative where full electrification is impractical or uneconomic. Early pilot programmes in Europe and elsewhere are exploring hydrogen rail in contexts where electrification would be costly and traffic levels too low to justify wires — for example, on non-electrified regional routes in parts of Scandinavia and the United Kingdom.

The logic is the same as with other modes: where electricity can be delivered directly and at scale, it remains the king; where it cannot, and existing operations rely on carbon-intensive diesel for continuous service, hydrogen or hydrogen-derived solutions may be a sensible complement until broader electrification or another low-carbon infrastructure arrives.

Aviation — Clear Boundaries, Incremental Change

Aviation is one of the sectors where the physical boundaries are clearest, and where over-promising would be particularly damaging to credibility. Energy density, safety requirements, and payload constraints mean that hydrogen is not a solution for long-haul aviation, and there is little value in pretending otherwise. Intercontinental flights operating for many hours at high altitude impose requirements that liquid hydrogen, at least with foreseeable aircraft architectures, is unlikely to meet.

Where hydrogen does enter the aviation discussion is at the other end of the spectrum. Short-haul and regional routes, with lower payloads, shorter ranges, and more frequent turnaround, present a fundamentally different operating environment. In these segments, hydrogen — whether used directly in fuel cells or combusted in modified turbines — becomes technically conceivable. Even here, the pathway is gradual, constrained, and highly dependent on aircraft design and infrastructure availability.

Alongside direct hydrogen use, sustainable aviation fuels (SAF) play a central role. SAF offers a route to emissions reduction that is compatible with today's aircraft and fuel infrastructure, making it particularly attractive for medium- and long-haul operations where new propulsion systems are hardest to introduce. Many

SAF pathways rely on green hydrogen as a feedstock, linking aviation decarbonisation indirectly to hydrogen even where aircraft never carry the molecule themselves.

This dual approach is reflected in the strategies of major players. Airbus has publicly explored hydrogen-based aircraft concepts focused on short- and medium-range applications, while explicitly recognising the limits imposed by physics and certification requirements. At the same time, companies such as ZeroAvia are working on hydrogen-electric propulsion for regional aircraft, targeting incremental entry into service rather than wholesale disruption of the sector.

Crucially, none of this should be understood as a distant or speculative vision. Aviation operates on long development and fleet-replacement cycles, often measured in decades. Against that backdrop, the prospect of hydrogen-enabled regional aircraft and widespread SAF adoption within the next 20 to 25 years is not unrealistic — it is consistent with how the industry has historically evolved.

As with shipping and rail, aviation reinforces a recurring theme. Hydrogen is unlikely to appear everywhere, nor all at once. Its role emerges where operational boundaries allow it, and where existing alternatives struggle to meet tightening emissions constraints without compromising safety, performance, or economics.

Power System Balancing & Storage — Function, Not Optimisation

Power system balancing and storage are among the most conceptually difficult aspects of the energy transition to discuss clearly. This is not because the technologies involved are poorly understood, but because the problem is too often framed as a single optimisation exercise. When viewed through the narrow lens of efficiency alone, many proposed solutions appear irrational. When viewed through the lens of system function, they often make sense.

A critical distinction that is frequently overlooked is that power system balancing is not a single problem. Short-term balancing — managing fluctuations over seconds, minutes, or hours — is fundamentally different from long-duration or seasonal balancing, where energy must be stored or made available across days, weeks, or even months. These challenges differ not only in scale, but in character. Solutions that perform exceptionally well in one domain are often poorly suited to the other.

Numerous technologies have attempted to bridge this gap, including compressed air and other forms of mechanical storage. In many cases, their limitations are not

primarily technical, but economic. They can function as designed, yet struggle to compete once system costs, utilisation rates, and infrastructure requirements are taken into account.

Daily balancing, weekly adequacy, and seasonal energy shifting are not variations of the same challenge; they are fundamentally different problems, and no single technology addresses all three.

Batteries illustrate this clearly. They are highly effective at short-duration balancing, frequency control, and daily cycling. As such, they are already becoming an integral part of modern electricity systems. But extending this logic to long-duration or seasonal storage quickly encounters physical and economic limits. Storing energy for prolonged periods, with minimal losses and at system-relevant scale, demands different characteristics altogether.

It is in this context that hydrogen begins to appear — not as a competitor to lithium-ion batteries, but as a response to a gap where clean alternatives remain limited. The debate is often framed as a comparison between batteries and hydrogen, yet this framing misrepresents the role hydrogen is being asked to play. In long-duration storage and system resilience applications, hydrogen is rarely competing against batteries. It is competing against the absence of viable solutions, or against fossil-based fallback mechanisms that already exist.

In practice, the comparison is not hydrogen versus batteries, but hydrogen versus diesel generators, fossil-fired backup plants, and— increasingly—curtailment.

This distinction matters because it reframes the efficiency debate. Conversion losses associated with electricity-to-hydrogen-to-electricity pathways are real and should not be ignored. But efficiency is a local metric. It describes how well a component performs in isolation. Power systems, by contrast, are judged by global metrics: stability, reliability, and the ability to meet demand under all conditions. A perfectly efficient system that fails when it is needed most is not fit for purpose.

Re-electrification sits squarely within this logic. Generating electricity from hydrogen is not about maximising efficiency; it is about ensuring availability. It addresses situations where electricity cannot be stored easily, transmitted reliably, or delivered continuously. In these cases, the relevant comparison is not against an idealised fully electrified system, but against what is currently used to guarantee power when the grid is unavailable or insufficient.

This includes roles such as backup generation, black-start capability, and system recovery after extended outages — functions that modern electricity systems cannot simply wish away.

Importantly, this does not imply that hydrogen-based power generation is a near-term solution for all such cases, nor that current economics are favourable across the board. The point is more fundamental. A decarbonised energy system must still perform all the functions that today's system performs, namely balancing supply and demand, absorbing shocks, and providing power during extended disruptions. If those functions are ignored because they are difficult to decarbonise, the transition remains incomplete.

These challenges are most visible in remote, islanded, or weak-grid contexts, but they also exist quietly within advanced energy systems.

Seen this way, power system balancing and long-duration storage reinforce a central theme of this book. Hydrogen's relevance does not arise from theoretical elegance or efficiency optimisation. It arises from the need to preserve system functionality as the underlying energy mix changes. Where electricity alone can meet those needs, it should. Where it cannot, alternative mechanisms must be considered — even if they challenge conventional efficiency-based thinking.

Understanding this distinction — between local efficiency and system stability, between optimisation and function — is essential to understanding why hydrogen continues to surface in discussions about the future energy system. It is not because it is the best solution everywhere, but because some problems remain unsolved without it.

Diesel Generators — A Big Missing Piece in Decarbonisation

One of the most striking absences in the current energy transition narrative is attention to diesel generators. These assets play a very large role in power reliability in many parts of the world, yet they receive far less policy and investment focus than many of the "obvious" decarbonisation use cases, such as industrial feedstock, mobility, and electricity sector transformation. That gap is increasingly hard to reconcile with climate goals.

Globally, the diesel generator market is substantial — various industry analyses place its current value in the tens of billions of U.S. dollars and project continued growth over the coming decade. Estimates suggest the market was worth roughly USD 19–25 billion in the early 2020s and is expected to continue expanding, with some projections indicating values approaching USD 30–40 billion by around

2030 and beyond, supported by increasing demand for reliable power supplies, infrastructure development, and backup capacity in data centres and industrial facilities.

Diesel generators are used widely — not only as emergency backup for hospitals, factories, telecoms, and critical infrastructure, but also as primary power sources where grid access is limited or unreliable. Their sheer scale means they contribute significant CO_2 emissions every year, both through direct fuel combustion and through lifecycle impacts associated with global fuel supply chains. While there is no single global tally for CO_2 from diesel generators alone, the combustion of diesel fuel produces roughly 2.7 kg of CO_2 per litre consumed, illustrating that even modest usage contributes cumulatively to energy-related emissions.

Despite this, most existing policy frameworks have not yet taken aim at this sector in anything like the way they have with transport or industrial emissions. Regulations and incentives have focused on areas considered "obvious" for reduction — renewables, electrification of vehicles, industrial decarbonisation pathways — leaving diesel generators largely in the background.

This omission matters. From a system function perspective, diesel generators represent a fallback in the absence of viable alternative solutions for certain reliability and resilience roles. In places where grid outages are frequent, where backup must be guaranteed, or where infrastructure is weak or absent, diesel gensets fill a real need. In transitioning away from fossil fuels, simply acknowledging that physics and current economics make hydrogen conversions inefficient today cannot be the endpoint; a long-term solution for replacing these generators is required if we are to genuinely decarbonise the broader energy system.

Addressing this gap does not mean rushing policy mandates without planning. It means recognising — explicitly — that diesel generators form a large, emissions-intensive segment of the current energy landscape that has been underdiscussed in transition planning, and that we need to begin thinking about credible pathways for their decarbonisation. This might include hydrogen alternatives, hybrid systems, or other low-carbon backup strategies that align with broader system balancing needs.

Digital Infrastructure — A New Driver of Energy Demand

One use case that has risen rapidly up the energy agenda in recent years barely featured in long-term planning only a decade ago: hyperscale data centres and artificial intelligence infrastructure. Unlike traditional commercial loads, these facilities require vast quantities of continuous, high-quality electricity, with avail-

ability measured not simply in annual uptime but in near-continuous operation. Individual campuses are now being designed with power requirements measured in hundreds of megawatts, while some proposed developments are approaching gigawatt scale. In the United Kingdom alone, data centres currently represent around 1.6 GW of installed capacity, yet projects seeking new transmission connections now total approximately 60 GW—equivalent to around three-quarters of the entire transmission-level demand connection queue. Such figures illustrate not only the extraordinary pace of digital growth, but how rapidly AI and data infrastructure are becoming major drivers of future electricity demand.

For hydrogen, the opportunity is unlikely to lie in replacing grid electricity, but in supporting the resilience and flexibility these facilities require. Data centres demand uninterrupted operation, making backup power, long-duration energy storage, microgrids and resilient on-site generation increasingly important considerations. Today these functions are still largely fulfilled by diesel generators. As operators seek to decarbonise without compromising reliability, hydrogen and fuel-cell technologies offer one potential pathway for replacing these fossil-based backup systems while preserving the resilience on which digital infrastructure depends.

Perhaps the more significant lesson, however, is what this trend reveals about the wider energy transition. Hyperscale data centres, electrolyser plants, battery gigafactories, semiconductor manufacturing and other electricity-intensive industries are increasingly competing for the same scarce resource: timely access to high-capacity grid connections. In many regions, the limiting factor is no longer the ability to generate clean electricity, but the ability to connect it where it is needed. The future energy transition will therefore be shaped not only by how much low-carbon electricity we can produce, but increasingly by how intelligently we expand, manage and prioritise the networks that deliver it. In this context, hydrogen is not competing with electrification; it is one of several technologies that must coexist within an energy system where access to electricity is becoming as strategically important as electricity itself.

Synthetic Fuels & Derivatives — Molecules That Travel

Synthetic fuels and hydrogen derivatives occupy a distinct yet closely related position within the hydrogen landscape. They are not best understood as standalone energy solutions, but as extensions of hydrogen's role as an industrial feedstock — adapted to contexts where hydrogen itself is difficult to store, transport, or use directly.

At their core, synthetic fuels, such as ammonia, methanol, and other e-fuels, are mechanisms for embedding hydrogen into more manageable molecules. These carriers trade conversion efficiency for practicality. They offer higher volumetric energy density, compatibility with existing infrastructure, and easier handling across long distances. In doing so, they allow hydrogen to move through parts of the energy system that would otherwise remain inaccessible.

This logic mirrors what has already been discussed in industrial contexts. Ammonia and methanol are not speculative products; they are already produced and consumed at scale today. The difference in a low-carbon system is not their function, but their input. When produced using green hydrogen rather than fossil-derived feedstocks, these molecules become vectors for decarbonisation rather than sources of emissions.

The relevance of synthetic fuels therefore emerges most clearly in sectors where direct electrification is constrained and where pure hydrogen faces practical limitations. Shipping and aviation are prominent examples, but they are not unique. In each case, the driver is not enthusiasm for complexity, but the need to balance energy density, storage, safety, and existing asset lifetimes.

It is also important to be clear about what this section is not claiming. Synthetic fuels are not an efficient way to use renewable electricity, nor are they likely to be deployed everywhere. Their role is conditional and selective. They exist because some parts of the energy system require molecules that can be stored, transported, and handled using familiar methods — even if that comes at the cost of additional conversion steps.

Seen this way, synthetic fuels and derivatives do not represent a divergence from hydrogen's core logic, but a continuation of it. They extend hydrogen's reach where direct use becomes impractical, reinforcing the broader theme of this chapter: hydrogen appears not where it is ideal, but where system constraints leave few viable alternatives.

These pathways are examined in greater detail later in the book, particularly in the chapter on Power-to-X and system integration.

Where Hydrogen Probably Does Not Belong

A credible discussion of hydrogen use cases must also be clear about where hydrogen is unlikely to play a meaningful role. This is not a judgement on the technology itself, but a recognition that energy systems benefit from choosing the simplest solution that meets the requirement.

In contexts where direct electrification already performs well, hydrogen adds little value. Passenger cars operating in dense, well-electrified regions are a clear example. Here, battery-electric solutions are already scalable, efficient, and increasingly cost-effective. Introducing hydrogen into these applications would duplicate infrastructure without delivering commensurate system benefits.

Hydrogen is also poorly suited to applications where system complexity outweighs functional gain. Wherever additional conversion steps introduce operational risk, maintenance burden, or coordination challenges without solving a binding constraint, simpler approaches tend to prevail. Energy transitions succeed not by maximising technological diversity everywhere, but by avoiding unnecessary complication where it offers no clear advantage.

Finally, hydrogen is unlikely to displace solutions that already align well with existing infrastructure, behaviour, and economics. Where alternatives are mature, familiar, and improving rapidly, hydrogen faces a high threshold to justify its inclusion. In such cases, restraint is not a weakness of the transition, but a sign of discipline.

Being explicit about these limits strengthens the broader argument. Hydrogen does not need to be everywhere to matter. Its value lies in addressing specific pressures within the system, not in replacing technologies that already work well.

Use Cases as a Portfolio, Not a Bet

Taken together, the use cases explored in this chapter point towards a common conclusion: hydrogen does not scale through conquest, but through sequencing. It enters the energy system where constraints are tightest, alternatives are limited, and system pressure is highest.

Different applications mature at different speeds. Industrial feedstocks move first because hydrogen is already embedded in the chemistry. Mobility applications follow selectively, shaped by distance, payload, and utilisation. Power system balancing and long-duration storage evolve more slowly, constrained by economics and infrastructure, but driven by the need for resilience. Synthetic fuels extend hydrogen's reach where direct use becomes impractical.

This progression is uneven by design. Progress will vary by region, by sector, and by policy environment. Some projects will advance quickly; others will stall or fail. That unevenness is not evidence of incoherence, but of adaptation. Energy systems are complex, and transitions unfold through accumulation rather than uniform rollout.

Viewed this way, hydrogen use cases are not predictions of scale, but indicators of pressure. They reveal where existing solutions strain under new objectives — decarbonisation, security of supply, and system resilience — and where new capabilities begin to form in response.

Hydrogen will not arrive everywhere at once. It will arrive where systems are under the greatest strain — and from there, it will spread outward, incrementally, as infrastructure, confidence, and capability accumulate.

4. Technology Pathways and Innovation

"The best way to predict the future is to invent it." — Alan Kay

The Polymer Frontier

On a forested hill just outside Vancouver's city centre sits Simon Fraser University, a place where the glass walls of modern laboratories look out over the Pacific and the distant blue silhouettes of the North Shore Mountains. The campus is a world unto itself, suspended between sea and sky, built high enough that on clear days you can see the entire city glinting far below like a mosaic of possibility. Fog moves quickly here, rolling over the concrete walkways and curling around the brutalist buildings, lending the place an almost monastic quiet in the early morning hours. It is in one of these labs, tucked into the heart of the chemistry department, that Steven Holdcroft begins most of his days.

Steven has been at this long enough to have become an elder statesman of polymer chemistry, yet something about him still seems defiantly youthful. Perhaps it is the way he leans forward when he talks, hands moving as though shaping invisible molecules in the air. Perhaps it is the pleasure he takes in small victories — an unexpected result on an impedance plot, a membrane sample that survives a longer stress cycle than predicted, a student's experiment that works on the first try. Or perhaps it is simply the sense that the next breakthrough could happen today, or tomorrow, or even this afternoon, if the team tries one more formulation before going home.

His domain is anion exchange membrane technology — AEMs — a field that, for many years, existed in a kind of scientific half-light. Too complex for mainstream attention, too immature for industrial adoption, neither alkaline nor PEM, neither familiar nor fully alien. It was a technology waiting for its moment, held back by the stubborn difficulties of polymer stability in alkaline environments and the intricate chemistry of ion transport through polymer backbones. But Steven saw something in it long before most others did. He saw an elegant simplicity beneath the complexity: a pathway to efficient electrolysis using cheaper materials, oper-

ating without precious-metal catalysts, opening the door to low-cost green hydrogen at scales the world would one day demand.

He had been working on AEM technology at SFU for years, refining polymer chemistries, exploring how chain architecture and substituent groups shaped conductivity, investigating how membranes swelled, degraded, cracked — or, on the best days, refused to crack at all. His lab became known internationally for its work on hydroxide-conducting polymers, producing some of the earliest candidates capable of functioning at the high pH and high currents needed for serious electrochemical applications. For decades, AEM was seen as charming but impractical. Now the world was finally beginning to understand what Steven had grasped much earlier: the future might depend on getting this right.

Back in 2017, long before hydrogen became a talking point in boardrooms and parliaments, Steven helped found a startup with a small group of collaborators and former students. They called it Ionomr Innovations — spelled with a quiet nod to "ionomer," the class of materials that had captured his imagination for most of his career. The company was built on membranes first developed in his lab, materials that proved unexpectedly tough, elegantly efficient, and capable of performance that had eluded AEM developers for years. They branded the materials Aemion and later Aemion+, polymers built not on the standard quaternary ammonium groups that had dominated the field, but on sturdier cationic architectures designed to resist alkaline attack.

Ionomr grew quickly. Funding arrived from Canada, the United States, and Europe. Hydrogen's moment — long predicted, long delayed — began to surface. As it did, the significance of what Steven and his team had built became harder to ignore. AEM electrolysis promised a bridge between cost and performance: lower materials cost than PEM, higher potential efficiency than traditional alkaline systems, and the possibility of producing hydrogen using catalysts far cheaper than platinum or iridium. What had once seemed academically interesting was beginning to look economically consequential.

As the company matured, Steven stepped back from operational roles. The transition felt natural. Ionomr was no longer an academic spin-out; it had become an industrial engine. It needed managers, financiers, and commercial strategists. Steven returned to his hilltop lab at SFU, where the work was quieter, slower, and somehow more fundamental. New students arrived. Others moved on. Polymers were synthesised, cast, tested, destroyed, and remade. Ideas emerged from late-night discussions. Problems appeared from nowhere, as they always do. Solutions followed days or weeks later, arriving with the same quiet suddenness.

Sometimes Steven walked the edge of the campus in the morning before heading to the lab. The view from Burnaby Mountain stretches across the inlet toward ships anchored offshore, waiting to be filled or emptied. The energy system of today and the energy system of tomorrow coexist in that view. Massive tankers and container vessels sit beneath power lines that snake into the city, while wind turbines on distant ridges turn slowly in shifting air. Somewhere below, electrolysers — still imperfect, still expensive — split water into hydrogen and oxygen. The future he works toward is not yet evenly distributed. But it is visible enough to guide the work.

SFU itself is preparing for a leap of its own. On land just beyond the main academic precinct, plans are taking shape for an electrolyser test field — a dedicated area where membranes developed in Steven's lab can be pushed far beyond the comfort of benchtop instrumentation. It will be a place of valves and stacks and cooling loops, where hardware breathes and groans in ways only real equipment can. Here, AEM membranes will meet practical realities: fluctuating temperatures, impurities, mechanical stresses, and operational currents that reveal whether a polymer can survive the journey from lab curiosity to industrial certainty. For Steven, this is not a departure from his work, but its continuation.

Inside the lab, the tools of his trade are scattered across benches: beakers clouded with polymer solutions, glass plates for casting films, humidity chambers, conductivity rigs, tensile testers, microscopes for examining the fine fractures that quietly narrate membrane lifetimes. To an outsider it may look chaotic. To Steven though, it is a landscape of clues. Every membrane cast is a hypothesis. Every failure carries information. Every success, however modest, moves the field forward.

He sometimes refers to what he is chasing as the "next great leap forward." Not because he seeks dramatic language, but because he understands that AEM electrolysis will not become mainstream through incrementalism alone. Something must shift — a chemistry, a structure, a way of integrating membranes into systems that tolerate real-world abuse. The leap might come from a new polymer backbone, from an overlooked stabilising group, or from the noisy middle ground where chemistry meets engineering. Steven does not pretend to know where it will come from. He only knows that it must be pursued.

When he speaks with his students, he listens more than he lectures. He challenges assumptions without hardening them into positions, nudging ideas forward rather than knocking them down. What animates him most is not that his generation might finish the work, but that the next one might see paths he never could.

At conferences, colleagues often find him surrounded by small clusters of researchers leaning in. They ask about polymer stability, C–N bond cleavage, water uptake, conductivity trade-offs, and the uneasy balance between swelling and mechanical strength. Steven listens carefully, replies with clarity, and never implies that his own answers are final. He has mastered something rare: the ability to be an expert without becoming an absolutist.

And so the work continues. On the hill above Vancouver, in a lab lined with instruments and shelves of samples, Steven and his team push forward — not loudly, not dramatically, but steadily. Their membranes will soon be cast not only in petri dishes and glass plates, but also into electrolyser stacks the size of furnaces. Their chemistry will meet steel, thermal gradients, and operators in hard hats. Their ideas will be tested under pressures no lab simulation can fully anticipate. That is precisely the point.

AEM technology still faces hurdles: durability thresholds, cost targets, manufacturing challenges. However each time Steven examines a new data plot or lifts a membrane that has lasted longer than any before it, the future feels a little closer. He knows he is not alone. Laboratories across Europe, Asia, and North America are pushing alongside him. The frontier is wide, and many hands are needed.

Yet on this particular hill, in this particular lab, one scientist continues to work patiently at the boundary between promise and proof. His efforts will not single-handedly deliver the hydrogen economy. But without work like his — exacting, inventive, and sustained — that economy may never reach its full potential.

Some revolutions begin with silence, like a hydrogen train gliding into a station. Others begin in places like Steven's lab, where the hum of equipment and the steady focus of a small team quietly shape the foundations of a new energy world. The membranes may be thin, almost fragile to the touch, but the future they represent is anything but.

If there is a front line in the hydrogen economy — or, if one were inclined to imagine a USS Enterprise, boldly going where no one has gone before — it is here: in universities, research institutes, and small start-ups, where chemists, engineers, and materials scientists work patiently at the edge of what is known. They are not chasing spectacle, but progress — testing new ideas, refining old ones, and accumulating the incremental gains that, taken together, make transformation possible.

But before we return to Steven, to AEM membranes, and to the polymer frontier that may yet reshape electrolysis, we need to step back. This is because innova-

tion in hydrogen does not begin with materials or membranes, but with a more fundamental question: how hydrogen is actually made at scale, and why different technological pathways exist in the first place.

To understand where hydrogen technology is going, we must first understand the core conversion processes that sit at its foundation — the electrochemical and thermochemical routes by which electricity, heat, and molecules are transformed into hydrogen.

Electrolysis — The Core Conversion Technologies

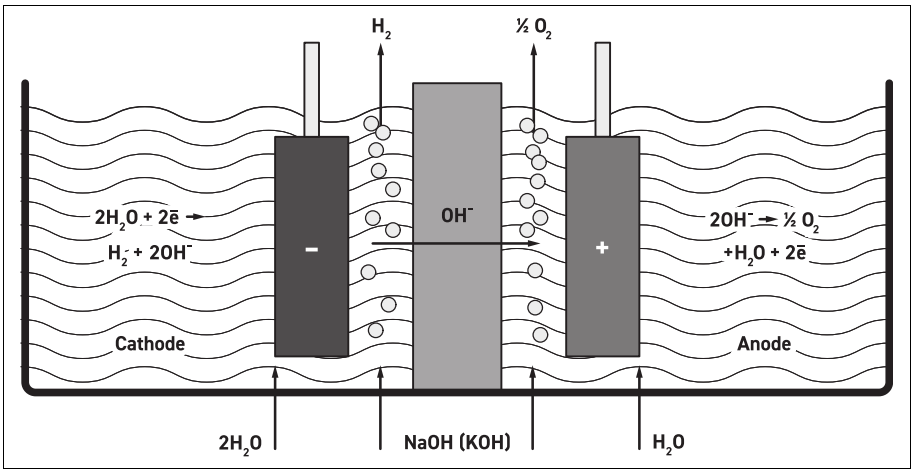
At its simplest, electrolysis is an act of reversal. Instead of releasing energy by combining hydrogen and oxygen, it consumes energy to pull them apart. Water is split into its elemental components using electricity, producing hydrogen at the cathode and oxygen at the anode. The chemistry itself is unremarkable — it has been understood for more than a century. What is new is not the reaction, but the role it is now being asked to play.

Electrolysis is emerging as one of the core conversion technologies of the energy transition. Just as steam engines once converted heat into motion, and internal combustion engines converted chemical energy into mobility, electrolysers convert electricity into molecules — dense, transportable, and versatile carriers of energy. They are the mechanism by which renewable power, nuclear heat, and industrial electricity are translated into fuels, feedstocks, and storage that can move across sectors and borders.

What distinguishes one electrolyser technology from another is not the final product, but the architecture of the system itself: how ions are transported, how current is distributed, how heat is managed, and where losses are tolerated or eliminated. These design choices shape everything that follows — efficiency, durability, material demand, operating flexibility, and ultimately cost. They determine whether an electrolyser behaves like a laboratory instrument, an industrial machine, or a piece of national infrastructure.

Over time, four principal electrolysis pathways have emerged. Each reflects a different optimisation logic, a different answer to the same fundamental challenge: how to turn electricity into hydrogen reliably, repeatedly, and at scale. None is universally superior. Each is suited to particular operating environments and system roles. Together, they form the technological foundation upon which the hydrogen economy is being built.

Alkaline Electrolysis — The Industrial Baseline



Schematic overview of alkaline water electrolysis
A simplified diagram showing an anode and cathode separated by a diaphragm, a liquid KOH electrolyte enabling hydroxide ion (OH^-) transport, electrons flowing through an external circuit, and separate collection of hydrogen and oxygen.

Alkaline electrolysis is the reference technology of hydrogen production. Long before hydrogen became entangled with energy transition narratives, alkaline systems were already operating quietly in refineries, chemical plants, and industrial gas facilities. Their longevity is not accidental. It is the result of a design philosophy that prioritises stability, tolerance, and incremental improvement over optimisation at the edge of performance.

At the heart of an alkaline electrolyser is a liquid electrolyte — most commonly an aqueous solution of potassium hydroxide (KOH). KOH plays a critical role: it enables the transport of hydroxide ions (OH^-) between the cathode and anode, while remaining chemically stable under strongly alkaline conditions. The electrolyte is contained within a porous separator or diaphragm that allows ionic conduction, but prevents the mixing of hydrogen and oxygen gases.

The electrochemical reactions themselves are simple. Water is reduced at the cathode to form hydrogen gas and hydroxide ions. Those hydroxide ions migrate through the electrolyte toward the anode, where they are oxidised to form oxygen gas and release electrons back into the external circuit. The process is mature, well characterised, and forgiving of real-world operating conditions.

This maturity translates into robustness. Alkaline systems tolerate a wider range of water qualities, impurities, and operating fluctuations than many newer technologies. Electrode materials are typically based on nickel or nickel alloys, meaning that alkaline electrolysis does not inherently depend on scarce or geopolitically sensitive precious metals. Some manufacturers are, however, exploring platinum-group-metal coatings and other advanced electrode materials as a means of improving efficiency and performance. Degradation rates are slow, lifetimes are long, and failure modes are well understood.

Strengths

- Low exposure to critical materials, with no inherent requirement for platinum-group metals
- Long operational lifetimes, often measured in decades
- Well-established supply chains and operating experience

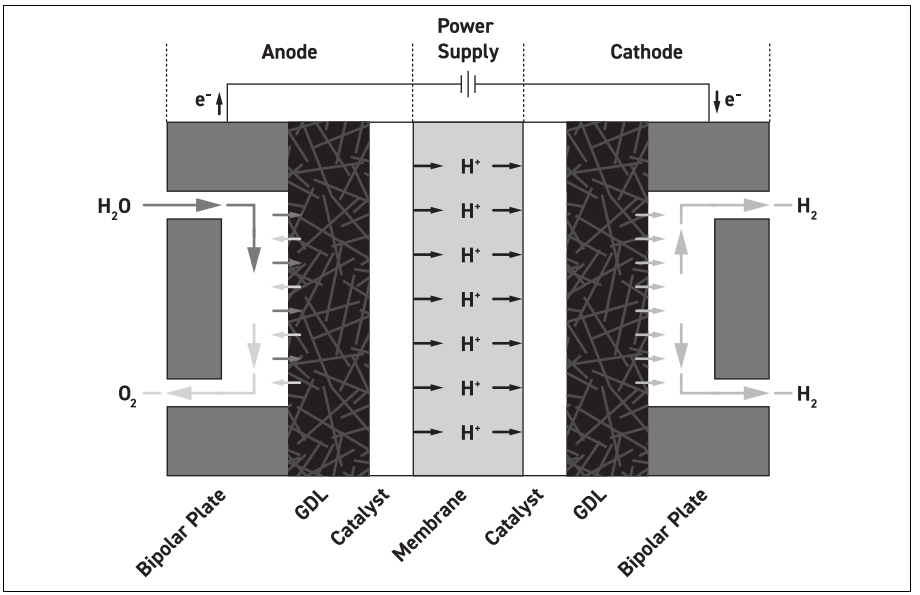
Constraints

- Limited dynamic operation, with slower ramp rates and reduced tolerance for frequent cycling
- Larger physical footprint, driven by lower current densities and liquid electrolyte systems

A simple process overview, however, can obscure where energy is lost. In alkaline electrolysis, losses occur in several places. Ionic losses arise from resistance within the liquid electrolyte and separator. Resistive losses occur in electrodes, current collectors, and external electrical connections. Thermal losses appear as heat generated during operation, which must be managed to prevent efficiency from declining further. None of these losses is unusual, but together they define the practical efficiency limits of the system.

For these reasons, alkaline electrolysis remains well suited to steady-state, industrial-scale hydrogen production, particularly in environments where space is available and flexibility is secondary to reliability. It is not the most compact, nor the most responsive technology. But as an industrial baseline — a proven way of turning electricity into hydrogen day after day, year after year — alkaline electrolysis continues to set the standard against which other pathways are measured.

PEM Electrolysis — Power Density and Flexibility



Schematic overview of PEM water electrolysis
A conceptual diagram showing a solid polymer membrane conducting protons (H^+), with water supplied to the anode, hydrogen produced at the cathode, oxygen at the anode, and electrons flowing through an external circuit. The membrane separates gases, while enabling ionic transport.

In a PEM electrolyser, the solid polymer membrane replaces the liquid electrolyte used in alkaline systems, allowing higher current densities, faster response, and a more compact architecture — at the cost of greater dependence on specialised materials, tighter operating requirements, and more demanding system control.

If alkaline electrolysis is defined by robustness, proton exchange membrane (PEM) electrolysis is defined by intensity. PEM systems are designed to do more in less space, respond faster to changing power conditions, and operate as an active component of a dynamic electricity system rather than as a steady industrial appliance.

At the centre of a PEM electrolyser is a solid polymer membrane, or ionomer, that conducts protons while remaining impermeable to gases and electrons. Water is supplied to the anode side of the membrane, where it is split into oxygen, protons,

and electrons. The protons migrate through the membrane to the cathode, while electrons travel through the external circuit. At the cathode, protons and electrons recombine to form hydrogen gas.

This solid-state architecture eliminates the need for a liquid electrolyte and allows electrodes to be pressed directly against the membrane. The result is short ionic pathways, low internal resistance, and the ability to operate at much higher current densities than alkaline systems. In practical terms, this means more hydrogen production per square centimetre of active area, and a far more compact system overall.

Two parameters dominate PEM operation: voltage and current density.

Voltage reflects how much electrical energy is required to drive the reaction. It includes the theoretical minimum set by thermodynamics, plus additional losses from reaction kinetics, ionic resistance within the membrane, electronic resistance in the electrodes, and mass-transport limitations. Lower voltage at a given output corresponds to higher efficiency — but achieving it requires careful materials design and precise control of operating conditions.

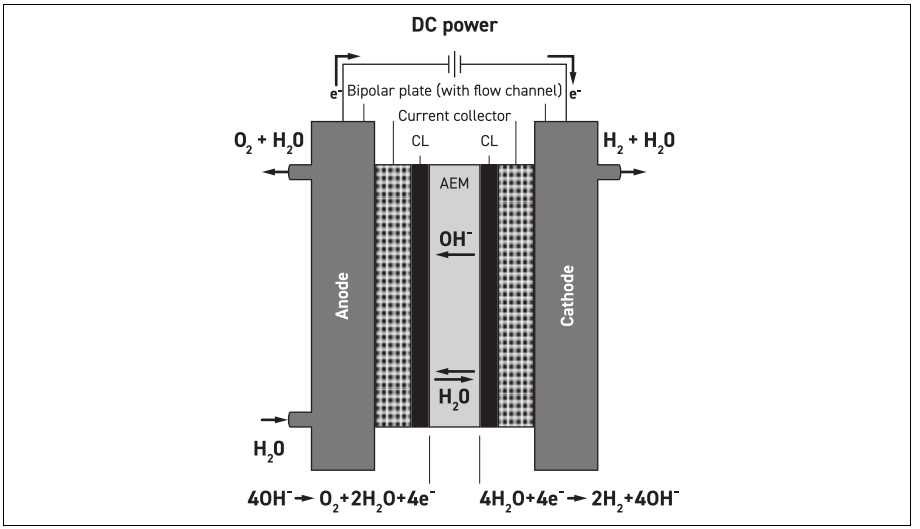
Current density, typically expressed in amps per square centimetre, indicates how hard the electrolyser is being driven. Higher current density increases hydrogen production and reduces system size, but it also accelerates degradation. Membranes, catalysts, and interfaces experience greater mechanical, chemical, and thermal stress. Operating at high current density is therefore a trade-off between capital efficiency and lifetime.

These trade-offs extend into operating pressure. PEM electrolysers can produce hydrogen at atmospheric pressure, moderate pressures around 10 bar, or higher pressures approaching 30 bar. Each choice shifts complexity within the system. Higher operating pressure reduces the energy required for downstream compression but increases mechanical stress on the membrane and stack components. Lower-pressure operation eases stack design and durability concerns, but transfers energy consumption and capital cost to compressors and balance-of-plant equipment.

There is no universally optimal pressure. The right choice depends on how the electrolyser is integrated into the wider system — whether hydrogen is used locally, stored, transported, or further processed. In this sense, PEM electrolysis illustrates a broader truth about hydrogen technologies: performance is inseparable from context.

PEM electrolysers excel where flexibility matters. They ramp quickly, tolerate frequent cycling, and integrate naturally with variable renewable power. But this flexibility comes with greater dependence on specialised materials, tighter operating windows, and more demanding control systems. PEM is not simply a faster version of alkaline electrolysis; it is a different machine, designed for a different role in the energy system.

AEM Electrolysis — The Emerging Hybrid



Schematic overview of AEM water electrolysis
A conceptual diagram showing a solid anion exchange membrane conducting hydroxide ions (OH^-), with water supplied to the cathode, hydrogen produced at the cathode, oxygen at the anode, and electrons flowing through an external circuit. The membrane separates gases while enabling ionic transport, combining alkaline chemistry with a solid-state, PEM-like architecture.

In an AEM electrolyser, the solid membrane replaces the liquid alkaline electrolyte, allowing compact system designs and flexible operation while retaining alkaline chemistry and the potential for non-precious metal catalysts.

Anion exchange membrane (AEM) electrolysis exists because the established pathways — alkaline and PEM — each solve part of the problem, but neither resolves it entirely. Alkaline systems offer durability and low material criticality, but struggle with compactness and dynamic operation. PEM systems deliver

power density and flexibility, but rely on expensive materials and environmentally persistent polymers. AEM electrolysis emerges from the attempt to bridge that gap.

At a fundamental level, AEM electrolysers conduct hydroxide ions through a solid polymer membrane, rather than through a liquid electrolyte as in alkaline systems. This allows AEM architectures to retain alkaline chemistry — and the ability to use non-precious metal catalysts — while adopting PEM-like, solid-state form factors. In principle, the result is a system that is compact, flexible, and materially less constrained.

The motivations behind AEM development are therefore clear. Reducing reliance on precious metals lowers cost exposure and supply-chain risk. Reducing or eliminating PFAS-based materials addresses a growing environmental and regulatory concern. Together, these pressures have pushed materials scientists and system designers toward alternative membrane chemistries and electrode architectures that would have seemed speculative only a decade ago.

Where AEM shows particular promise is in its material optionality. By decoupling high performance from platinum-group catalysts and fluorinated membranes, AEM opens the door to a wider range of polymer backbones, cationic groups, and catalyst formulations. This flexibility creates room for innovation not only in electrochemistry, but also in manufacturing, recycling, and long-term sustainability.

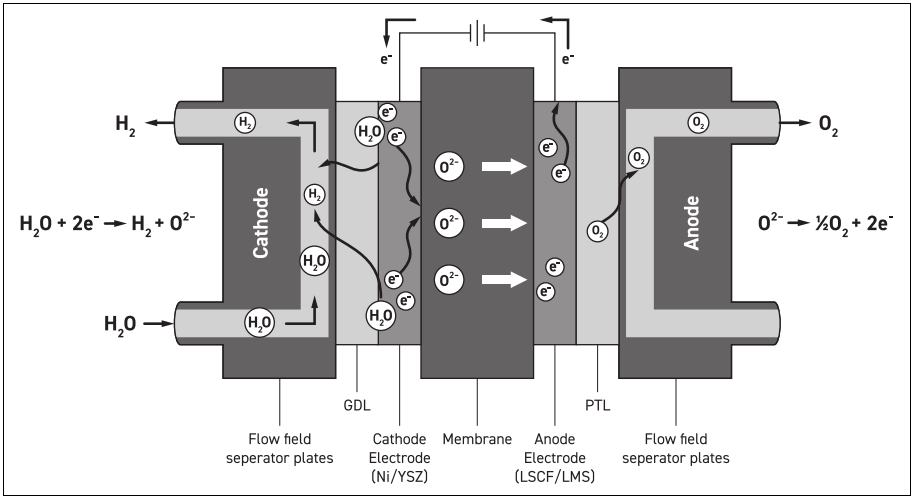
Yet the challenges are significant. Durability remains the primary constraint. Hydroxide-conducting membranes must survive highly alkaline environments while maintaining mechanical integrity and ionic conductivity over long operating lifetimes. Carbonate formation, driven by exposure to carbon dioxide in feed water or ambient air, can degrade performance and complicate system operation. And beyond the laboratory, AEM systems remain comparatively immature, with limited field data and few large-scale deployments.

These challenges do not invalidate the technology — they define its current position. AEM electrolysis should be understood as a technology in transition, moving steadily from academic promise toward industrial relevance. Its progress will depend less on dramatic breakthroughs than on accumulated improvements in membrane chemistry, electrode design, system integration, and operational learning.

For this reason, AEM is best viewed as a promising pathway, not yet a default architecture. It may not replace alkaline or PEM systems wholesale. Instead, it is likely to carve out its own role, particularly in applications where material sus-

tainability, cost pressure, and compact design intersect. As the polymer frontier advances, the balance between these pathways may shift — and AEM may yet become a central pillar of hydrogen production. But that outcome will be earned through performance in the field, not assumed in advance.

SOEC — Efficiency Through Heat



Schematic overview of solid oxide electrolysis
Conceptual diagram showing steam input, oxygen ion (O^{2-}) transport through a solid ceramic electrolyte at high temperature, hydrogen and oxygen production at separate electrodes, and integration with an external heat source.

Solid oxide electrolysis cells (SOECs) approach hydrogen production from a fundamentally different direction. Rather than optimising materials and architecture to handle higher current densities or greater flexibility, SOECs shift the problem into another domain entirely: temperature.

SOECs operate at high temperatures, typically between 700 and 850 °C. At these temperatures, the energy required to split water is partly supplied as heat rather than electricity. From a thermodynamic perspective, this is highly attractive. Electrical efficiency increases, and when suitable heat sources are available, overall system efficiency can exceed that of low-temperature electrolysis by a considerable margin.

In an SOEC, water is introduced as steam. Oxygen ions migrate through a solid ceramic electrolyte — commonly yttria-stabilised zirconia — while electrons flow through the external circuit. Hydrogen is produced at one electrode, oxygen at the other, in a process that is electrochemically elegant and energetically efficient.

But efficiency here is conditional. High-temperature operation imposes stringent materials requirements and introduces sensitivities that low-temperature systems largely avoid. Thermal expansion, mechanical stress, and chemical compatibility become dominant design constraints. Cycling — particularly frequent start-stop operation — accelerates degradation. SOEC systems therefore favour steady-state operation, where temperatures and loads change slowly, if at all.

For this reason, SOECs are best understood as integrated system technologies rather than modular components. They align naturally with environments where high-grade heat is already present: nuclear power plants, steel mills, cement kilns, and other industrial processes with continuous thermal output. In such contexts, SOECs can convert electricity and heat into hydrogen with exceptional efficiency, turning what would otherwise be waste energy into a valuable product.

That same integration requirement limits their flexibility. SOECs are poorly suited to highly variable power input, rapid cycling, or distributed deployment. They demand careful thermal management, sophisticated balance-of-plant systems, and operating philosophies that prioritise stability over responsiveness.

SOEC technology therefore illustrates a recurring theme in the hydrogen landscape: there is no free efficiency. Gains achieved in one domain are paid for in another — through complexity, sensitivity, and tighter coupling to surrounding infrastructure. Where those trade-offs align with the system context, SOECs can be compelling. Where they do not, their advantages quickly erode.

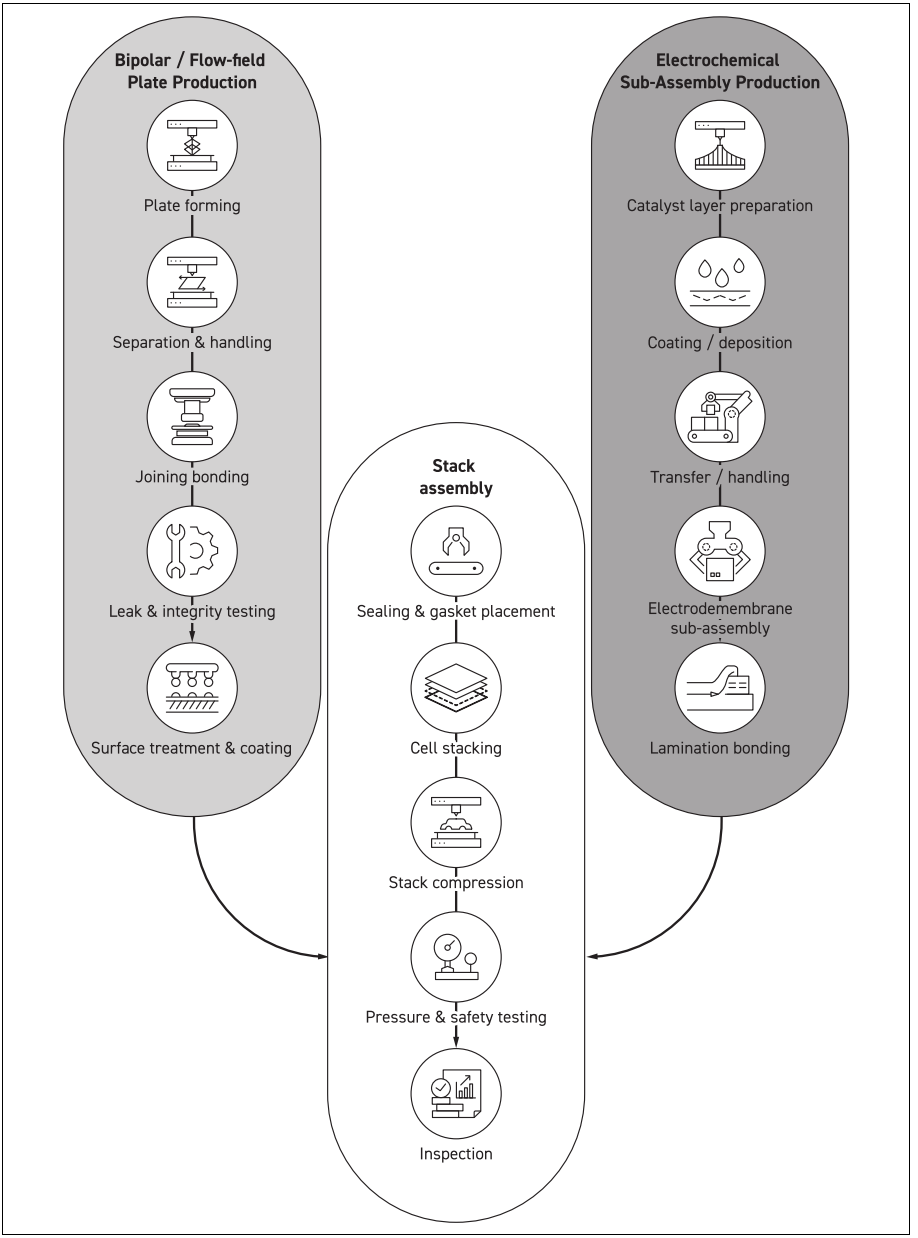
In the broader electrolysis landscape, SOECs do not compete directly with alkaline, PEM, or AEM systems. Instead, they occupy a distinct niche, one defined by integration rather than modularity. They remind us that hydrogen production is not a single technology problem, but a portfolio of solutions, each shaped by the environments in which it is expected to operate.

If electrolysis is the chemistry of conversion, the next bottleneck is industrial: how those cells become products that can be built at volume. In hydrogen, the factory is part of the technology: it determines yield, cost, reliability, and how fast learning curves actually arrive.

Stack Design and Manufacturing Reality

Electrolysis technologies are often discussed as if performance emerges directly from chemistry. In practice, performance is just as often determined by how that chemistry is assembled into hardware, and how consistently that hardware can be produced. The transition from laboratory cell to industrial stack is not a scaling exercise in the abstract; it is a manufacturing problem with tight tolerances, unforgiving interfaces, and little patience for theoretical elegance.

An electrolyser stack is, at its core, a repeated structure. Individual electrochemical cells are layered together, separated by bipolar plates, sealed by gaskets, compressed into a mechanically stable unit, and connected electrically and hydraulically so that current, water, gases, and heat all move where they are supposed to. Each interface is a potential point of loss, degradation, or failure. Each repetition magnifies both strengths and weaknesses.



Simplified electrolyser stack build process

A step-by-step schematic illustrating how an electrolyser stack is assembled, from individual components to a compressed, functional unit.

In practice, stack performance emerges not from any single component, but from the precision with which hundreds of identical layers are assembled, sealed, and compressed into a repeatable structure.

The stack must perform several functions simultaneously. It must maintain uniform current distribution across large active areas. It must keep hydrogen and oxygen safely separated under pressure. It must manage heat generated during operation. Lastly, it must do all of this while tolerating thermal expansion, mechanical stress, and chemical attack over thousands of operating hours. None of these requirements is optional.

From a manufacturing perspective, this creates immediate challenges. Tolerances matter. Small deviations in plate flatness, gasket thickness, or compression force can lead to uneven current density, local hot spots, accelerated degradation, or gas crossover. What works in a single-cell test rig may fail when replicated hundreds of times in a commercial stack.

Yield therefore becomes a defining constraint. Unlike many energy technologies, electrolysers do not scale by simply making components larger. They scale by making more units, more consistently. This places electrolysers closer to advanced manufacturing sectors — such as fuel cells, batteries, or semiconductors — than to traditional energy infrastructure. Learning curves are driven not just by materials cost, but by automation, quality control, and process repeatability.

Stack design choices also lock in operational behaviour. Decisions about cell size, number of cells, flow-field geometry, and compression strategy influence pressure drop, thermal gradients, and degradation rates. Once a stack architecture is established, changing it is slow and expensive. This is one reason why early design decisions exert such long influence over cost and performance trajectories.

Losses accumulate here as well. Electrical resistance in bipolar plates and current collectors adds to system voltage. Non-uniform flow distribution increases pumping power and reduces effective utilisation of active area. Thermal gradients create inefficiencies and shorten component life. None of these losses is dramatic in isolation, but together they shape real-world efficiency far more than headline figures suggest.

Perhaps most importantly, stack manufacturing exposes a truth that is easy to overlook in policy and strategy discussions: electrolysers are not just technologies; they are products. They must be designed to be built, tested, shipped,

installed, operated, serviced, and eventually replaced. A design that cannot be manufactured reliably will not scale, regardless of its theoretical advantages.

As hydrogen moves from pilot projects to industrial deployment, stack manufacturing becomes a bottleneck not because it is unsolved, but because it is exacting. Progress here is measured not in breakthroughs, but in incremental gains: higher yields, tighter tolerances, better quality control, and designs that accept the realities of mass production. This is where hydrogen technologies begin to leave the laboratory behind and enter the discipline of industry.

For much of the past decade, this manufacturing reality was masked by low volumes. Most electrolyser original equipment manufacturers were producing only a handful of stacks per year, often assembled through highly manual processes. Cells were aligned by hand, gaskets placed individually, compression adjusted manually, and quality checks performed through inspection rather than instrumentation. These methods were adequate for pilot projects and early demonstrations, but they do not translate into industrial scale. They rely heavily on skilled labour, tolerate variability, and embed cost structures that cannot survive high-volume deployment.

Over the last one to five years, that picture has begun to change. As deployment expectations have risen, electrolyser OEMs have started taking the first steps toward automation. Assembly lines are becoming more structured. Subprocesses, such as plate handling, membrane placement, and compression are increasingly mechanised. Inline testing and digital quality control are beginning to replace manual inspection. The pace and depth of this transition vary widely between manufacturers, reflecting different starting points, capital availability, and technology choices.

This shift is still incomplete. The market has not yet demanded sustained, high-volume output, and there remains unused manufacturing capacity across much of the sector. Full automation therefore has not been economically justified in most cases. However, the direction of travel is clear. If electrolyser production is to move from hundreds of units to thousands — and eventually tens of thousands — manual assembly will become a hard constraint rather than a temporary inconvenience.

The next phase of scale-up will be defined less by chemistry than by factories. As demand firms and volumes rise, manufacturers will be forced to pursue deeper automation, tighter process control, and higher capital intensity in production. Those who succeed will not necessarily be the ones with the most advanced electrochemistry, but those who can translate designs into repeatable, automated

manufacturing systems. In the hydrogen economy, industrial learning will increasingly be written not in laboratory notebooks, but in factory layouts.

One important consequence of this manufacturing reality is that electrolyser stacks are not an industrial novelty. In structural terms, the build process for power stacks closely resembles that of other electrochemical energy devices — most notably lithium-ion batteries and fuel cells. Layered architectures, repeated cell units, tight tolerances, sealing, compression, and the need for uniform current distribution are common challenges across all three.

This matters because it means the hydrogen industry is not reinventing the wheel. The expertise required to design and operate high-throughput electrochemical manufacturing lines already exists. Decades of learning from battery gigafactories — in automation, inline quality control, materials handling, and yield optimisation — are directly applicable to electrolyser production. While electrolyser stacks introduce their own specifics, particularly around gas management, corrosion, and pressure containment, these differences sit on top of a familiar industrial foundation.

In practical terms, this lowers execution risk. It shortens learning curves. It allows electrolyser OEMs to draw on an existing workforce, established suppliers, and proven factory design principles. The constraint is therefore not the absence of manufacturing knowledge, but the timing and confidence to deploy it at scale. As demand becomes more predictable, the transition from semi-automated assembly to fully industrialised production will increasingly resemble a continuation of the battery manufacturing story, rather than the invention of an entirely new one.

Monitoring, Control, and Degradation

In the early stages of any industrial technology, the biggest problem is not failure — it is ignorance. Electrolysers degrade slowly, often invisibly. Performance drifts by fractions of a percent, a temperature distribution changes across the stack, a seal begins to creep, a local hotspot forms and disappears again. By the time output falls far enough to notice, the real damage has usually already been done.

This is why monitoring matters. Not as an optional layer of sophistication, but as the mechanism that turns an electrolyser from a prototype into an asset. The transition from manual assembly to industrial manufacturing only accelerates this need. As stacks become higher power density and more tightly optimised, the operating window narrows. The tolerance for uncertainty shrinks. If electrolysers

are to operate reliably for years, under variable loads and real-world conditions, they need a way to observe themselves.

Cell Voltage Measurement

A widely used monitoring approach in modern electrolysers is cell voltage measurement (CVM): measuring voltage not only across the full stack, but across individual cells or defined cell groups. Instead of treating the stack as a single black box, CVM reveals internal distribution — which regions are healthy, which are drifting, and which are beginning to fail.

Voltage is not just an electrical number. In an electrolyser, it is a proxy for multiple physical realities: reaction kinetics, ionic resistance, contact resistance, hydration state, temperature distribution, gas transport, and degradation at electrodes and interfaces. When a cell voltage rises relative to its neighbours, something inside that cell has changed. CVM makes those changes visible early — often long before they show up in stack-level performance.

In practice, CVM requires an integrated measurement architecture, robust acquisition electronics, and software capable of operating continuously in a noisy, high-current industrial environment. In practice, this provides internal visibility that would otherwise be unavailable.

From Monitoring to Decisions

The value of CVM is not the data itself, but what it enables. It allows operators to move from reactive maintenance to informed intervention. Outliers can be identified before they become shutdown events. Patterns across the stack can indicate maldistribution of flow, uneven cooling, or compression issues. It can also support safer operation by providing early warning of conditions associated with increased gas crossover risk or local overheating.

In practice, this shifts monitoring from passive observation to active diagnosis. It enables earlier identification of abnormal cells before degradation becomes visible at stack level and supports more informed operational adjustments — including changes to load, temperature, or water management — that reduce local stress and slow deterioration. It allows maintenance to move from reactive intervention to targeted diagnosis, helping to distinguish between electrical, thermal, mechanical, and chemical causes of performance drift. It also plays a role during commissioning and service, where uniform cell behaviour can be verified after installation or refurbishment. Beyond individual assets, the same data provides feedback into design and manufacturing processes, exposing recurring patterns

that point to tolerances, compression uniformity, interface quality, or material variability.

CVM also links factory conditions with field performance. By exposing internal variation, it provides feedback on manufacturing quality. Consistent voltage gradients or recurring outliers can often be traced back to tolerances, compression uniformity, interface contact quality, gasket behaviour, or material variability.

This becomes increasingly important as the sector transitions from manual assembly toward automation. Automation reduces variability, but it also makes systematic errors repeatable. CVM helps detect whether process changes improve uniformity or introduce new bias, accelerating industrial learning and reducing lifetime risk.

CVM as Part of Wider Condition Monitoring

CVM is most powerful when it becomes one layer within a wider condition monitoring approach. Electrolysers are multi-physics machines: electrical, thermal, hydraulic, and mechanical domains interact continuously. A mature monitoring solution therefore combines CVM with temperature, pressure and differential pressure, flow stability indicators, gas purity or crossover signals, and electrical integrity checks. CVM provides internal electrochemical visibility; surrounding sensors explain why the voltage is changing

Monitoring does not need to be “advanced” to be valuable. Even simple safeguards — derating when certain regions drift, tightening operating limits when gradients increase — can materially extend life and reduce unplanned downtime.

Electrolysers will degrade; that is unavoidable. What determines economic success is whether degradation is understood and managed, or whether it emerges as surprise failure. CVM is one of the clearest examples of the industry shifting from “build and hope” toward “build, observe, learn, and improve” — a shift that is as much about industrialisation as it is about technology.

Materials, Ionomers, and the PFAS Question

If electrolysis technologies define how hydrogen is produced, materials define whether that production can scale with credibility. Few issues illustrate this more clearly than the growing scrutiny of PFAS-based materials within the hydrogen value chain.

PFAS — per- and polyfluoroalkyl substances — are a broad family of fluorinated compounds valued for their chemical stability, thermal resistance, and durability. These properties are precisely why they have been widely adopted in demanding industrial applications, including membranes, ionomers, seals, and coatings used in PEM electrolysers and fuel cells. In electrochemical systems, PFAS-based polymers have delivered decades of reliable performance under harsh conditions where few alternatives could survive.

However, those same properties have also created a problem. PFAS compounds are environmentally persistent, resistant to natural degradation, and increasingly detected across water systems, ecosystems, and human tissue. As production volumes rise, what once appeared as a contained industrial material challenge becomes a broader societal one. Regulatory scrutiny has intensified accordingly, particularly in Europe, where restrictions and bans are now being proposed or implemented across entire classes of PFAS compounds.

This is not a debate about intent or past decisions. PFAS materials were adopted because they worked — and because, for a long time, there were no viable substitutes. The question now is different. As hydrogen moves from niche deployments toward industrial scale, does the performance advantage of PFAS still outweigh its long-term environmental and regulatory risk?

Answering that question requires moving beyond binary thinking. The choice is not between “PFAS everywhere” and “PFAS nowhere.” It is about recognising that scale changes the equation. A few kilograms of material used in pilot systems pose a very different risk profile than thousands of tonnes embedded across gigawatt-scale infrastructure. What was once a reasonable engineering compromise becomes a strategic decision with long-term consequences.

This is where materials science re-enters the story — and where innovation looks very different from the popular image of sudden breakthroughs. The development of new ionomers, membranes, and polymer architectures is slow, iterative, and often frustrating. Progress is measured in incremental gains in conductivity, durability, chemical stability, and manufacturability. This is the kind of innovation that unfolds over decades, not funding cycles.

The work of Steven Holdcroft and others in the polymer frontier exemplifies this long-cycle innovation. The pursuit of PFAS-free or PFAS-reduced ionomers has required deep rethinking of polymer backbones, cation stability, water management, and degradation pathways. Early candidates failed. Many still do. But each failure narrowed the search space, gradually revealing architectures that could

survive alkaline environments, sustain high current densities, and be manufactured at scale.

That same spirit of persistence can be found today in a new generation of smaller companies pushing at the boundaries of what is possible. One such example is Ionysis, a young startup based in Freiburg, Germany. With a team numbering only a few dozen, Ionysis operates far from the scale of multinational OEMs — but very close to the cutting edge of materials innovation.

Ionysis is focused on developing PFAS-free ionomer and membrane solutions for electrochemical applications, including hydrogen technologies. Their work is not framed as a rejection of existing systems, but as an attempt to expand the option space. By exploring alternative polymer chemistries and production routes, they are testing whether performance traditionally associated with fluorinated materials can be approached — or eventually matched — using architectures that carry a lighter environmental footprint.

What makes companies like Ionysis significant is not that they have solved the PFAS challenge outright. They have not. What matters is that they demonstrate a credible pathway forward, grounded in chemistry rather than rhetoric. Their work acknowledges the trade-offs openly: PFAS-free materials today may sacrifice some durability, require tighter operating windows, or demand new manufacturing approaches. Learning curves are real, and costs will not fall overnight.

But this is precisely how responsible industrial transitions unfold. New materials are proven first in controlled environments, then in limited deployments, and only later at scale. Feedback from the field informs the next iteration. Over time, performance gaps narrow, confidence grows, and what once seemed risky becomes routine.

Seen in this light, the PFAS question is not an obstacle to the hydrogen economy — it is a test of its maturity. Does the industry wait for regulation to force change, or does it invest early in alternatives, accepting near-term complexity in exchange for long-term resilience? Does it optimise for today's balance sheet, or for infrastructure that society will still trust decades from now?

Hydrogen's promise has always rested on more than efficiency or cost. It rests on credibility. The choices being made now — in materials labs, pilot lines, and small startups pushing against the limits of polymer chemistry — will determine whether hydrogen industrialises with foresight rather than hindsight. In that sense, the quiet work being done in places like Vancouver and Freiburg may prove just as consequential as the gigawatt-scale projects that dominate the headlines.

Blue Hydrogen Technologies — SMR and ATR

If electrolysis represents the future-facing edge of hydrogen production, reforming technologies represent its present reality. Today, the global hydrogen economy exists largely because steam reforming works — reliably, at scale, and at cost levels that industry has been willing to accept for decades. Refineries, fertiliser plants, and chemical complexes around the world depend on reforming-based hydrogen as a foundational input, not a speculative alternative.

For this reason alone, any serious discussion of hydrogen's role in the energy transition must engage directly with steam methane reforming (SMR) and autothermal reforming (ATR). These are not marginal technologies. They are among the most mature, industrialised chemical processes in operation today. The question facing the transition is therefore not whether reforming should exist, but whether it can be adapted to meet a decarbonising world.

Blue hydrogen is the result of that adaptation. By integrating carbon capture and storage into reforming processes, blue hydrogen seeks to retain the strengths of existing hydrogen production — scale, reliability, and integration with industrial demand — while sharply reducing associated emissions. Whether it succeeds depends not on colour labels, but on process design, capture efficiency, system boundaries, and governance.

Two reforming pathways dominate the discussion: Steam Methane Reforming (SMR) and Autothermal Reforming (ATR).

Steam Methane Reforming (SMR) — The Industrial Workhorse

Steam methane reforming has been the backbone of industrial hydrogen production for decades. Methane from natural gas is mixed with steam and passed over a nickel-based catalyst at high temperature. The primary reaction produces hydrogen and carbon monoxide, followed by a water–gas shift reaction that converts carbon monoxide into additional hydrogen and carbon dioxide.

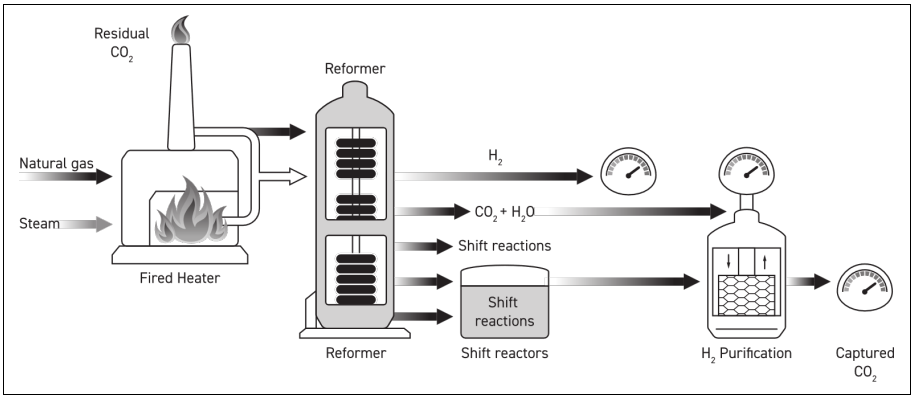
The chemistry is well understood, the equipment is proven, and the economics are favourable. This is why SMR became the default. Yet from a decarbonisation perspective, its weakness is structural: carbon dioxide is produced in multiple places, not all of which are easily captured.

In a conventional SMR system, emissions arise from:

- the reforming reactions themselves
- the combustion of natural gas to supply process heat

- downstream purification and auxiliary systems

Carbon capture can be applied to the process gas stream, where CO_2 concentrations are relatively high and capture is comparatively efficient. However, emissions associated with furnace combustion are more diffuse and harder to capture. As a result, SMR-based blue hydrogen typically achieves partial capture, often in the range of 60–90%, depending on system design and boundary definition.



Simplified SMR process with carbon capture
A schematic showing natural gas and steam entering the reformer, hydrogen production via reforming and shift reactions, CO_2 generation at multiple points, and capture applied primarily to the process gas stream, with residual emissions from the fired heater.

SMR with CCS can significantly reduce emissions relative to grey hydrogen, but it does not eliminate them entirely. Its role, therefore, is conditional: valuable where infrastructure already exists and rapid emissions reduction is needed but constrained by its inherent capture limitations.

Autothermal Reforming (ATR) — Designed for Higher Capture

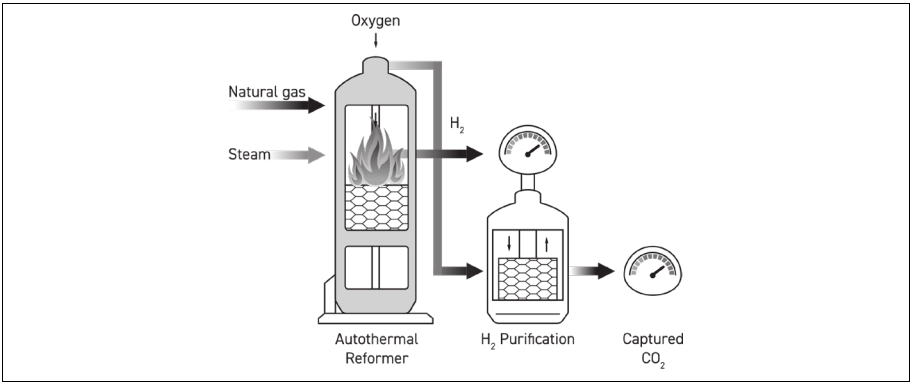
Autothermal reforming approaches the same chemistry from a different configuration. Instead of supplying heat externally via a fired furnace, ATR combines partial oxidation and steam reforming in a single reactor. Oxygen is introduced directly, allowing the reaction to be thermally self-sustaining.

This architectural choice has important consequences. Since the reforming reactions and heat generation occur within a single vessel, the resulting CO_2 stream is more concentrated and easier to capture. In principle, ATR systems can achieve

very high capture rates, often exceeding 95%, when integrated with appropriate CCS infrastructure.

The trade-off is complexity. ATR requires:

- an oxygen supply, typically from an air separation unit
- tighter integration between reaction, heat management, and capture
- higher capital intensity and more sophisticated control



Simplified ATR process with integrated carbon capture
A schematic showing natural gas, steam, and oxygen entering a single reforming reactor, hydrogen production, a concentrated CO_2 stream routed to capture, and hydrogen purification downstream.

ATR is therefore less flexible as a retrofit option, but more attractive for new-build plants designed explicitly for low-carbon operation. Where capture efficiency is paramount and scale justifies complexity, ATR becomes compelling.

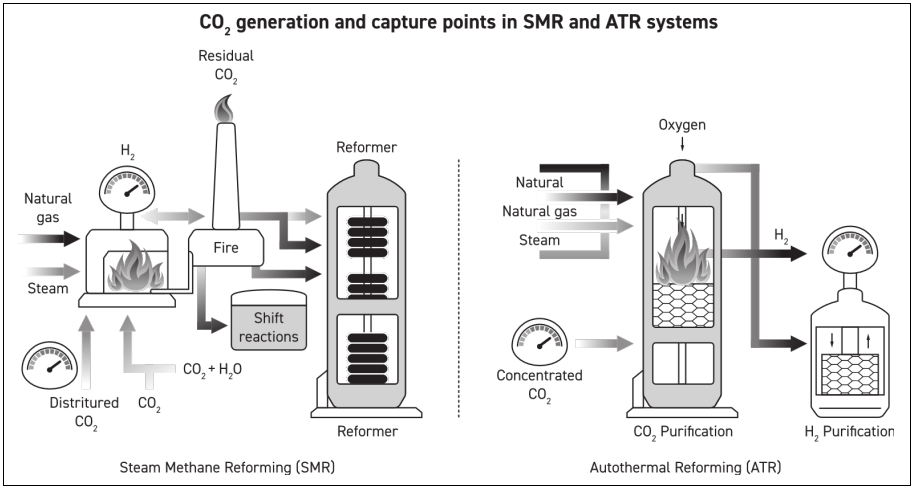
Carbon Capture and Storage — The Defining System of Blue Hydrogen

In blue hydrogen, carbon capture and storage is not an optional enhancement. It is the defining system. In practice, blue hydrogen is less a hydrogen technology than a CO_2 logistics technology that happens to produce hydrogen. Reforming technologies have produced hydrogen reliably for decades; what distinguishes blue hydrogen from its grey predecessor is not how hydrogen is made, but what happens to the carbon that is released in the process. Without effective capture,

transport, and permanent storage, blue hydrogen simply reverts to conventional fossil-based production under a different name.

Carbon capture in reforming systems is shaped as much by process architecture as by separation technology. In steam methane reforming, carbon dioxide is generated at several points: within the reforming and shift reactions themselves, and in the combustion systems that supply high-temperature heat. The CO_2 present in the shifted process gas is relatively concentrated and well suited to capture using established technologies. This is where most SMR-based capture systems focus. By contrast, emissions from the fired heater are more diffuse, and unless specific design measures are taken, they are often left uncaptured. The result is that SMR-based blue hydrogen typically achieves partial capture across the full system boundary, even if capture efficiency at individual points is high.

Autothermal reforming addresses this challenge through design. By integrating heat generation and reforming reactions within a single reactor, ATR produces a more concentrated carbon dioxide stream, simplifying separation and enabling significantly higher overall capture rates. This structural advantage is the primary reason ATR is frequently favoured for new blue hydrogen projects targeting capture rates above ninety percent. The trade-off is increased complexity: oxygen supply, tighter process integration, and higher capital intensity. In effect, ATR represents a reforming pathway designed from the outset around the requirements of carbon capture.



CO_2 generation and capture points in SMR and ATR systems
Comparative schematic illustrating distributed emission sources in SMR versus concentrated capture streams in ATR.

Capturing carbon dioxide is only the first step. While much of the current focus is on permanent storage, carbon capture also plays an important role as a feedstock enabler. In pathways such as e-fuels — including synthetic aviation fuels (e-SAF) — captured CO_2 becomes a valuable input, combined with low-carbon hydrogen to produce liquid hydrocarbons. In this context, carbon is not simply a waste to be disposed of, but a resource to be reused. The distinction matters: storage addresses emissions that must be permanently removed, while utilisation supports the creation of drop-in fuels for sectors where direct electrification remains difficult. Both rely on the same fundamental capability — the ability to capture carbon dioxide reliably and at scale.

Once separated, CO_2 must be dried, compressed, and transported to a suitable storage site. This introduces a set of constraints that are often invisible in high-level discussions, but decisive in practice. Blue hydrogen projects are therefore closely tied to geography. They tend to emerge in regions with existing pipeline infrastructure, proximity to geological storage formations, and industrial clusters capable of sharing transport networks. Where these conditions do not exist, blue hydrogen becomes far harder — and far more expensive — to justify.

Storage itself is not a one-time action, but a long-term commitment. Injected carbon dioxide must remain isolated from the atmosphere for decades to centuries. This requires carefully characterised geological formations, continuous monitor-

ing, and regulatory oversight. Pressure behaviour, seismic response, and potential leakage pathways must all be tracked over time. In this sense, carbon storage is not simply a technical exercise but a societal one. Public confidence in blue hydrogen depends on trust that captured carbon remains permanently stored, not just temporarily displaced.

This is why monitoring, verification, and transparent accounting are central to the credibility of blue hydrogen. Capture rates must be measured across clearly defined system boundaries. Transport losses must be understood. Storage integrity must be demonstrated continuously, not assumed. Upstream methane leakage must be included in lifecycle assessments. Without this discipline, claims of low-carbon performance cannot be substantiated.

Here, blue hydrogen converges with a theme that runs throughout this chapter: observability creates trust. Just as electrolysers require insight into stack behaviour over time, carbon capture systems require continuous visibility into performance and containment. In both cases, credibility is earned through data, not declarations.

It is tempting to reduce carbon capture to a single headline figure — ninety percent, ninety-five percent — but such numbers can obscure more than they reveal. What matters is not the capture rate at a single point in the process, but the emissions profile of the entire system once all inputs, outputs, and losses are accounted for. Blue hydrogen succeeds or fails not on slogans, but on whether carbon capture functions as a coherent, verifiable system.

For this reason, the future of blue hydrogen is inseparable from the future of CCS itself. If carbon capture can be deployed at scale with high reliability, robust monitoring, and public trust, blue hydrogen can play a meaningful transitional role — reducing emissions rapidly while renewable-based hydrogen continues to scale. If it cannot, blue hydrogen loses the very characteristic that defines it.

In that sense, carbon capture is not an accessory to blue hydrogen. It is the technology on which the entire pathway stands or falls.

Transitional, Not Timeless

It is tempting to treat blue hydrogen as either a distraction or a silver bullet. Neither view holds up under technical scrutiny. Blue hydrogen is best understood as transitional and context dependent.

It can deliver meaningful emissions reductions quickly, particularly in industrial clusters where hydrogen demand already exists. It can support early market development, infrastructure build-out, and learning-by-doing. Howeverl it is not a permanent substitute for hydrogen produced without fossil feedstocks.

The long-term trajectory of the hydrogen economy will be shaped by declining renewable electricity costs, improved electrolysis technologies, and tightening environmental constraints. In that future, the role of blue hydrogen may shrink — but that does not invalidate its near-term contribution.

What matters is clarity. Blue hydrogen is not "clean" by default, nor is it inherently flawed. It is a system whose value depends on design choices, capture performance, and governance. Treated honestly, it can serve as a bridge. Treated casually, it risks becoming a detour.

Plasma Pyrolysis — Hydrogen Without CO_2

One emerging hydrogen pathway sits slightly outside the usual colour debate because it reframes the carbon question entirely. Instead of producing hydrogen through electrolysis or reforming — both of which ultimately deal in oxygen chemistry and carbon dioxide — plasma-based methane pyrolysis aims to split methane directly into hydrogen and solid carbon.

The underlying principle is simple. Methane (CH_4) contains hydrogen and carbon. If the molecule can be cracked without introducing oxygen, the reaction pathway avoids forming CO_2 in the first place. In its idealised form, methane pyrolysis produces hydrogen (H_2) and solid carbon (C). This makes the process conceptually different from SMR and ATR, where carbon ends up primarily as CO_2 and must then be captured, compressed, transported, and stored. Here, carbon is not a waste gas stream — it is a material output.

In practice, achieving this split reliably requires very high energy densities and precise control. Plasma pyrolysis uses an electrical plasma — an ionised gas with extremely high local temperatures — to provide the energy needed to break methane's molecular bonds rapidly. Companies such as HiiROC in the UK are exploring this approach, sometimes describing it as "plasma electrolysis" because electricity drives the conversion, even though the chemistry differs from water electrolysis.

The resulting solid carbon is often described as carbon black (or related forms of solid carbon). Carbon black is already an industrial commodity used widely in tyres and rubber products, pigments and inks, and certain polymer applications.

Depending on purity and morphology, solid carbon can also have higher-value potential in specialty materials. This is where the pathway's credibility is won or lost: the economics improve dramatically if the carbon product can be sold into stable markets at consistent quality but weaken if the carbon becomes a disposal problem.

As with all hydrogen pathways, system boundaries still matter. Methane pyrolysis does not eliminate upstream emissions associated with natural gas production and transport, and its climate value depends on methane leakage rates and the carbon intensity of the electricity driving the plasma. But it offers an intriguing alternative design logic: rather than building a CO_2 capture-and-storage system, it converts carbon into a storable solid from the outset.

For now, plasma pyrolysis remains niche — early-stage, scaling-limited, and highly dependent on feedstock, electricity price, and product markets. Yet it illustrates a wider point that runs through this chapter: hydrogen will not be built through one pathway, but through multiple conversion routes, each trading one constraint for another. This is less a new "colour" of hydrogen than a different engineering bargain — one that may become valuable in the right contexts.

5. Operating Reality and Industrial Readiness

"Plans are nothing; planning is everything." — Dwight D. Eisenhower

If the previous chapter explored what hydrogen technologies can do, this chapter addresses what they actually do once they leave the laboratory, the pilot plant, or the slide deck.

Between technical possibility and industrial reality lies a difficult middle ground. This is where systems encounter losses, materials age, tolerances tighten, operators intervene, financiers ask uncomfortable questions, and designs are forced to justify themselves over years rather than hours. It is here — not at the frontier of innovation — that many promising technologies stumble.

Hydrogen is no exception. Electrolysers do not operate in isolation; they operate under pressure, under variable loads, within constrained grids, alongside compressors, storage systems, pipelines, ships, safety systems, and standards bodies. Small design decisions made early — regarding operating pressure, loading philosophy, monitoring depth, manufacturability, or interface definition — quietly compound over time, shaping cost, reliability, and trust.

This chapter focuses on those compounding realities. It examines how hydrogen technologies behave when scaled, manufactured, integrated, monitored, certified, and operated as infrastructure rather than experiments. It explores where losses accumulate, why operating philosophy matters as much as chemistry, how modularity and standardisation determine deployment speed, and why observability is the difference between learning systems and fragile ones.

Industrial readiness is not achieved by perfect technology. It is achieved when imperfect technologies are assembled into coherent systems that can be built repeatedly, operated predictably, financed confidently, and improved over time.

That transition — from machine to infrastructure — is where the hydrogen economy will ultimately succeed or fail.

Hydrogen Handling, Losses, and Operational Reality

Up to this point, the discussion of hydrogen has focused on how it is produced — through electrolysis, reforming, and increasingly sophisticated material and process design. But production is only the beginning. Once hydrogen leaves the plant boundary, it enters the realm of operational reality, where physical behaviour, losses, and system inefficiencies quietly shape economics and performance.

Hydrogen is an unusually demanding molecule to handle. Its low molecular weight, high diffusivity, and wide flammability range require careful engineering at every stage of the value chain. Compression, storage, transport, and reconversion all impose energy penalties. Individually, these losses may appear modest. Collectively, they compound rapidly.

Electrical losses occur first, as power is converted into chemical energy during electrolysis or consumed during reforming and capture processes. Additional losses arise during compression, where hydrogen must often be pressurised for storage, transport, or downstream use. Thermal losses follow in cooling systems, heat exchangers, and — where relevant — liquefaction. Transport introduces further inefficiencies through pressure drops, auxiliary energy use, and inevitable losses associated with containment.

These trade-offs become particularly apparent when considering how hydrogen is stored. Unlike natural gas, hydrogen has a very low volumetric energy density under ambient conditions. Simply put, a large amount of space is required to store a useful quantity of energy. As a result, hydrogen is typically either compressed, liquefied, or converted into another carrier molecule.

Compression is often the first step. For industrial storage, pressures of tens of bar may be sufficient. For mobility applications, however, pressures typically rise to 350 bar or even 700 bar. This significantly reduces the storage volume required, but it comes at a cost. Compressors consume energy, generate heat, require maintenance, and add complexity to the overall system. The higher the pressure, the greater these penalties become.

At some point, the industry begins asking a different question: rather than continuing to compress hydrogen ever further, would it be more efficient to liquefy it instead?

Liquefaction dramatically increases volumetric energy density and can make long-distance transport more practical. However, it introduces a different set of challenges. Cooling hydrogen to approximately -253°C requires substantial energy, sophisticated cryogenic equipment, and ongoing management of heat ingress and boil-off losses. The energy penalty associated with liquefaction is significant and often exceeds the energy required for compression alone.

The choice between compression and liquefaction is therefore not a matter of identifying a superior technology, but of selecting the least costly compromise for a given application. Compression tends to dominate shorter-distance and smaller-scale applications. Liquefaction becomes increasingly attractive where energy density and transport distance become critical. Both approaches involve losses. The engineering challenge lies in determining which losses are most acceptable within the wider system.

These effects become particularly visible in cryogenic storage systems. Even well-designed hydrogen storage is not passive. Hydrogen gradually absorbs heat from its surroundings and evaporates, producing boil-off gas (BOG). Rather than venting this gas — which would waste energy and raise safety concerns — industrial systems rely on BOG compressors to recover and reintroduce it into the process. This adds complexity, cost, and parasitic energy demand, but is essential for safe and efficient operation. The lesson, familiar from LNG, is that storage systems must actively manage losses rather than assume they can be eliminated.

The same reality applies — even more forcefully — to long-distance shipping. Marine transport of cryogenic liquids inevitably incurs boil-off as heat penetrates insulation over multi-week voyages. In LNG shipping, despite decades of optimisation, typical boil-off rates on modern tankers are on the order of 0.1–0.25% of cargo per day, meaning that a long journey, such as from Qatar to Europe, can result in several percent of the original cargo being consumed or reliquefied during transit. This loss is managed operationally, but it is never zero.

For liquid hydrogen, the challenge is more severe. Hydrogen must be stored at far lower temperatures than LNG, making insulation harder and heat ingress more consequential. Without aggressive mitigation, boil-off losses for cryogenic hydrogen would be higher than those seen in LNG transport. This has important implications for global hydrogen trade. While it is technically possible to ship liquid hydrogen over long distances, ultra-long routes — such as Australia to Europe — are unlikely to be economically or energetically sensible. Losses accumulate, energy penalties rise, and system complexity increases.

As a result, if hydrogen shipping occurs at scale, it is far more likely to be regional rather than intercontinental: Australia supplying Asia, or perhaps the west coast of the Americas, rather than Europe. Alternatively, hydrogen will be converted into more stable carriers, such as ammonia or other hydrogen-derived molecules, that can be transported with far lower boil-off risk. These are not secondary considerations; they fundamentally shape which trade routes and business models make sense.

The key point is not that hydrogen is inefficient or impractical. It is that hydrogen systems reward honesty. Every conversion step, every pressure change and every kilometre of transport has a cost. Ignoring these realities does not make them disappear; it simply shifts them downstream, where they reappear as higher costs, lower availability, or reduced confidence from investors and operators.

In this sense, losses are not a flaw of hydrogen — they are a design constraint. Systems that acknowledge them early can optimise around them. Systems that assume them away cannot. As hydrogen moves from demonstration to deployment, success will increasingly depend on how well these operational realities are understood, quantified, and integrated into design and operation.

Pressure, Loading, and Operating Philosophy

If technology defines what hydrogen systems can do, operating philosophy determines what they really do. Long before optimisation routines or cost models come into play, project developers must make a small number of foundational choices about how electrolysers and associated systems are intended to operate. These choices — often made early and rarely revisited — shape durability, availability, and risk across the full lifetime of an asset.

Two of these choices dominate: how the system is loaded over time, and at what pressure hydrogen is produced.

Electrolyser Loading Philosophies — A Machine's Perspective

From the outside, electrolyser operation is often described in terms of megawatts, operating hours, or capacity factors. Inside the stack, however, loading strategy manifests very differently. What matters is not simply how much power is consumed over a year, but how frequently and how abruptly operating conditions change.

A baseload operating philosophy exposes the electrolyser to relatively stable conditions. Current density, temperature, pressure, and water management remain

within narrow bands for extended periods. From an electrochemical perspective, this stability simplifies control and slows many degradation mechanisms. Catalyst layers experience fewer transient stresses, membranes remain closer to steady-state hydration, and thermal gradients are easier to manage. Baseload operation aligns naturally with industrial off-takers, nuclear supply, or firm grid connections, but assumes a downstream system capable of absorbing constant hydrogen output.

By contrast, variable or renewable-following operation subjects the electrolyser to frequent changes in load. Current density ramps up and down, temperatures fluctuate, and pressure dynamics become more complex. These transients introduce additional stress: membranes expand and contract, catalyst layers experience non-uniform reaction rates, and balance-of-plant components cycle more frequently. Modern electrolysers — particularly PEM systems — are designed to tolerate a degree of this behaviour, but it alters ageing mechanisms and complicates lifetime prediction.

Between these extremes lie band-limited or hybrid operating strategies, where electrolysers vary output within defined ranges while avoiding the most aggressive transients. From a machine perspective, this approach attempts to preserve many of the benefits of steady operation while still providing system flexibility. Achieving this requires coordination between power electronics, control systems, storage buffers, and downstream demand — and increasingly, digital optimisation.

Crucially, an electrolyser designed primarily for steady operation will not behave optimally if forced into highly dynamic service, just as a system optimised for flexibility may carry unnecessary complexity if operated rigidly. Loading philosophy is therefore not a dispatch decision layered on top of a finished design — it is a core design assumption that influences materials selection, cooling strategies, control architecture, and warranty frameworks.

Pressure — Where Energy, Stress, and Integration Meet

Pressure is often introduced as a technical parameter, but in practice it is a system-level decision that determines where energy is consumed, where stress accumulates, and where complexity resides.

As discussed earlier in the context of PEM electrolysis, hydrogen can be produced at atmospheric pressure, at moderate pressures around 10 bar, or at higher pressures approaching 30 bar. Each option shifts trade-offs across the stack and the balance of plant.

Operating at atmospheric pressure minimises mechanical stress on membranes, seals, and stack hardware. It simplifies stack design and generally supports longer component lifetimes. However, hydrogen produced this way must be compressed downstream for storage, transport, or further processing. The energy penalty does not disappear — it is simply transferred into compressors, cooling systems, and auxiliary equipment.

Operating at moderate pressure, typically around 10 bar, represents a pragmatic compromise. Some of the compression work is absorbed electrochemically within the stack, reducing downstream requirements without dramatically increasing mechanical stress. For many systems, this balance offers a sensible midpoint between durability, efficiency, and system simplicity, particularly where hydrogen is stored locally or fed into medium-pressure distribution systems.

Higher-pressure operation, such as 20–30 bar, further reduces or even eliminates the need for initial compression stages downstream. This can simplify plant layout and reduce equipment count. However, the benefit comes at the cost of increased mechanical and electrochemical stress within the stack. Membrane reinforcement, sealing strategies, and material selection become more critical, and degradation sensitivity increases. Pressure cycling — particularly when combined with dynamic loading — can become a dominant lifetime driver.

What matters is that pressure cannot be optimised in isolation. A pressure choice that looks attractive at the electrolyser boundary may impose hidden costs elsewhere in the system, just as a conservative stack design may push complexity into compressors and storage. The “right” pressure is therefore inseparable from downstream integration: whether hydrogen is consumed on-site, injected into pipelines, stored in caverns, liquefied, or converted into derivatives, such as ammonia.

In this sense, pressure is not merely an operating setpoint. It is an architectural decision that locks in where energy is spent, how assets age, and how the system responds to variability.

Operating Philosophy as a Design Commitment

Taken together, loading strategy and pressure define how an electrolyser behaves over time. They influence degradation rates, maintenance intervals, availability, and predictability. They also shape how confidently performance can be guaranteed, insured, and financed.

The costliest failures in hydrogen projects rarely stem from flawed chemistry. They arise when systems are asked to operate in regimes they were never designed for — chasing flexibility without accounting for material stress, or prioritising simplicity while ignoring downstream constraints.

Successful hydrogen systems are not those that maximise a single parameter. They are those that align technology choice, operating philosophy, and system integration into a coherent whole. Pressure and loading are where that alignment is either achieved — or quietly undermined.

In the next section, we turn from how systems are operated to how they are assembled, examining modularity, integration, and the standards required to move hydrogen from bespoke projects to repeatable infrastructure.

Modularity and Standardisation: The Building Blocks of Energy Tomorrow

One of the most underestimated forces shaping the future of energy is not a fuel or a technology—it is modularity. In the fossil era, energy infrastructure was defined by scale. Bigger was cheaper, more efficient, and more reliable. Every major project was a bespoke engineering exercise, meticulously optimised for a specific location and a specific fuel stream. Standardisation existed, but only within narrow technological families: turbines looked like turbines, pipelines looked like pipelines, and refineries looked like refineries. The system grew through replication, not modularity.

The new energy system evolves differently. It grows through modules, not megaprojects. Hydrogen production skids, containerised refuelling stations, battery packs, compression units, ammonia crackers, fuel cell generators—each arrives as a pre-engineered black box, designed to plug into a wider system with minimal customisation. Instead of spending years designing an asset from scratch, companies assemble systems from interoperable components. The energy transition accelerates not through heroic engineering feats, but through repeatable, scalable design patterns.

It is worth noting that this modular logic applies primarily to electricity-driven hydrogen pathways—green, pink, and emerging plasma-based concepts—where production, compression, and storage assets are naturally deployed in discrete units. Blue hydrogen, by contrast, inherits much of its structure from conventional reforming and CO_2 capture plants, where economies of scale and centralised optimisation dominate. The distinction matters, but it does not invalidate modularity where the system architecture supports it.

This shift matters because modularity is what enables the hub-and-island architecture to exist in practice. Energy islands rely on assets that can be sized to local needs, deployed quickly, expanded incrementally, and replaced or upgraded without tearing the entire system apart. A 2 MW electrolyser today can become 6 MW tomorrow simply by adding new skids. A hydrogen refuelling station can expand its dispensing, compression, or storage capacity by adding modular equipment rather than redesigning the entire forecourt. A microgrid can integrate new generation or storage assets through standardised digital interfaces. Modularity turns infrastructure from a fixed asset into a living system.

Standardisation is the second pillar of this transformation. A modular system collapses under its own weight unless the modules can speak a common language. This is where frameworks such as Modular Type Package (MTP) become revolutionary. MTP allows components from different manufacturers—electrolysers, batteries, fuel cells, compressors, pumps, control systems—to integrate through a shared automation interface. Instead of bespoke coding for every project, systems assemble like Lego — not in the sense of simplicity, but in the consistency of how interfaces are defined. EPCs shift from bespoke engineering to system integration. Start-ups can innovate around the edges without destabilising the core. This is how complexity becomes manageable.

It is precisely at this point that a common misconception appears—one I have heard many times from the DCS world:

“MTP was created for pharmaceuticals. Hydrogen plants are one-time builds. MTP doesn’t apply.”

This interpretation reflects a narrow reading of MTP’s origins and an incomplete view of how hydrogen infrastructure is likely to evolve.

MTP does not standardise the physical process itself; it standardises how modular assets are described, controlled, and integrated at the automation layer.

It is true that MTP emerged from industries such as pharmaceuticals, where production lines constantly adapt and equipment is reconfigured to support different “recipes.” But the assumption that hydrogen facilities are static, one-off installations could not be further from reality. Hydrogen plants in many cases will evolve continuously:

- Electrolysers are added as demand grows.
- Storage expands when seasonal or security needs increase.
- Compression and purification skids change as purity requirements shift.

- Battery systems are upgraded, resized, or relocated.
- New carriers—ammonia, methanol, LOHCs—introduce entirely new modules.
- Digital optimisation tools require periodic integration as forecasting improves.
- Individual modules can be replaced or upgraded as needed, allowing containerised or skid-based units to be exchanged without disrupting the wider system.

A hydrogen site today is rarely the hydrogen site it will be five years from now. These are not static chemical plants—they are modular energy platforms.

Where pharmaceuticals use MTP to change recipes, hydrogen systems use MTP to change scale, change function, and change architecture over time. The logic differs, but the need for flexible integration is just as strong—arguably stronger.

Indeed, the hydrogen sector may ultimately benefit more from MTP than pharma ever did, because hydrogen is not confined to one industry. It touches mobility, industry, ports, grids, storage, backup power, and synthetic fuels. Its modularity needs are far broader.

Yet it is also important to acknowledge a reality: MTP is not finished.

The automation layer is mature and powerful, but the safety layer—crucial for high-risk energy environments—has not yet been fully standardised. This limits the degree to which hydrogen systems can rely on MTP alone. For now, additional engineering layers, bespoke safety logic, and site-specific certifications must fill the gap. This slows deployment and adds cost.

But when the safety standard arrives—and it will—the impact will be profound.

Hydrogen systems will move from being project-based to product-based. Entire subsystems will be deployable as certified, safety-integrated modules. Integration times will fall dramatically. EPCs will shift towards orchestrating interoperable packages rather than engineering bespoke solutions. Manufacturers will focus on mass-producible infrastructure, not custom builds. And the market will expand because modular hydrogen will finally become accessible to organisations that lack deep chemical or process-engineering expertise.

In my view, once the safety layer is in place, the sky is the limit.

Hydrogen infrastructure will be deployed with the speed and reliability seen in modern industrial automation. Costs will drop not through scale but through

standardisation. Iteration—not gigantism—will drive improvement. As a result, the energy system will become upgradable, adaptable, and fundamentally more resilient.

This transformation makes hydrogen far more accessible to sectors that would never build a bespoke chemical plant. Supermarkets, logistics hubs, campuses, industrial estates, ports, and municipalities can adopt hydrogen-ready systems without building a chemical engineering department. Hydrogen ceases to be niche—it becomes infrastructural, and infrastructure, unlike bespoke projects, is something financiers know how to price, insure, and replicate.

Modularity also reshapes the economics of hydrogen. In the fossil system, efficiency depended on large units. In the modular system, efficiency depends on learning curves, iteration, and manufacturability. A future electrolyser champion will not win by building the largest electrolyser. They will win by building the most manufacturable electrolyser—one that is standardised, automated, and mass-produced.

Once modularity takes hold, the entire energy system becomes upgradable. In the fossil world, upgrading a plant meant shutting it down, spending heavily, and performing surgical engineering on fixed, monolithic assets. In a modular world, upgrading becomes something entirely different. Capacity can be increased simply by adding more modules. Ageing components can be swapped out without disturbing the rest of the system. Control systems can be renewed through software updates rather than hardware overhauls. Even the dominant energy vector on a site can evolve over time as new modules—ammonia crackers, fuel cells, batteries, heat pumps—are added or reconfigured. The system evolves continuously rather than episodically, adapting to new technologies, new economics, and new operational demands with far greater fluidity.

It begins to resemble computing infrastructure more than civil infrastructure.

But modularity is not a magic bullet. It requires regulatory adaptation, harmonised certification, and cross-manufacturer interoperability. It demands new skills—not in welding pipelines, but in systems engineering, digital automation, and integration. It also requires mindset change: the energy system must shift from bespoke projects to modular platforms.

The deeper point is this: modularity is what allows the energy transition to scale at the pace required. Without it, every project is a one-off challenge, every integration a puzzle, every deployment a multi-year endeavour. With it, the energy

system becomes a global industrial platform—configurable, repeatable, efficient, adaptable.

Modularity is not a feature.

It is the architectural foundation of Energy Tomorrow.

Standardisation, Norms, and Industrial Readiness

One of the most overlooked yet foundational enablers of industrialisation is standardisation — the quiet architecture that allows technologies from different makers, countries, and sectors to interoperate, to be certified, to be financed, and ultimately to scale. In a nascent industry like hydrogen, where every company seems to be designing everything from first principles, the absence of common guidance creates islands of engineering rather than a connected infrastructure. This slows deployment, increases risk, and multiplies cost.

Standards and norms are not just technical checklists. They are trusted reference points that reduce uncertainty for manufacturers, operators, regulators, insurers, and financiers. They define what "safe," "compatible," and "verified" mean across the supply chain — from production and measurement methods to materials, safety protocols, and performance criteria. In mature industries, standards create competition, not conformity: they allow multiple suppliers to compete on quality, price, and innovation while guaranteeing that components and systems work together. Hydrogen urgently needs the same coherence if it is to scale at pace.

Recent progress in Germany illustrates how norms can shift from abstract idea to industrial tool. Under the leadership of Dr. Lydia Vogt, Head of Strategic Development Hydrogen Technologies at the German Institute for Standardization (DIN), a coordinated effort has produced the Standardisation Roadmap for Hydrogen Technologies — the first national strategic framework addressing hydrogen standards in a structured way, spanning production, infrastructure, applications, quality infrastructure, training, safety, and certification.

The roadmap is not an academic exercise. It represents the collective effort of hundreds of experts across industry, research, government, and civil society — organised into dozens of working groups to map existing standards, identify gaps, and prioritise action across the hydrogen value chain. It provides a publicly accessible overview of existing and emerging norms, outlines hundreds of recommendations for future work, and kick-starts implementation projects to close critical gaps. This kind of visibility and coordination is a material milestone for the industry.

The importance of this initiative becomes clear when viewed against the alternative: without norms, every project negotiates its own technical interfaces, safety criteria, measurement methods, and certification pathways. That uncertainty adds cost and risk at every handoff — from OEM to EPC, from builder to operator, and from operator to insurer. It also makes systematic performance comparison — and therefore financing — far more difficult.

Importantly, this work is not happening in isolation. Parallel efforts are underway at the international level. ISO Technical Committee 197 (ISO/TC 197) has been developing global hydrogen standards for decades, covering areas such as hydrogen quality, production, storage, transport, and refuelling infrastructure. In Europe, CEN/TC 268 focuses on hydrogen technologies and is increasingly active as hydrogen moves from research into deployment. In the United States, standards bodies such as ASTM and CSA are contributing guidance on materials compatibility, pressure systems, and safety, while Japan and South Korea have advanced national hydrogen codes closely aligned with early commercial deployment.

What is still missing, however, is coherence. Many of these standards exist in parallel, at different levels of maturity, and often without a clear hierarchy that engineers and project developers can easily navigate. This is why national roadmaps like the one developed in Germany matter so much. They do not replace international standards; they translate them into actionable guidance that can be used today by designers, EPCs, regulators, insurers, and financiers.

At this early stage, many technical rules must still be developed, adapted, or harmonised — for example in areas such as underground storage, mobility applications, hydrogen derivatives, measurement protocols, and cross-sector data exchange. The fact that these gaps have been systematically identified rather than ignored is itself a breakthrough, and it provides a concrete menu of work for the coming years.

Ultimately, standardisation and norms do something that innovation alone cannot: they turn bespoke engineering into repeatable, certifiable, scalable infrastructure. They reduce risk, build trust, lower barriers to entry, and make hydrogen technologies bankable and insurable. They are what allows hydrogen to escape pilot projects and become industrialised infrastructure.

In a world where so much design still happens in isolation, norms are what make shared progress possible. They are undervalued not because they are small, but precisely because they underpin everything else. Hydrogen innovation cannot be fully realised without them.

Convergence and Coherence

If there is one lesson that emerges from the technologies explored across the preceding chapters, it is that hydrogen will not be defined by a single winning pathway. There is no universal electrolyser architecture, no dominant pressure regime, no perfect material set, and no one-size-fits-all operating philosophy. Each of these technologies exists for a reason — shaped by context, constraints, and system-level trade-offs.

Alkaline, PEM, AEM, and solid oxide electrolysis do not compete in a vacuum; they serve different roles within a broader energy system. Blue hydrogen, likewise, is not an ideology but a technical pathway whose relevance depends on geography, infrastructure, and timing. Materials choices, operating strategies, and manufacturing approaches all reflect the same truth: optimisation at the component level only matters if it supports coherence at the system level.

This is why integration matters more than perfection. The energy transition will not be won by the most elegant stack design or the lowest theoretical efficiency. It will be won by technologies that work together — that can be manufactured, certified, financed, operated, upgraded, and trusted at scale. Systems thinking must take precedence over purity tests.

Innovation, in this context, no longer means invention alone. It means better materials that balance performance with environmental responsibility. Smarter operation that embraces variability rather than resisting it. Modular systems that can scale incrementally instead of demanding monumental bets. Digital coordination that turns complexity into insight. And standardisation that allows all of this to move faster, not slower.

The hydrogen economy will succeed not because its technologies are advanced, but because they are assembled coherently into systems that society is willing to trust.

6. The Role of Nuclear (and also Oil & Gas)

"The test of a first-rate intelligence is the ability to hold two opposed ideas in mind at the same time and still retain the ability to function." — F. Scott Fitzgerald

Nuclear Energy: The Underestimated Hydrogen Workhorse

I very nearly called this chapter "The Chapter of Uncomfortable Truths", as the energy transition unfortunately contains quite a few—and nuclear power is one of the biggest. As I showed earlier, achieving full carbon neutrality means replacing around 620 exajoules of energy. Today, renewables provide about 13% of the global mix and nuclear another 5%. That leaves 82% still supplied by oil, gas, and coal.

This simple reality forces us to confront a fact that makes many people uncomfortable: renewables cannot carry the entire load alone. If we are truly serious about reducing CO_2 at the scale required, nuclear power must play a larger role. There is no credible scenario where we achieve deep decarbonisation without it.

Part of the challenge is not technical but ideological. I discuss geopolitics in a later chapter, but here I want to touch on something more fundamental—the way political narratives can distort rational decision-making in the energy transition.

Having lived in Central Europe, I've seen first-hand the strong aversion to nuclear power, especially in Germany. In my view, the debate there became far larger, louder, and more emotional than the underlying risks or technical issues justified. When I moved back to Germany in 2007, it was common to see the yellow sticker on the back of cars: "Atomkraft? Nein Danke!" ("Atomic energy? No thanks!").

A colleague of mine, equally invested in the energy sector, joked that we should print a counter-sticker featuring a prehistoric Fred Flintstone style stone wheel and the slogan: "Steinzeit? Nein Danke!" ("Stone Age? No thanks!"). His point was

that the conversation had drifted so far into symbolism that it no longer resembled a real energy discussion.



I focus on Germany here because this is, frankly, a very German problem. And the difficulty is that when Germany develops a problem—especially in energy—its sheer weight inside the EU means everyone else suddenly inherits it. That's just how the political and economic balance of the union works.

This is one of the uncomfortable truths: in many countries, the energy transition has been treated as a kind of Green, liberal, almost nostalgic political movement—something associated with lifestyle and identity rather than engineering, economics, and physics. At the same time, ironically, some of these same Green political currents are among the most resistant to the technologies that we really need to achieve the transition: nuclear power, blue hydrogen, and carbon capture.

In other words, the energy transition is sometimes being attacked from both sides.

For me, one of the worst energy decisions made by any European government in the past 15 years was Germany's choice to shut down its nuclear power plants after Fukushima. Angela Merkel had a remarkable ability to read public sentiment and reposition herself to occupy the political centre. My interpretation is that, in this case, she saw a wave of fear rising, chose to ride it, and in the process created a long-term structural problem for Germany's energy system. As I said at the time to many colleagues: as a physicist, she should have known better. I remain con-

vinced that keeping nuclear online would have put Germany in a far stronger position today. Politically, however, the decision allowed her to respond to the public mood, win favour among voters sympathetic to the anti-nuclear position, and strengthen her position ahead of the next election.

What makes this even more fascinating is the mismatch between political identity and hydrogen leadership. Some of the strongest, most passionate advocates for hydrogen in Europe are not from Green parties at all. Listen to Hubert Aiwanger, leader of the Freie Wähler in Bavaria, or Jorgo Chatzimarkakis, the former MEP from the FDP, who now leads Hydrogen Europe as their CEO. They speak with clarity, urgency, and—in Jorgo's case—refreshing brutal honesty about what needs to happen. These are exactly the voices you might expect from the Green parties, yet they are noticeably absent. The Greens' deep scepticism of what they consider "dirty industries" often overshadows the practical steps required to decarbonise those very sectors.

It would be unfair, however, to suggest that pragmatic voices do not exist within the Green movement itself. Anyone who has had the pleasure of speaking with Dr Danyal Bayaz, the Finance Minister of Baden-Württemberg, will recognise a politician deeply committed to the energy transition but also focused on making it work in the real world—cutting bureaucracy, enabling investment, and approaching the challenge with a practical rather than ideological mindset.

As Germany is such a central actor within the EU, its anti-nuclear stance doesn't remain a purely domestic matter. It spills over into EU hydrogen policy, especially in determining what counts as "low-carbon" hydrogen for funding and regulation. Hydrogen produced from natural gas with carbon capture (blue) or from electrolysis powered by nuclear (pink) has, in my view, historically not been sufficiently acknowledged or supported. This has led to a narrower, more constrained policy framework than the energy system needs.

Nuclear energy—and the pink hydrogen it enables—is not just a technical tool; it is a strategic asset. Nations with nuclear capabilities, including the US, France, the UK, South Korea, Japan, and China, view nuclear-powered hydrogen as a lever for energy sovereignty, industrial competitiveness, and global influence.

Producing hydrogen domestically allows countries to reduce dependence on imports, leverage existing nuclear supply chains, and develop export markets for low-carbon hydrogen. Pink hydrogen becomes a tool for industrial power, technology leadership, and secure energy supply, reinforcing a nation's position in emerging clean-energy markets.

In short, I believe Germany's ideological position has shaped the whole EU's approach—to its detriment.

France has always taken a very different view of nuclear energy than many of its neighbours. Where others see nuclear as a political liability, France sees it as an asset—one of the foundations of its industrial strength and national identity. So, when hydrogen entered the conversation as the next big piece of the energy transition, France didn't hesitate for a second. Why on earth wouldn't you use your largest, most stable source of low-carbon electricity to produce it? For the French, the idea that nuclear-powered hydrogen—what we now call pink hydrogen—should somehow be treated differently from renewable-based hydrogen was simply absurd. It would be like owning a vineyard and insisting on importing your wine.

And this, in essence, is where the long EU debate begins; not with technology, but with philosophy. For years, Europe has been arguing not about whether pink hydrogen is low-carbon, but whether it fits the political narrative certain member states prefer to tell. France was ready to move. Brussels, less so.

Part of the problem lay in the EU's long struggle to translate the concept of "low-carbon hydrogen" into a clear and operational regulatory framework. In principle, defining hydrogen based on lifecycle emissions should have been relatively straightforward. In practice, the process became drawn out and politically contested, with multiple draft definitions circulating for years before a formal methodology was finally adopted in 2025.

Even now, while the framework is in place, industry is still working through the practical details of implementation, certification, and how different production pathways will ultimately be treated. In energy infrastructure terms, that period of uncertainty has been costly: projects stall, investment hesitates, and industrial momentum shifts elsewhere.

For France, this delay has never been just bureaucratic irritation; it has been a competitive handicap. While green hydrogen secured an early and privileged position within EU funding and regulatory frameworks, pink hydrogen long sat in a grey zone—technically viable, politically tolerated, but not fully recognised. You could build it, but you could not reliably monetise it. And in the EU, definitions determine the flow of capital. To the French, the entire situation felt less like caution and more like an attempt to slow-roll nuclear by burying it under paperwork.

If you ask the French, they will politely suggest that the long road toward defining "low-carbon hydrogen" did not unfold purely by accident. Countries where anti-

nuclear sentiment runs deep—Germany foremost among them—were widely seen as exerting quiet but persistent resistance during the years when these definitions were being debated. Here is where the topic becomes uncomfortably political, as in the EU, when Germany objects to something, the impact radiates outward. Berlin does not always need to veto policies outright; sustained reluctance is often enough to slow the machinery. When Germany has an issue, it has a tendency to become an EU issue, whether the rest of the bloc likes it or not.

But then something shifted. In 2025, after years of tension around the role of nuclear power in Europe's hydrogen strategy, Germany softened its stance. The new government signalled that it would no longer actively resist recognising nuclear-derived hydrogen within the EU's emerging low-carbon framework. What had been a largely French monologue began to resemble a Franco-German dialogue. In Brussels, that changes everything. A united front from Paris and Berlin remains the closest thing Europe has to an unstoppable force.

This shift didn't come out of nowhere. In my view, Germany is wrestling with the realities created by its own decision to shut down its nuclear plants, and the consequences are becoming clearer every year: higher electricity prices, tensions around supply security, and a renewable build-out that—even at its most optimistic—cannot cover every sector. Germany may not be rushing back to nuclear power, but its changing position suggests a growing recognition that blocking others from using it is strategically self-defeating.

France, meanwhile, has moved ahead with characteristic determination. It has positioned pink hydrogen as a central pillar of its industrial recovery plan, tying it directly to its nuclear renaissance. New targets aim for 4.5 GW of electrolysis by 2030 and roughly double that by 2035. The logic is straightforward: France already has abundant low-carbon electricity, so why wait?

The real issue, however, is not France's enthusiasm but Europe's hesitation. Even though a formal framework for low-carbon hydrogen has now begun to emerge, the years spent debating definitions created a long period of uncertainty for industry. Major infrastructure investments rarely move forward without regulatory clarity, and during that time companies struggled to commit capital with confidence. It is difficult to build a market when one of its most promising production pathways spends years trapped between technical viability and political hesitation.

And this matters far beyond France. Pink hydrogen is not just a French preference; it is one of the few ways Europe can produce hydrogen competitively on its own territory. If the EU insists—politically or ideologically—on green hydrogen

only, it risks becoming dependent on imports from regions with better sun and wind resources. That is the uncomfortable truth at the heart of the debate: pink hydrogen is not just about emissions; it is about sovereignty.

Where we stand today is a moment of transition. France is pushing hard. Germany is shifting. Brussels is slowly awakening to the risk that delay equals decline. As these forces converge, the position of pink hydrogen in the European energy landscape is finally beginning to clarify, even if the regulatory details and certification systems are still catching up.

In many ways, the entire saga symbolises the broader challenge of the energy transition in Europe: a tension between ideology and pragmatism, between national sensitivities and collective ambition, and between what looks good politically and what works industrially. Pink hydrogen sits right at that crossroads, and what Europe decides over the next few years will determine whether hydrogen becomes a pillar of European competitiveness—or yet another area where Europe invents the concept while someone else scales the industry.

I have spoken a lot about the European Union here, but it would be wrong to think this is solely a European problem. Across the Atlantic, the United States faces a very similar challenge — California, in particular, is grappling with what to do with its nuclear power, once again shaped heavily by ideology. The state's laws give preference to renewable sources being placed on the grid, often to the detriment of other low-carbon options. Now, in principle, prioritising renewables isn't a bad thing. But when you also have a baseload of nuclear power available, as California does, the picture becomes more complicated.

A few years ago, in a conversation with a colleague in the US energy sector, I learned something that still makes me shake my head: they were often forced to dissipate excess nuclear-generated electricity from the system — in some cases by routing it into resistor banks to dissipate as heat — because the grid had no mechanism to use it productively. Low-carbon electrons turned into heat. Literally wasted. Imagine it: a perfectly clean, steady supply of electricity simply being destroyed because the system has no legal or practical mechanism to put it to productive use. Surely there is a better option.

And there is. That same energy could be used to produce hydrogen. Electrolysis doesn't care whether the electricity comes from wind, solar, or nuclear — it just needs electrons. California has the baseload capacity to create hydrogen continuously, even when solar panels are silent at night. That hydrogen could then be stored, transported, and used across industry, transport, or even exported to other states. In short, it could turn wasted energy into a strategic asset.

In fact, pilot projects in the U.S. have already demonstrated this. Most notably, the Nine Mile Point Nuclear Power Station in Oswego, New York, began producing hydrogen on-site in 2023 using low-temperature electrolysis powered by nuclear electricity. The plant has been able to operate continuously, producing hydrogen with very low carbon intensity — exactly the kind of "pink hydrogen" France is promoting in Europe. The concept is proven; the barrier is political, not technological. It's almost comical, if it weren't so costly, that in a world racing toward decarbonisation, we are literally throwing away perfectly good, low-carbon electrons.

Nuclear energy is often thought of as just electricity, but in a hydrogen system, it offers unique advantages. Its baseload output means hydrogen can be produced continuously, day and night, without worrying about clouds or windless days. Unlike solar or wind, nuclear electricity is predictable and stable, which translates into a stable cost structure for hydrogen production—a critical factor for industries like steel, chemicals, and refining that require uninterrupted supply.

Moreover, nuclear reactors generate high-temperature heat, which can be directly used to boost the efficiency of electrolysis. This combination of low-carbon electricity and high-grade thermal energy makes nuclear uniquely suited to produce pink hydrogen at scale. While renewables will dominate the bulk of green hydrogen, nuclear fills the gaps in reliability, consistency, and 24/7 operation, which are essential for industrial off-takers who cannot tolerate interruptions.

Small Modular Reactors, or SMRs, take nuclear-powered hydrogen production to the next level. Unlike traditional reactors, SMRs are modular and scalable, meaning they can be installed incrementally to match demand. They produce both electricity and high-temperature heat (300–900°C), which can feed electrolysis systems efficiently.

The real advantage is proximity to demand. By locating SMRs near industrial hubs or hydrogen "clusters," hydrogen can be produced right where it will be consumed. This reduces the need for long-distance pipelines or imports, cuts transmission losses, and strengthens regional energy security. SMRs transform nuclear from a centralised, one-size-fits-all solution into a flexible tool that can support local industrial economies and hydrogen networks.

California's dilemma mirrors what we've seen in Europe: when ideology dominates energy policy, even abundant, clean resources can be sidelined. And while the politics differ on either side of the Atlantic, the effect is the same: low-carbon energy that could be powering the transition is instead being constrained or wasted. This brings us back, inevitably, to the central theme of this chapter: the

critical role of nuclear, and alongside it, oil and gas, in providing the reliable, low-carbon energy we need — not just in theory, but in practice. The tools exist; the challenge is choosing to use them.

There is one further point about nuclear energy for the future that needs to be addressed, and if it works out as planned, it could hopefully change the entire debate around nuclear energy.

Back in 1995, as I was waiting for the train at Ashford Station, I met a young woman who, little did I know at the time, would become one of the most intelligent and inspiring people I would ever know. Her name is Kate Lancaster (now Dr. Lancaster), a lecturer at the University of York whose passion is physics — and within that, an even deeper fascination with lasers (she has some excellent videos on YouTube for the Royal Institution that I highly recommend).

So, where do nuclear energy and lasers come together? The answer is nuclear fusion. In some fusion experiments, lasers act like tiny, ultra-powerful hammers. They fire at a minuscule pellet of hydrogen fuel, causing it to implode under tremendous heat and pressure. This violent squeeze forces the hydrogen atoms to fuse, releasing energy hotter than the Sun — all from a pellet smaller than a pinhead.

One of the perks of knowing someone in this field is that I occasionally get a peek behind the curtain of these developments. Originally, when I was working in oil and gas, my first instinct was to ask when she was going to make me redundant. Now, it's to understand how progress is unfolding — and progress, I can tell you, is definitely being made. Breakthroughs in laser-driven fusion are slow and painstaking, but the pace of innovation is accelerating. If scientists can scale this from a tiny laboratory pellet to something that produces consistent, usable energy, fusion could offer a limitless, zero-carbon source of power — potentially rewriting the rules of energy entirely.

Beyond the laboratory curiosities, there are real, large-scale experiments pushing nuclear fusion toward practical reality. One of the most prominent is the Wendelstein 7-X in Germany, a stellarator designed to demonstrate that plasma — the hot, charged gas where fusion occurs — can be confined stably for longer periods. Unlike earlier experiments that struggled with instability, Wendelstein aims to show that fusion reactions can be sustained in a controlled way, a key step toward usable energy. Meanwhile, in the United States, projects like the National Ignition Facility are steadily improving laser-driven fusion, bringing us closer to the point where more energy is output than is input — the long-sought milestone called ignition. Optimistically, many in the field suggest that small pilot

fusion plants could begin producing usable electricity within the next 15–20 years, though commercial-scale deployment is likely to take longer.

What makes fusion fundamentally different from the nuclear power that most people are familiar with — fission — is safety. Fission reactors, like the ones that supply much of today's low-carbon electricity, split heavy atoms, such as uranium, producing radioactive waste and requiring complex safety systems to prevent accidents. Fusion, by contrast, involves combining light atoms, typically isotopes of hydrogen, and the reaction is intrinsically self-limiting. If anything goes wrong — if the plasma becomes unstable or cooling fails — the reaction simply shuts itself down. There is no risk of a runaway chain reaction, no meltdown, and the only byproducts are helium and, in some designs, short-lived radioactive materials.

This safety feature alone could revolutionise the conversation about nuclear power. The fear and political resistance that have slowed fission development around the world would be far less relevant for fusion. If fusion proves scalable and reliable, it may finally remove one of the largest ideological obstacles to nuclear energy — and with that, open the door for countries to consider it as a mainstream, zero-carbon pillar of their energy system, alongside renewables and potentially even hydrogen production.

The implications for hydrogen are immediate. Fusion, once scaled, could provide a virtually limitless, zero-carbon source of electricity and high-temperature heat — exactly what is needed for pink hydrogen production. Unlike fission, fusion's inherent safety means it could be deployed closer to industrial hubs without the same political or regulatory constraints, enabling continuous, reliable hydrogen generation. In addition, the abundance and predictability of fusion power should dramatically reduce the cost of electricity, which in turn lowers the production cost of hydrogen, making low-carbon hydrogen not only more abundant and politically acceptable, but also more affordable. In other words, successful fusion doesn't just promise a new type of electricity; it could transform the economics and politics of the entire hydrogen economy.

However, nuclear alone cannot shoulder the energy transition. Oil and gas still provide the lion's share of the world's energy and possess vast infrastructure, technical expertise, and capital that can be redirected toward low-carbon hydrogen. Just as nuclear offers reliable electrons and heat for hydrogen production, oil and gas can pivot from fossil fuels to blue hydrogen, leveraging existing pipelines, storage facilities, and industrial know-how. Ignoring these sectors would be a strategic misstep: the transition needs all hands-on deck, combining

the reliability of nuclear with the flexibility and scale of hydrocarbon-based infrastructure.

Oil & Gas: From Hydrocarbons to Molecule Management

My background, as I have already said, is in oil and gas. I will, however, point out that I am very much in favour of electrolysis-based hydrogen (green and pink) and believe that should be the focus. At the same time, I am enough of a realist to acknowledge that we still need blue hydrogen — hydrogen produced from natural gas with carbon capture. The reason is simple: the infrastructure, technical expertise, and industrial know-how built over decades in the oil and gas sector cannot be replaced overnight. Pipelines, storage facilities, and industrial clusters designed for hydrocarbons can be repurposed to handle hydrogen, making the transition more practical and cost-effective.

Oil and gas companies are often portrayed as obstacles in the energy transition, but the reality is more nuanced. These companies already manage molecules, pipelines, storage caverns, ports, and global energy distribution — hydrogen is simply another molecule. Their decades of operational experience mean that scaling hydrogen, whether blue, pink, or green, is largely a logistical challenge rather than a technological one.

My personal belief, looking forward over a timeline of decades, is that blue hydrogen will serve as one of the essential bridges until fusion-based hydrogen becomes commercially viable — hopefully within 20–30 years. During that period, blue hydrogen can provide reliable low-carbon supply, help scale demand, and ensure that industrial and transport sectors have the continuous energy they need. Importantly, when combined with rigorous carbon capture, it allows us to manage emissions responsibly while we build the infrastructure for a cleaner, nuclear- and fusion-powered future.

Beyond being a bridge, blue hydrogen is strategic: by producing large-scale low-carbon hydrogen today, it can anchor early hydrogen markets, reduce emissions from refineries and ammonia plants, and provide the industrial experience and supply chains needed for the longer-term green and pink hydrogen transition.

Why blue hydrogen remains important in the near term is relatively straightforward. We already have a mature oil and gas industry, established natural gas supply chains, and extensive infrastructure for transporting and handling gas. We also have mature hydrogen production technologies: as discussed in the previous chapter, SMR and ATR remain the principal methods used to produce hydrogen from natural gas today. In principle, therefore, the pathway to blue hydrogen

appears simple: combine these established processes with carbon capture. But is it really that straightforward?

Much of the hydrogen infrastructure we need already exists: gas pipelines can be converted or blended with hydrogen, LNG terminals can evolve into hydrogen carriers, geological formations are ideal for hydrogen or CO_2 storage, and industrial ports can become hydrogen export hubs. Leveraging these existing assets allows the transition to be faster, cheaper, and safer than building everything from scratch.

Unfortunately, while the principle is straightforward, the practicalities are more complex. Capturing a high proportion of the CO_2 produced requires additional energy and sophisticated equipment, which increases costs and reduces overall efficiency. Then comes the question of storage: the captured carbon must be securely injected underground or otherwise permanently sequestered, and the geological, regulatory, and safety aspects of that are far from trivial. Finally, even with all this in place, blue hydrogen remains dependent on fossil fuels, meaning it can only ever be a transitional solution rather than a long-term zero-carbon answer.

All of this means that while green hydrogen can be developed and deployed at small scale, blue hydrogen generally requires much larger, highly integrated industrial facilities, translating into higher capital intensity and more complex project development.

The economics are revealing. Even with existing infrastructure, adding carbon capture to SMR or ATR systems increases the cost of hydrogen production significantly. Estimates vary, but while blue hydrogen is often cheaper than green hydrogen today, that advantage is highly sensitive to natural gas prices, capture efficiency, methane leakage, and access to CO_2 transport and storage. Under conditions of high gas prices, stringent capture requirements, or very low-cost renewable electricity, the cost gap can narrow sharply or even invert. In other words, while blue hydrogen leverages what we already have, it is not a cheap shortcut — it is a carefully managed bridge, bought at a premium in complexity, risk, and long-term dependency, to ensure continuity and reliability as cleaner alternatives scale.

This cost sensitivity underscores why clear policy frameworks and investment signals are required. At the same time, cost alone does not negate the strategic importance of blue hydrogen. It remains essential for scaling the hydrogen economy in the near term, particularly because of its location within existing industrial clusters that already consume hydrogen as feedstock — refining, ammonia

production, steelmaking, and other heavy industries. This proximity reduces transport requirements, lowers distribution costs, and ensures a reliable, continuous supply where it matters most.

Blue hydrogen enables the build-out of industrial demand, supply chains, and the technical and regulatory frameworks required for a low-carbon hydrogen market. Without it, there is a risk of a widening gap between ambition and reality — industries that depend on uninterrupted hydrogen supply would be left exposed while green and pink hydrogen scale. In that sense, blue hydrogen functions as a transitional backbone, buying time and operational experience as renewables, nuclear, and eventually fusion expand to meaningful capacity.

Carbon capture, utilisation, and storage (CCUS) is not just for hydrogen; it is already critical for heavy industries such as cement, steel, and chemicals. These sectors provide real-world laboratories for scaling capture technology, improving solvents, compression methods, transport, and long-term storage solutions. Lessons learned here can be directly applied to blue hydrogen, accelerating deployment and reducing risk.

Take cement, for example: producing one ton of cement typically emits around 0.7–0.9 tons of CO_2, much of it unavoidable because it comes from the chemical transformation of limestone into clinker. In steelmaking, especially in blast furnaces, emissions are similarly embedded in the process itself, not just in the energy used. For these sectors, CCUS is not optional; it is a requirement under increasingly strict emissions constraints.

Several real-world projects illustrate this progression. The Norcem Brevik cement plant in Norway has implemented full-scale carbon capture on part of its production, sending CO_2 to permanent storage under the North Sea. In steel, the ArcelorMittal plant in Ghent, Belgium is piloting capture to reduce blast furnace emissions, with plans to expand as technology and economics evolve. Chemical plants producing ammonia or methanol are also increasingly integrating capture solutions, taking advantage of relatively concentrated CO_2 streams.

The important insight for hydrogen is that these developments in heavy industry act as scaling platforms for CCUS. As these sectors improve capture efficiency, compression, transport, and storage integrity, they build the knowledge, supply chains, and economies of scale that can be transferred to blue hydrogen production.

This cross-industry dynamic is particularly important because hydrogen alone is unlikely to drive CCUS to breakthrough cost levels. By integrating lessons from

cement, steel, and chemicals, the hydrogen sector can benefit from faster technology maturation. As CCUS becomes routine in these established industries, the equipment, skills, and regulatory frameworks will already be in place — enabling blue hydrogen to fulfil its transitional role effectively, bridging today's emissions gap while green, pink, and eventually fusion-based hydrogen scale.

The Strategic Pivot of Oil & Gas Companies

Hydrogen represents a strategic path for oil and gas companies to evolve rather than disappear. These firms are uniquely positioned to take a leadership role in the hydrogen economy because they already possess capabilities that are directly transferable. Decades of managing multibillion-dollar assets, running complex safety and regulatory frameworks, and overseeing deep subsurface engineering projects give them a technical foundation that few other sectors can match. Beyond engineering, their experience in global commodity trading, long-term contracting, and risk management enables them to structure hydrogen markets, secure supply chains, and manage the uncertainties inherent in scaling a new energy system.

The operational expertise of oil and gas companies means they can approach hydrogen production and distribution with a level of scale and reliability that newcomers would struggle to achieve. Pipelines, storage facilities, and terminals—originally built for hydrocarbons—can be repurposed to carry and store hydrogen, reducing costs and accelerating deployment. Their existing industrial clusters, often located near chemical plants, refineries, and steelworks, provide a natural network for hydrogen offtake. In practical terms, this means that when oil and gas companies move into hydrogen, they are not creating an entirely new market from scratch—they are adapting an enormous, global infrastructure that already works.

Recognising the inevitability of the energy transition, most major players—BP, Shell, Equinor, TotalEnergies, and Aramco—have actively begun pivoting toward low-carbon molecules. Their investments are broad, spanning blue hydrogen, green hydrogen, ammonia supply chains, CCS hubs, and synthetic fuels. Each of these represents a strand of the transition. Blue hydrogen leverages existing natural gas assets and expertise in carbon capture; green hydrogen demonstrates a commitment to renewable integration; ammonia provides a hydrogen carrier that can be traded internationally; CCS hubs not only capture emissions from hydrogen but also from adjacent industrial sectors; and synthetic fuels allow for the continuation of energy-intensive transport while decarbonizing upstream pro-

cesses. For these companies, hydrogen is not a threat, but simply the next molecule in a portfolio that they have spent decades mastering.

Economics are another compelling driver. The oil and gas industry generates over $2 trillion in annual revenue. Even a modest redirection of capital toward low-carbon hydrogen can dramatically accelerate the sector's growth. Large-scale deployment reduces costs, improves efficiency, and strengthens the market, creating a virtuous cycle: investment enables scale, scale drives cost reduction, and lower costs attract further investment. Their financial capacity also provides confidence to governments, private investors, and industrial off-takers that hydrogen projects can be executed reliably and at industrial scale. In other words, the sector's financial muscle is a lever that can catalyse the broader hydrogen economy.

Beyond technical skill and financial power, oil and gas companies bring project management sophistication that is crucial for hydrogen development. Hydrogen hubs, spanning blue and green hydrogen production — and increasingly involving nuclear-linked pink hydrogen through shared infrastructure — require careful coordination across multiple stakeholders, including utilities, industrial off takers, regulators, and transport operators. O&G firms have decades of experience delivering multi-billion-dollar projects on time and on budget under highly regulated and complex environments. They understand permitting, environmental compliance, and risk management—lessons that can now be applied to hydrogen infrastructure development.

There is also a geopolitical dimension. Many oil and gas firms operate globally, with expertise in cross-border energy trade, long-term contracting, and international markets. This allows them to not only supply hydrogen domestically but also to participate in global hydrogen trade, creating networks that can support countries without domestic production capacity. Their experience navigating energy geopolitics ensures that hydrogen can be produced, moved, and sold safely, reliably, and profitably, essential factors if hydrogen is to become a true global commodity rather than a niche industrial input.

Finally, this pivot is not solely about adapting current assets; it is also about shaping the future energy system. Oil and gas firms are uniquely capable of bridging the current gap between fossil fuels and fully renewable or nuclear-powered hydrogen. Blue hydrogen can provide a reliable, low-carbon supply in the near term while the renewable and nuclear sectors scale up. At the same time, investments in green hydrogen, ammonia, and synthetic fuels demonstrate a long-term commitment to decarbonisation. In effect, these companies are creating a bridge across the transition—a bridge that combines technical expertise, financial power, and industrial scale.

In short, the strategic pivot of oil and gas companies is not merely a defensive manoeuvre to protect market share. It is a proactive, multidimensional strategy to leverage existing competences, infrastructure, and capital to build a new hydrogen economy. By translating project management, engineering, regulatory, trading, and financial capabilities into hydrogen production, storage, and distribution, these companies can deliver scale, reliability, and efficiency that will be essential for the energy transition. Their participation ensures that hydrogen is not only technically feasible, but also economically and industrially viable, complementing the role of nuclear, renewables, and emerging fusion technologies in achieving a low-carbon future.

Why Oil & Gas Will Remain Essential

Even as hydrogen technologies ramp up and the energy transition accelerates, oil and gas will remain deeply embedded in our society, economy, and industrial systems for decades to come. The global energy system is massive: in 2025, we still consume roughly 600 exajoules of energy annually, with hydrocarbons supplying more than 80% of that demand. Hydrogen—despite its transformative potential—cannot instantly replace all this energy. While hydrogen can scale rapidly in industrial clusters, transport, and select power applications, fossil fuels continue to supply critical energy in sectors where alternatives are technically or economically challenging. The scale of global demand ensures that oil and gas will continue to power many facets of modern life, not as a convenience but as a necessity.

Yet the role of hydrocarbons extends far beyond powering electricity grids or vehicles. Oil and gas provide the raw materials for countless products that define our daily lives. Plastics, synthetic fibres, fertilisers, pharmaceuticals, cosmetics, lubricants, and industrial chemicals—all of these rely on hydrocarbons as feedstocks. The chemistry behind these materials is deeply intertwined with petroleum and natural gas; creating alternatives is neither simple nor fast. Even with green chemistry breakthroughs, scaling bio-based or recycled substitutes to meet global demand will take decades. Hydrogen, valuable as it is for energy, cannot yet substitute for these essential industrial and consumer applications.

Consider the case of fertilisers, where ammonia production depends almost entirely on hydrogen derived from natural gas. While blue or green hydrogen can eventually provide low-carbon alternatives, the infrastructure required to produce, store, and distribute hydrogen at the scale needed for global agriculture is immense. Likewise, the plastics and chemical industries depend on hydrocarbons not only as energy, but as building blocks. Converting these supply chains to

alternative molecules would require massive investment, new chemical processes, and regulatory approval, all while maintaining affordability and reliability. It's not merely a question of switching fuels; it's a question of sustaining global food, medicine, and industrial systems.

This dual reality—hydrogen scaling alongside, but not replacing, hydrocarbons—underscores the strategic importance of oil and gas companies. Their operational expertise, global infrastructure, and financial resources are not just useful for producing hydrogen; they remain indispensable for meeting ongoing energy and material needs. Pipelines, storage caverns, ports, industrial clusters, and refining capacity represent decades of investment. These assets cannot simply vanish; instead, they can be repurposed, adapted, and optimised to support the low-carbon transition. By investing in hydrogen while continuing to supply essential products and energy, these companies act as the bridge between today's world and a decarbonised future.



Adapted from the International Association of Oil & Gas Producers (IOGP), Oil in Everyday Life; illustration redesigned for this book.

Moreover, oil and gas firms possess technical and organisational capabilities that are highly transferable. Project management for multibillion-dollar assets, risk

assessment, global commodity trading, regulatory compliance, and subsurface engineering are not skills easily replicated by newcomers. When these capabilities are applied to hydrogen infrastructure, carbon capture projects, or synthetic fuel production, the learning curve accelerates, costs come down, and scale can be reached faster. Essentially, the same systems that delivered oil and gas to the world can now deliver low-carbon molecules, supporting industrial decarbonisation at scale.

There is also a geopolitical dimension. Hydrocarbons are entwined with global trade, finance, and national energy security. Countries and industries dependent on continuous energy flows cannot afford abrupt disruption. Hydrogen is growing, but supply chains, storage, and transport networks remain nascent. In the near term, oil and gas provide reliability and flexibility, ensuring that energy is available where and when it is needed, especially during peak demand periods, extreme weather, or grid stress. Ignoring this role would be not just impractical but potentially destabilising.

Even within the hydrogen ecosystem itself, oil and gas will play a continued role. Blue hydrogen, for example, relies on existing gas infrastructure combined with carbon capture. Scaling green hydrogen quickly is essential, but it cannot yet match the volume, reliability, and cost-effectiveness of hydrocarbons for large industrial off-takers. In steel, refining, and ammonia production, hydrogen demand is measured in megawatts or gigawatts, not kilowatts. Blue hydrogen, enabled by oil and gas expertise, fills this gap, providing an early backbone while pink hydrogen from nuclear and green hydrogen from renewables scale up.

Furthermore, oil and gas's continued presence ensures that society has time to adapt technologically and socially. Energy transitions are rarely instantaneous; they unfold over decades, shaped by infrastructure, capital, regulation, and behaviour. Abruptly removing hydrocarbons from the system would risk supply shocks, price volatility, and industrial disruption. By contrast, an integrated approach—where hydrocarbons coexist with hydrogen and renewables—allows for a controlled, resilient transition, giving governments, businesses, and consumers the time to innovate, invest, and adapt.

Finally, oil and gas are essential simply because of the breadth of human activity that depends on them. Beyond power generation and transport, hydrocarbons feed the medical, chemical, agricultural, and manufacturing sectors. Pharmaceuticals, fertilisers, plastics, cosmetics, and lubricants all originate from oil and gas. Even as hydrogen grows in importance, these sectors cannot immediately pivot away from fossil-based feedstocks without significant technological

breakthroughs. For now, hydrocarbons remain indispensable—not as a relic of the past, but as a foundation supporting the new energy era.

In short, even in a hydrogen-powered future, oil and gas remain central. Their infrastructure, expertise, and capital support both the energy transition and society's ongoing energy and material needs. Hydrogen and hydrocarbons are not mutually exclusive—they are complementary. Understanding this coexistence is crucial for realistic energy policy, strategic planning, and industrial investment. The transition is not about discarding the old, instead upgrading and repurposing it. Only by recognising the continued role of oil and gas can the world ensure a smooth, reliable, and scalable transition to a low-carbon hydrogen economy.

Ultimately, while hydrogen, nuclear, and renewables are reshaping the energy landscape, oil and gas continue to underpin both our energy and material needs. Plastics, pharmaceuticals, fertilisers, and countless industrial products still rely on hydrocarbons, and global energy demand remains immense. Even as we scale low-carbon alternatives, the reality is that oil and gas provide the flexibility, reliability, and infrastructure that cannot be replaced overnight. The question is no longer whether we will use these resources—they will remain part of the system for decades—but whether we will ever fully remove them from the energy mix. The answer is far from certain. What is clear, however, is that any credible path toward decarbonisation must recognise the enduring presence of oil and gas, and leverage it strategically rather than ignore it.

The BBQ Quandary: Expectations vs Reality

One of the challenges when discussing the energy transition is that we naturally focus on the largest and most visible parts of the system: power stations, wind farms, solar parks, hydrogen pipelines, industrial plants, and electric vehicles. Yet many of the most interesting lessons can be found in far smaller and more familiar places.

Consider the humble barbecue.

As spring turns to summer, millions of people across the world wheel their gas barbecue out into the garden, connect a propane cylinder, and think little more about where the fuel came from. The propane bottle is simply there, available at the local hardware store, ready to provide a few hours of cooking and enjoyment.

Yet that small cylinder tells an important story about the complexity of the energy transition.

Today, most propane is not produced because society specifically needs barbecue fuel. Instead, it emerges as a by-product of natural gas processing and oil refining. When hydrocarbons are extracted and processed, they are separated into a range of useful products: gasoline, diesel, jet fuel, lubricants, petrochemical feedstocks, butane, propane, and many others. The propane used in barbecues, caravans, motorhomes, small boats, rural homes, camping equipment, forklifts, and off-grid properties emerges as part of a much larger hydrocarbon value chain.

This raises an interesting question.

Imagine it is 2050. Oil and gas consumption has fallen dramatically, renewable electricity dominates much of the grid, and hydrogen has become a major industrial energy carrier. You walk into a store to buy a propane cylinder for your barbecue.

Where did it come from?

Today, the answer is simple: it is a consequence of producing oil and natural gas. Since propane is largely produced as part of a much larger hydrocarbon system, consumers rarely see its true value. In many ways it arrives as a convenient by-product of activities that society is already undertaking for other reasons. In a future where far less oil and gas are produced, the answer becomes much more complicated.

One possibility is bio-propane, sometimes referred to as renewable propane. Produced from waste cooking oils, vegetable oils, animal fats, biomass-derived feedstocks, and renewable fuel processes, it is chemically identical to conventional propane. Existing cylinders, burners, heaters, cookers, and distribution networks continue to work exactly as they do today. For the consumer, there would be virtually no noticeable difference.

Another possibility is synthetic propane. In principle, hydrogen produced from water can be combined with captured carbon dioxide to create hydrocarbon molecules, including propane. The chemistry is well understood and aligns closely with many of the Power-to-X concepts discussed elsewhere in this book. The challenge is economics. Producing synthetic propane requires renewable electricity, electrolysis, carbon capture, synthesis processes, and distribution infrastructure. It is entirely possible, but almost certainly more expensive than obtaining propane as a by-product of conventional hydrocarbon production.

A third possibility is that many propane applications simply become electrified. Home heating may increasingly rely on heat pumps. Cooking may transition to

induction systems. Water heating may become electric. For many consumers, this will be a perfectly reasonable outcome.

However, electrification is not always the complete answer.

Consider a remote cabin in northern Canada, a tiny home in rural Scotland, a sailing yacht, a caravan, or an off-grid property. In many of these situations, a propane cylinder is an exceptionally effective energy storage device. It stores energy for months, requires no charging infrastructure, operates regardless of weather conditions, can be transported easily, and can be replaced in minutes. The simplicity of the solution is often its greatest strength.

This naturally leads some people to ask whether hydrogen could eventually replace propane cylinders.

For most domestic and portable applications, the answer is probably not.

Hydrogen excels in many applications, but portable cooking and heating are unlikely to be among them. Propane liquefies under relatively modest pressure and can be stored safely in inexpensive cylinders at ambient temperatures. Hydrogen, by contrast, requires either very high-pressure storage or cryogenic temperatures approaching -253°C. While these challenges can be overcome in industrial settings, they make little practical sense for a family barbecue or a camping stove.

Ironically, propane is one of those fuels that is already exceptionally well matched to the applications it serves. It is energy dense, easy to transport, easy to store, relatively inexpensive, and supported by mature infrastructure. Replacing it with hydrogen would solve few problems while introducing several new ones.

The broader lesson is not really about barbecues at all.

It is about expectations and reality.

Many aspects of modern life depend on energy systems that most people rarely think about. We want energy to be clean, affordable, secure, and always available. We want transitions to happen quickly, but we often do not want higher costs, new infrastructure nearby, or disruptions to familiar ways of living. The reality is that every service provided by today's energy system must ultimately be replaced by something else, and that replacement is not always straightforward.

Cleaner does not automatically mean cheaper. Sustainable does not automatically mean simpler. And replacing a fuel often means replacing entire supply chains, industries, and supporting infrastructure.

The propane cylinder beside a barbecue, caravan, boat, or rural home is therefore a useful reminder that the energy transition is not simply about replacing energy. It is also about replacing molecules. Understanding where those future molecules come from may prove just as important as understanding where future electrons come from.

In many ways, the barbecue question captures the essence of the wider transition. The challenge is not merely to build a cleaner energy system. It is to ensure that the countless services provided by the current system continue to exist in the future one. That requires pragmatism, engineering realism, and a recognition that the transition is ultimately a systems challenge rather than a simple technology swap.

Why the Transition Needs Both Nuclear and Oil & Gas

The energy transition is often presented as a clean moral binary—renewables or nothing, a total break with the past, a heroic leap into a fossil-free future. It is an inspiring story, but not an accurate one. The reality is more complex, and, in many ways, more uncomfortable: we cannot build the world we want without using the world we already have.

Yes, renewables bring extraordinary advantages. Their cost curves continue to fall, and their scalability is unmatched. But renewables alone cannot provide the constant, high-temperature, round-the-clock energy that industry depends on. That is where nuclear comes in, offering reliability, stability, and process heat that no other low-carbon source can match. And then there is oil and gas—often cast as the villain—yet indispensable for their infrastructure, capital, and operational skill. These companies understand molecules, subsurface engineering, global logistics, and long-term risk management in a way no other sector does. Hydrogen is simply another molecule in a system that they already know how to operate.

Hydrogen sits precisely at the point where these worlds converge. It requires the scale of renewables, the endurance of nuclear, and the engineering and commercial backbone of oil and gas. A realistic transition is not a replacement of one system with another—it is an evolution. By integrating what already exists with what must be built, we can develop hydrogen faster, reduce costs, ensure reliability, and create a transition that is smoother and less volatile than the simplistic "renewables-only" vision suggests.

These are some of the uncomfortable truths of the transition. We will still need oil and gas—perhaps for longer than many would like to admit—and not only for

energy. Even a fully electrified world will require petrochemicals, polymers, pharmaceuticals, lubricants, and materials that cannot simply be wished away. At the same time, nuclear energy must be part of the mix, despite political sensitivities, because no country can maintain industrial competitiveness on intermittent energy alone.

And perhaps the most uncomfortable truth of all: energy systems do not fit neatly within political boundaries or ideological preferences. The transition will require cooperation between countries that disagree on almost everything else. It demands a level of pragmatism, engineering realism, and geopolitical maturity that our current political climate rarely encourages. Climate molecules do not carry passports. Hydrogen, electrons, and CO_2 do not vote. The physics will not change to fit political slogans.

What this chapter ultimately highlights is that the energy transition is not about discarding the old system. It is about upgrading it—using the strengths of renewables, nuclear, and oil and gas together to build something better. The transition will be judged not by purity of intent, but by its ability to deliver real, reliable, global decarbonisation and energy security. To achieve that, we must build bridges, not burn them. And above all, we must remember the simplest, most universal truth of all:

We're all in the same boat.



7. Integrating the System — Hydrogen and the Wider Energy Mix

"The whole is greater than the sum of its parts." — Aristotle

Why Integration Now → Why Power-to-X

The energy system is entering a phase where its greatest challenge is no longer generation, but coordination. Over the past decade, vast amounts of renewable capacity have been added across the world. Wind and solar are now mature technologies, deployed at scale, and increasingly cost-competitive. Yet as their share of the energy mix grows, the limitations of the traditional energy system are becoming increasingly visible.

For most of the fossil-fuel era, energy systems were organised around separation. Electricity, heat, transport, and industry evolved as largely independent domains, each optimised around a dominant fuel and a predictable flow of energy. Coal, oil, and gas arrived when needed, could be stored easily, and were dispatched according to demand. The system worked because supply was controllable and abundant, and because each sector could be planned in isolation without destabilising the others.

That logic no longer holds. Renewable electricity is fundamentally different. Its output is variable, weather-dependent, and only partially predictable. Wind and solar generation do not respond to demand signals in the way fossil plants once did; instead, demand must increasingly adapt to supply. As renewable penetration rises, mismatches between when energy is produced and when it is needed become structural rather than exceptional. Periods of surplus and scarcity are no longer anomalies to be smoothed at the margins — they are defining features of the system.

In a siloed energy architecture, these mismatches quickly become problems. Excess electricity leads to curtailment, wasted investment, and grid congestion. Shortfalls increase price volatility and strain reliability. Each sector, acting inde-

pendently, lacks the flexibility to respond effectively to conditions elsewhere in the system. The result is an energy landscape that is technically capable of deep decarbonisation, but operationally fragile.

Electricity is the most efficient carrier of energy — but not the most versatile. This is the integration challenge. The task ahead is not simply to replace fossil fuels with clean ones, but to link energy carriers, sectors, and time scales into a coherent whole. Electricity must be able to move beyond the grid. Energy must be stored not just for hours, but for days, weeks, and seasons. Industrial demand, transport fuels, and heating systems must become participants in system balancing rather than passive consumers.

Over the past decade, this challenge has often been described using the language of sector coupling — the idea that electricity, heat, transport, and industry must be increasingly connected. While conceptually useful, that terminology has remained abstract, and at times obscure. What matters in practice is not the coupling of sectors as an end in itself, but the concrete mechanisms through which energy is converted, stored, and re-used across the system.

This is where Power-to-X becomes the more useful frame.

Power-to-X describes a family of processes that convert electrical energy into other usable forms: hydrogen, synthetic fuels, chemicals, or heat. Unlike broader integration concepts, it speaks directly to what actually happens in the system — electrons being transformed into molecules, stored, transported, and ultimately used where electricity alone cannot easily go. Power-to-X is not a technology in its own right, but a system function: it is how surplus electricity acquires reach, duration, and flexibility.

Crucially, Power-to-X shifts the focus from individual assets to system behaviour. It reframes renewable electricity not as an end product, but as a primary input into a wider energy network. In doing so, it allows the energy system to operate across multiple dimensions simultaneously: short-term and long-term balancing, local and global energy flows, direct electrification and indirect decarbonisation.

The urgency of this shift is not ideological — it is practical. As renewable capacity continues to expand, systems that rely solely on electrification will face growing constraints in storage, grid reinforcement, and peak management. Power-to-X does not compete with electrification; it completes it. It provides the means to absorb excess generation, stabilise the system during scarcity, and extend decarbonisation into sectors that electricity alone cannot readily serve.

Integration, then, is no longer a theoretical ambition. It is an operational necessity. And Power-to-X is the language through which that necessity is increasingly being met.

Power-to-X as the Language of System Integration

Having established why integration is necessary, the question becomes how it is achieved in practice. Power-to-X is the language through which that integration becomes operational. It provides a practical way to describe how renewable electricity can be converted into other usable forms, acquire flexibility in time and space, and support sectors that are difficult to serve through direct electrification. Rather than treating electricity, fuels, and chemicals as separate domains, Power-to-X frames them as interconnected expressions of the same underlying energy resource.

At its simplest, Power-to-X refers to the conversion of electrical energy into another form — the "X" — that is better suited to a particular function. That X may be hydrogen, a synthetic liquid fuel, a chemical feedstock, or heat. What matters is not the label, but the role these conversions play in allowing energy to flow through the system in forms better suited to different functions.

This marks a shift in perspective. Traditional energy systems were built around fuels that arrived already packaged for use. Electricity, by contrast, was treated as an endpoint — generated and consumed in near real time. In a renewable system, its value increasingly depends on what can be done with it when it is not immediately needed.

Power-to-X turns surplus electricity into optionality.

The most familiar example is Power-to-Hydrogen, where electricity drives electrolysis to split water into hydrogen and oxygen. Hydrogen, in this context, is not primarily a competing fuel, but a way of storing and redistributing electrical energy. It can absorb excess generation, sit in storage for extended periods, and later be used to produce heat, power, or feedstock when conditions demand. In doing so, it allows electricity to be stored and used over days, weeks, or even seasons.

Yet hydrogen is only one branch of the Power-to-X family. In many cases, hydrogen itself is not the most practical endpoint. Its physical properties can limit where and how it is used, which is where Power-to-Liquids and Power-to-Chemicals enter the picture.

Through additional conversion steps, hydrogen can be combined with carbon or nitrogen to produce molecules such as methanol, ammonia, or synthetic hydrocarbons — including synthetic aviation fuels (e-SAF). These molecules trade efficiency for compatibility. They are easier to store, easier to transport, and often directly compatible with existing engines, ships, pipelines, tanks, and industrial processes. As discussed later, these pathways sacrifice efficiency, but they dramatically expand the number of places where renewable electricity can displace fossil fuels.

Seen this way, Power-to-X is not a menu of competing technologies, but a toolkit for system design. Each pathway offers a different balance between efficiency, flexibility, storage duration, and infrastructure reuse. The choice of X is not about maximising conversion efficiency in isolation, but about maximising system value under real-world constraints.

Power-to-X also reframes how we think about energy flows. Instead of a one-directional chain from generation to consumption, the system begins to behave more like a network. Electricity can flow directly into end uses when conditions allow. When it cannot, it can be converted into molecules, stored, transported, and reintroduced elsewhere in the system. Energy moves not just across sectors, but across time and geography.

Importantly, this does not diminish the role of direct electrification. Electrification remains the first and best option wherever it is technically and economically viable. Power-to-X exists to serve the remainder — the applications where those constraints become binding. The two approaches are complementary by design: electrification maximises efficiency; Power-to-X maximises reach.

Framing integration in Power-to-X terms also clarifies responsibility. Grid operators, industrial players, fuel producers, and policymakers can engage with concrete conversion pathways rather than abstract coordination goals. Questions become practical: where does surplus electricity appear, what form should it take to be useful elsewhere, and which conversions justify their losses through system benefits?

As renewable penetration increases, these questions become unavoidable. A system built solely around electrons will encounter physical and economic limits. A system that treats Power-to-X as an integral function gains degrees of freedom: more ways to balance, more ways to store, and more ways to decarbonise.

Power-to-X, then, is not a detour from electrification, nor a hedge against its success. It is the mechanism by which electrification scales beyond the grid — turn-

ing renewable electricity into a flexible, system-wide resource capable of supporting a fully decarbonised energy economy.

From Hydrogen to System Molecules

Hydrogen occupies a unique position in the Power-to-X landscape. It is both a destination and a departure point — an energy carrier in its own right, and the gateway to a wider family of system molecules. Understanding this dual role is essential to making sense of how integration unfolds in practice.

In much of the public debate, hydrogen is presented as an end state. In an integrated energy system, however, its most important function is often as an intermediate — the first step in converting renewable electricity into forms that can travel further, last longer, or fit more easily into existing infrastructure.

This distinction reframes how hydrogen is used. Its physical characteristics make it well suited to some applications, particularly large-scale storage and industrial use, but less convenient for others. Treating hydrogen as the final product in every context risks forcing it into roles where it is not the best system solution.

When hydrogen is understood instead as a platform molecule, its role expands. It becomes the starting point for a wider set of pathways: it can be combined with carbon to form liquid fuels and chemical feedstocks — including synthetic aviation fuels (e-SAF), which are critical for decarbonising long-haul aviation where direct electrification is not viable. It can also be bound to carrier liquids and released again where needed. Each pathway adds complexity, but each solves a specific system constraint.

This is the logic that leads from hydrogen to system molecules.

Methanol, ammonia, and liquid organic hydrogen carriers are not alternatives to hydrogen; they are extensions of it. They exist because real energy systems must contend with geography, infrastructure, safety, and existing assets. Liquids are easier to transport than gases, and global systems are already built around molecules rather than electrons. System molecules make hydrogen usable within the infrastructures and trade systems that already exist.

As discussed later, these pathways involve additional conversion losses. However, the relevant question is not whether conversion should occur, but where it enables decarbonisation that would otherwise remain out of reach.

The central question, therefore, is not whether hydrogen should be used directly or converted further, but where each option delivers the greatest system value.

In some contexts, hydrogen remains the right endpoint. In others, converting it into a different molecule unlocks applications that would otherwise remain dependent on fossil fuels.

The sections that follow explore these molecules not as fuels in competition, but as tools of integration — each chosen not for theoretical elegance, but for its ability to solve real constraints in a complex, evolving energy system.

Hydrogen — The System Molecule

Within a Power-to-X framework, hydrogen's value lies in what it enables. It is the first and simplest molecular carrier produced from renewable electricity — and it is this position at the boundary between power and matter that gives hydrogen its systemic importance.

Hydrogen's defining strength is temporal flexibility. Unlike batteries, which excel at balancing short-term fluctuations, hydrogen allows energy to be shifted across much longer time horizons. Electricity generated during periods of surplus can be converted into hydrogen and stored for days, weeks, or entire seasons. In systems with high renewable penetration, this ability to move energy through time becomes critical, allowing excess generation to be captured rather than curtailed and scarcity to be met without reverting to fossil backup.

This long-duration storage role operates at the scale of the system rather than the asset. Seasonal mismatches — summer solar surpluses and winter heating demand, for example — are difficult to resolve with batteries alone. Hydrogen provides a buffer that does not need to respond instantly, but must be available reliably when conditions demand it.

Hydrogen also introduces spatial flexibility. As a molecule, it can be transported in ways electricity cannot — through pipelines, ships, or trucks — enabling energy produced in one region to be used in another. This becomes increasingly important in a world where renewable resources are unevenly distributed, allowing regions rich in wind or solar to convert local electricity into hydrogen and move energy across distance without relying solely on grid expansion.

Together, these roles make hydrogen a balancing layer within the energy system. It absorbs excess electricity when supply is high, stores it over extended periods, and reintroduces it when and where it is needed. In doing so, it reduces curtailment, stabilises variability, and extends the operational range of renewable energy.

Hydrogen's role is therefore not universal, but selective. It is deployed where storage duration becomes too long, distances too great, or system constraints too binding for direct electrification alone. Used in this way, it strengthens the overall architecture of an integrated energy system.

In this sense, hydrogen is best understood as an enabling layer — the first molecular step that allows renewable electricity to be stored, transported, and redeployed as a flexible, system-wide resource.

Methanol — The Quiet Integrator (Compatibility)

If hydrogen is the gateway molecule of a Power-to-X system, methanol is often where that system becomes practical. It rarely attracts the same attention as hydrogen itself, yet in many real-world contexts methanol is the molecule that allows renewable electricity to be absorbed into existing industrial and transport systems.

Methanol's defining strength is compatibility. As a liquid at ambient conditions, it can be stored, transported, and handled using infrastructure that already exists. Tanks, pipelines, ships, and engines designed for liquid fuels can often be adapted rather than replaced. In a transition where speed, cost, and continuity matter, this characteristic carries enormous system value.

From a Power-to-X perspective, methanol packages hydrogen — and crucially carbon — into a form that fits the modern economy. Produced by combining hydrogen with captured carbon dioxide, it becomes both an energy carrier and a chemical feedstock, allowing renewable electricity to displace fossil inputs not only in energy use, but deep within industrial value chains.

Shipping illustrates this clearly. Direct electrification is impractical for long-distance maritime transport, and hydrogen presents storage and handling challenges. Methanol, by contrast, aligns with operational reality: it can be bunkered as a liquid, used in adapted engines or fuel cells, and integrated into port infrastructure with relatively modest changes.

As discussed earlier, these pathways trade efficiency for system compatibility. Their value lies in enabling decarbonisation in sectors where alternatives are limited, not in maximising conversion efficiency in isolation.

Methanol also introduces geographic flexibility. It can be produced wherever low-cost renewable electricity and a source of carbon are available, then transported globally using established trade routes. In doing so, it allows renewable

energy to enter the global economy in a form that is already widely understood and utilised.

Methanol's role is therefore not universal, but pragmatic. It allows renewable electricity to integrate into existing systems without requiring wholesale replacement — accelerating transition by working with the world as it is.

Ammonia — Hydrogen at Global Scale (Scale)

Ammonia occupies a distinctive position within the Power-to-X landscape. Where methanol excels at compatibility, ammonia's strength lies in scale. It is one of the few hydrogen-based molecules already produced, transported, and traded globally in enormous quantities.

At its core, ammonia functions as a hydrogen carrier. By combining hydrogen with nitrogen, it forms a dense, stable molecule that can be stored and shipped far more easily than hydrogen itself. Decades of experience in the fertiliser industry have created a mature global infrastructure of ports, pipelines, storage systems, and safety protocols that Power-to-Ammonia pathways can directly leverage.

This existing infrastructure enables energy to move at scale. Renewable electricity can be converted into hydrogen, synthesised into ammonia, and transported across continents using systems that already exist. In a world of unevenly distributed renewable resources, this allows energy to flow from regions of abundance to regions of demand without requiring vast new electrical interconnections.

Ammonia's primary system value lies in this ability to move and store energy at scale. While its direct use as a fuel may grow in specific applications, its broader role is as a transport and storage medium — enabling long-distance, long-duration balancing across the energy system.

As with other Power-to-X pathways, this comes with trade-offs. But as discussed earlier, the relevant comparison is not with idealised efficiency, but with the fossil systems these pathways replace. Where ammonia displaces imported coal, oil, or gas, its system value becomes clear.

Ammonia also provides optionality. It can be transported as a carrier, then cracked back into hydrogen, used directly in industrial processes, or in some cases employed as a fuel. This flexibility allows energy systems to evolve over time, adapting as infrastructure and demand develop.

In an integrated energy system, ammonia functions as a backbone molecule — enabling renewable energy to operate at global scale.

LOHC — Hydrogen Where Infrastructure Is Missing (Flexibility / Gaps)

Liquid Organic Hydrogen Carriers (LOHCs) occupy a more specialised yet important niche within the Power-to-X framework. Their value lies not in scale or efficiency, but in their ability to move hydrogen through environments where dedicated infrastructure is absent.

LOHCs bind hydrogen chemically to an organic liquid, allowing it to be transported and stored under ambient conditions. The hydrogen can later be released at the point of use. This approach trades efficiency for logistical simplicity.

That trade-off defines their role. Unlike compressed or liquefied hydrogen, LOHCs avoid high pressures, extremely low temperatures, and complex handling requirements. They can be managed using conventional liquid fuel infrastructure — tanks, trucks, rail, and ships — making them particularly useful in early-stage or distributed markets.

LOHCs are well suited to situations where hydrogen demand is geographically dispersed or infrastructure is incomplete: supplying industrial sites, refuelling stations, or remote energy systems. In these contexts, the alternative is often not a more efficient pathway, but no viable hydrogen pathway at all.

As discussed earlier, these systems involve additional conversion losses, which limits their role in large-scale, high-throughput energy transport. They are not a substitute for pipelines where pipelines are feasible, nor for ammonia where global transport is required.

Their value instead lies in flexibility. LOHCs enable incremental deployment, allowing hydrogen systems to develop without requiring full infrastructure build-out from the outset. They reduce the “all-or-nothing” nature of hydrogen adoption.

Within an integrated energy system, LOHCs function as connective tissue. They do not dominate energy flows, but they close gaps — extending the reach of hydrogen into environments where it would otherwise remain inaccessible.

Losses as the Cost of Reach

Every integrated energy system must confront an uncomfortable reality: converting energy from one form to another always involves losses. Electricity turned into hydrogen, hydrogen converted into methanol or ammonia, molecules transported and reconverted into heat, motion, or power — each step reduces the amount of usable energy that reaches the final application. In isolation, these losses can appear damning. In context, they are unavoidable.

Much of the criticism directed at hydrogen and Power-to-X pathways focuses on this inefficiency. Comparisons are drawn between direct electrification and multi-step conversion chains, often showing stark differences in end-to-end efficiency. These comparisons are technically correct — and systemically incomplete. They assume that direct electrification is always a feasible alternative. In many of the applications Power-to-X is designed to serve, it is not.

In aviation, for example, synthetic fuels such as e-SAF may represent the only scalable pathway to deep decarbonisation — even if the energy cost of producing them is high.

Importantly, this increase in fuel cost does not translate directly into proportional increases in end-user cost. Fuel represents only one component of an airline's operating expenses. Even where synthetic fuels are significantly more expensive to produce, the impact on ticket prices is typically far smaller — often amounting to a modest single-digit percentage increase. In other words, what appears costly at the level of fuel production becomes far more manageable when viewed across the full system.

Efficiency is only one dimension of system performance. Time, distance, scale, and compatibility matter just as much. Electricity is exceptionally efficient when it can be used directly, but it is difficult to store over long periods, challenging to transport over intercontinental distances, and constrained by grid capacity and reinforcement costs. Power-to-X pathways exist precisely because these constraints are real and persistent.

Seen through this lens, conversion losses are not a flaw in the system — they are the price paid to access new dimensions of flexibility. Hydrogen allows energy to move through time. Methanol and ammonia allow it to move through existing industrial and transport systems. LOHCs allow it to reach places where dedicated infrastructure does not yet exist. Each pathway sacrifices efficiency to gain reach, reliability, or speed of deployment.

The critical design question, therefore, is not how to eliminate losses, but when they are justified. A simple principle emerges convert energy as few times as possible, and only as far as necessary. If a demand can be met efficiently with direct electricity, the system should stop there. Introducing hydrogen or synthetic fuels where electrification is viable adds cost and complexity without delivering commensurate benefit. Discipline matters.

Where that discipline is applied, Power-to-X becomes a targeted tool rather than a blunt instrument. Hydrogen is reserved for long-duration storage, high-temperature processes, and grid support. Methanol and ammonia are deployed where liquid fuels or global transport are unavoidable. LOHCs fill logistical gaps rather than attempting to dominate energy flows. Losses are accepted only where they unlock decarbonisation that would otherwise remain out of reach.

This perspective also reframes the comparison with fossil systems. Fossil fuels appear efficient at the point of use, but their apparent simplicity masks externalities and system costs that are rarely accounted for in efficiency debates. Power-to-X losses are visible and measurable; fossil losses are embedded and deferred. When conversion pathways replace coal, oil, or gas in sectors that lack other options, the net system outcome remains strongly positive, even if the energy arithmetic looks unfavourable in isolation.

Ultimately, the success of an integrated energy system will not be judged by theoretical efficiency alone, but by its ability to deliver reliability, affordability, and deep decarbonisation at scale. Losses are the cost of reach — across time, across distance, and across sectors. When managed deliberately and applied selectively, they become not a liability, but a design choice that enables the energy transition to proceed where it otherwise could not.

A Living Power-to-X System — Chile

Chile offers one of the clearest real-world demonstrations of Power-to-X at system level. Its significance lies not in scale — not yet — but in alignment. Geography, resource availability, and strategic intent converge in a way that makes integration unavoidable.

Chile's renewable energy endowment is extreme. In the north, the Atacama Desert receives some of the highest solar irradiation on the planet. In the far south, Patagonia experiences some of the world's most consistent and powerful winds. These resources are vast but remote — distant from population centres, weakly connected to national grids, and disconnected from global energy demand. Exporting electricity directly is neither practical nor economical.

This constraint defines the system. Renewable electricity is treated not as an endpoint, but as a feedstock — converted into hydrogen and then into fuels and chemicals that can be stored, transported, and exported.

One of the earliest demonstrations of this approach is the Haru Oni project in southern Chile. Its value lies not in scale, but in integration. Wind power from Patagonia feeds electrolysers that produce hydrogen, which is then combined with captured carbon dioxide to synthesise e-methanol and e-gasoline — fuels that can be handled using existing infrastructure.

What matters is the system it represents. Haru Oni links variable generation, hydrogen production, carbon capture, fuel synthesis, and export logistics into a single operational chain. This is Power-to-X not as a concept, but as a working system.

Time is central to this configuration. Wind output fluctuates by the hour and season, and electrolysers respond accordingly, absorbing energy when it is available. Fuel synthesis and storage smooth this variability further, decoupling production from use. Energy that would otherwise be curtailed becomes an exportable product, with value realised over longer timescales.

Geography reinforces the model. Chile's distance from major energy markets limits its electricity exports but favours synthetic fuels. Liquids can be transported using existing maritime routes, allowing renewable energy to enter global trade in molecular form.

As discussed earlier, these pathways involve additional conversion losses. However, without them, much of Chile's renewable potential would remain stranded. With them, renewable electricity gains reach, duration, and economic relevance.

Chile's trajectory illustrates a broader point: integrated energy systems are built incrementally. Projects like Haru Oni test interfaces, expose constraints, and build operational knowledge. They are not templates to be copied wholesale, but proofs that the system logic works.

As Chile scales from pilots to industrial clusters, the same pattern will hold. Electricity becomes hydrogen; hydrogen becomes fuels and chemicals; value is realised far beyond the point of generation. In doing so, Chile offers a clear view of how Power-to-X enables integration — not through theoretical optimisation, but through practical application.

What This Means for the Energy System as a Whole

By this point, one conclusion should feel unavoidable: energy systems built for control must give way to energy systems built for adaptation. High-renewable futures are not defined by a single optimal pathway, but by the ability to respond — to weather, to demand, to geography, and to constraint. Power-to-X is less a technology choice than an admission of that reality.

What changes most is not the destination, but the design logic. In a fossil system, fuels arrived ready to use and the task was distribution. In a renewable system, energy arrives when it arrives, and the task becomes transformation. The central engineering challenge shifts from extraction to orchestration — deciding when energy should flow as electrons, when it should pause as molecules, and when it should be reshaped to fit existing infrastructure.

This has implications well beyond technology. Planning horizons lengthen. Value moves from individual assets to interfaces. Success depends less on perfect efficiency than on resilience, optionality, and the ability to learn under real operating conditions. Projects like Haru Oni matter not because they scale today, but because they teach the system how to behave tomorrow.

It also reframes the role of compromise. Conversion losses, hybrid solutions, and imperfect fits are not signs of failure; they are the fingerprints of a system in transition. The danger lies not in accepting those imperfections, but in pretending they can be engineered away entirely.

Perhaps most importantly, this perspective restores humility to the transition. No single pathway will dominate everywhere. Systems will evolve unevenly, shaped by local resources, existing assets, and institutional capacity. Power-to-X provides a way to navigate that diversity without fragmenting the whole.

The chapters that follow step into that diversity directly — into the places where these system choices become concrete, contested, and consequential. Integration, here, is no longer a concept. It is a practice, shaped by economics as much as engineering.

8. Understanding LCOH (Levelized Cost of Hydrogen)

"An expert is someone who has made all the mistakes possible in a very narrow field." — Niels Bohr

"In God we trust. All others must bring data." — W. Edwards Deming

The Cost of Something New

In northern Sweden, just outside the town of Boden, the land stretches wide and unhurried. Forests dominate the horizon. Rivers cut clean lines through rock and snow. For generations, this has been a place defined by raw materials—iron ore pulled from the earth, timber harvested and shipped south, energy flowing outward to feed industries elsewhere. It is a landscape shaped by extraction, distance, and patience.

Today, something new is taking shape here.

From the road, the construction site does not look revolutionary. Cranes arc slowly across the sky. Concrete foundations spread outward in measured grids. Steel frames rise, one section at a time. Trucks arrive and leave with predictable regularity. There is no spectacle, no dramatic architecture announcing a break with the past. And yet, beneath the surface, this project is quietly attempting one of the most ambitious industrial transformations of the modern era.

This is Stegra—formerly known as H2 Green Steel—and its goal is deceptively simple to state: to produce steel without coal, using green hydrogen instead of coke, and to do so at industrial scale. Not as a pilot. Not as a demonstration. But as a fully integrated, commercially viable steel plant designed to compete head-to-head with conventional production.

Steel is one of the hardest materials in the world to decarbonise. Every tonne produced today carries with it the weight of carbon emissions—roughly two tonnes of CO_2 per tonne of steel. The chemistry is unforgiving. Iron ore must be reduced, oxygen stripped away, and for more than a century coal has been the cheapest, most reliable way to do it. The global steel industry is therefore not just large; it is deeply entrenched. Its processes, supply chains, and cost structures are optimised down to the last cent.

Stegra's bet is that this optimisation can be rewritten.

Instead of coal, the plant will use hydrogen-based direct reduced iron (DRI). Iron ore pellets will be exposed to hydrogen, producing sponge iron and water vapour instead of carbon dioxide. That iron will then be melted and refined using fossil-free electricity. The hydrogen itself will be produced on site, through electrolysis powered by northern Sweden's abundant renewable energy—hydropower, wind, and a grid already among the cleanest in the world.

When fully built, the Boden facility is designed to include hundreds of megawatts of electrolyser capacity, making it one of the largest integrated green hydrogen systems in Europe. The ambition is not marginal decarbonisation, but structural change: a new template for how heavy industry can function in a low-carbon world.

From the outside, the logic is compelling. Sweden has cheap, clean electricity. It has iron ore supply chains. It has political stability and long-term industrial policy. If green steel can work anywhere, it should work here.

But logic alone does not make projects bankable.

This is where cost enters the story.

Hydrogen, for all its promise, is not cheap—at least not yet. Electrolysers require capital. Renewable power must be built at scale. Infrastructure must be oversized to handle variability. And in steelmaking, margins are thin and competition global. Every additional euro per tonne matters.

Stegra's existence, therefore, is inseparable from a question that sits at the heart of the hydrogen economy: what does it really cost to change the system?

The answer is not found in a single number. It emerges from trade-offs.

Hydrogen costs more than coal, but avoids carbon pricing exposure. Renewable electricity requires upfront investment, but offers long-term price stability. Building integrated hydrogen and steel infrastructure is capital-intensive, but

removes dependence on volatile fossil fuel markets. The economics only make sense when viewed across decades rather than quarters.

It is also important to recognise that today's comparison is rarely made on a perfectly level playing field. In many parts of the world, grey hydrogen continues to benefit from the absence—or only partial application—of meaningful carbon pricing. As a result, its market price often reflects only the cost of production, rather than the wider environmental cost of the associated emissions. This can make green hydrogen appear disproportionately expensive, when in reality part of the difference arises because the external costs of carbon remain only partially accounted for. As carbon pricing mechanisms mature and become more widely adopted, the economic gap between grey and low-carbon hydrogen is likely to narrow, not solely because green hydrogen becomes cheaper, but because the market begins to reflect a more complete picture of the true cost of production.

This is not an incremental upgrade. It is a re-pricing of risk.

At Stegra, the strategy has never been to wait for hydrogen to become cheap before acting. The strategy has been to build scale early, learning faster than competitors, locking in supply chains, and positioning the plant for a world where carbon constraints are no longer optional. Long-term offtake agreements, public financing, and strategic partnerships all play a role in bridging the gap between today's costs and tomorrow's equilibrium.

In that sense, Stegra is not just a steel project. It is a living case study in how the levelized cost of hydrogen intersects with real industry.

The hydrogen produced here will not be traded on spot markets or piped across borders. It will be consumed immediately, continuously, predictably, turned into iron, then steel, then cars, buildings, and infrastructure. This kind of demand profile is exactly what hydrogen economics need: steady load, high utilisation, minimal transport losses. The closer hydrogen production is to hydrogen use, the less cost leaks out of the system.

Seen this way, the Boden site becomes a lens through which LCOH stops being abstract. Cost is no longer just euros per kilogram. It becomes cost per tonne of steel avoided, cost per supply chain secured, cost per regulatory risk mitigated.

Walking the perimeter of the site today, it is easy to forget how unconventional this all still is. The cranes move slowly. The foundations cure. Snow settles and melts. Progress is measured in months, not headlines. Yet the choices embedded in this concrete will shape industrial economics for decades.

If Stegra succeeds, it will prove that hydrogen does not need to be cheap everywhere—only cheap enough where it matters most. It will show that early movers can absorb higher costs in exchange for strategic advantage. And it will demonstrate that the true cost of hydrogen cannot be separated from the cost of inaction.

This is why Stegra belongs at the opening of a chapter on LCOH.

Because levelised cost is not just about accounting. It is about intent. About deciding which costs we are willing to pay upfront, and which costs we are content to pass to the future. Coal looks cheap until carbon is priced. Hydrogen looks expensive until the system around it is redesigned.

In northern Sweden, that redesign is underway.

From a distance, it looks like any other industrial construction site. But beneath the cranes and concrete lies a quiet recalculation — one that asks not whether hydrogen is expensive today, but whether continuing the old way is affordable at all.

And that is where every serious discussion of hydrogen cost must begin.

What makes Stegra such a compelling case is not only what it is attempting, but the sequence in which it is attempting it. The ambition is undeniable. Producing green hydrogen at prices acceptable to the steel industry is an even greater ambition. Of all the sectors that must eventually be decarbonised with hydrogen, steel sits among the most demanding. It is capital-intensive, fiercely cost-competitive, and historically optimised around the cheapest possible inputs.

My concern is not with the destination, but with the order of travel.

As I discuss later, the hydrogen transition may be more robust if it progresses from sectors that can tolerate higher energy costs toward those that cannot. In that framing, steel would normally appear later in the sequence, once hydrogen pricing has been proven in less cost-sensitive applications. Stegra, by contrast, is attempting to pull that future forward.

That boldness carries both promise and risk.

The question, then, is not whether Stegra succeeds in every metric it has set for itself, but how its outcome is interpreted. Even a project that delivers partial success — learning, scale-up experience, and technical validation — can be invaluable to those actively building the transition. Yet there is also a wider audience watching. If such a high-profile effort is perceived externally as falling short, it risks

being misused by those already sceptical of hydrogen as evidence that the transition should slow, rather than adapt.

In that sense, projects like Stegra do not exist in isolation. They carry symbolic weight beyond their balance sheets. Their outcomes will shape narratives as much as markets.

It is for this reason that cost becomes so central.

This brings us back to the levelised cost of hydrogen. There are many variables that influence the energy transition, but few exert as much leverage across so many sectors as this one. If there is a single concept in this book that underpins everything else — technology choices, financing structures, infrastructure deployment, and industrial adoption — it is this.

Understanding how hydrogen reaches a price that markets can accept is not a technical footnote. It is the foundation. And to understand that, we must first understand how the levelised cost of hydrogen is built.

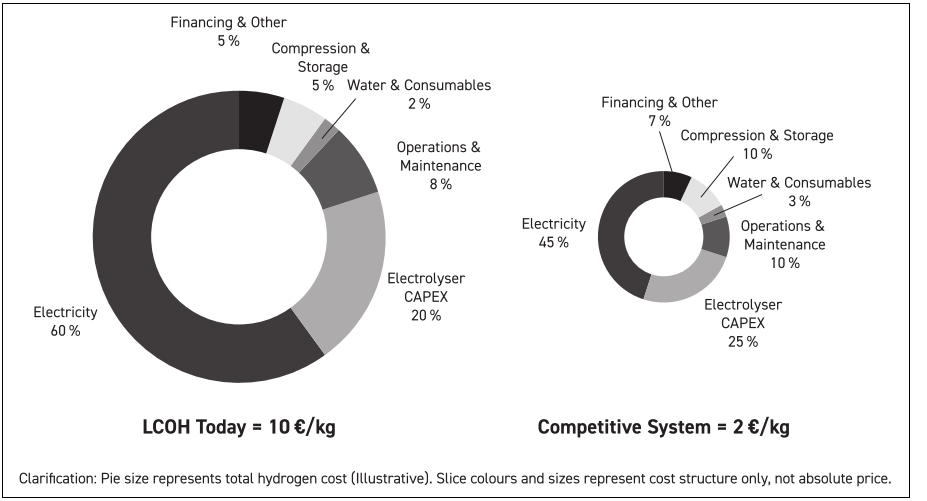
Stegra, then, should not be read as a verdict on hydrogen's readiness, but as a stress test of its economics. It is an attempt to compress decades of industrial evolution into a single project, exposing every assumption about cost, scale, power prices, financing, and risk. In doing so, it forces uncomfortable but necessary questions into the open — not only about whether hydrogen can decarbonise heavy industry, but about what must be true for it to do so competitively. Those questions do not diminish the project's ambition; they define its importance.

What Is LCOH — and Why It Matters

At its simplest, LCOH answers one question:

"What does it cost, on average, to produce one kilogram of hydrogen over the lifetime of an asset?"

Formally, LCOH aggregates all of the major costs associated with producing hydrogen over the lifetime of an asset and spreads them across the total hydrogen output. This includes upfront capital expenditure, ongoing operating costs, electricity or fuel input, asset lifetime assumptions, utilisation rates, and the cost of financing. Since all these elements are folded into a single number, LCOH is a powerful comparison tool — but also a potentially misleading one. Two projects using the same technology can report very different LCOH values depending on where they are built, how often they operate, and how their risks are financed.

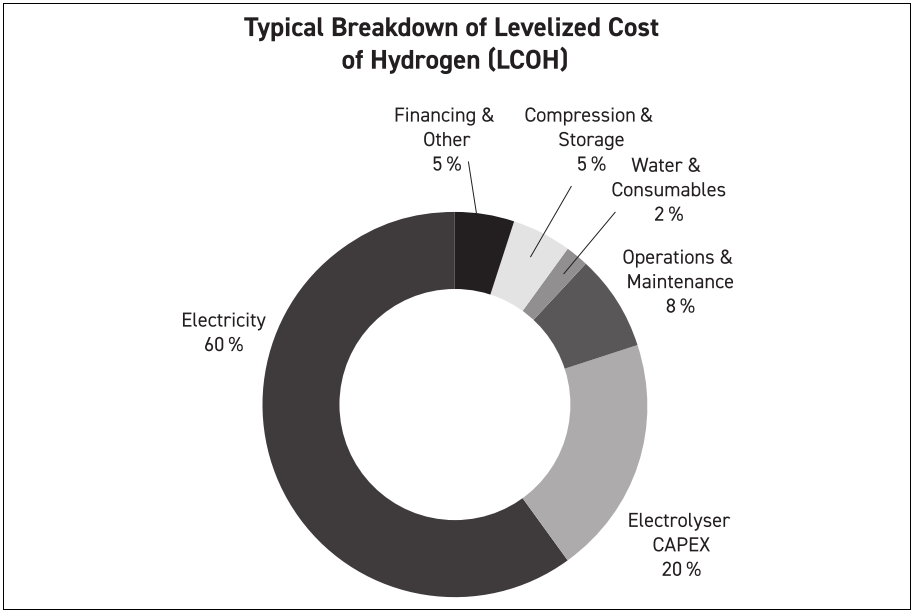


This figure shows why the levelised cost of hydrogen is best understood as a system outcome rather than a single cost input. The individual components of hydrogen cost — electricity, capital, operations, and supporting infrastructure — are broadly the same in every project. What changes is not whether electrolysers are operated intensively, but how efficiently the surrounding system supports them over time. Early projects often run electrolysers hard when power and offtake are available yet still carry high costs due to first-of-a-kind capital, conservative financing assumptions, limited integration, and immature demand structures. As systems mature, the same components produce more hydrogen over their lifetime, spreading fixed costs more effectively and reducing overall LCOH.

This does not imply that individual inputs stand still. Over time, electricity is expected to become cheaper and more predictable in many regions, electrolyser performance and durability will continue to improve, and a host of incremental gains — better system design, improved power electronics, higher operating temperatures, smarter control strategies, and learning-by-doing across supply chains — will all contribute to lower costs. These improvements matter. But their impact is magnified, or muted, by the system they operate within. The figure separates cost structure from cost scale to make this visible: the slices show what costs exist, while the size of the system shows how effectively those costs are ultimately converted into hydrogen.

The Five Core Cost Drivers of LCOH

Nearly every hydrogen cost model, regardless of complexity, boils down to five dominant variables.



The figure above shows a simplified breakdown of the levelised cost of hydrogen for a typical large-scale green hydrogen project. While the exact shares vary by location, technology choice, and system design, the hierarchy is remarkably consistent. Electricity dominates total cost, often accounting for around 60 per cent of LCOH. Capital expenditure on electrolysers follows, with all other components — operations, compression, storage, water, and financing — playing secondary but still meaningful roles.

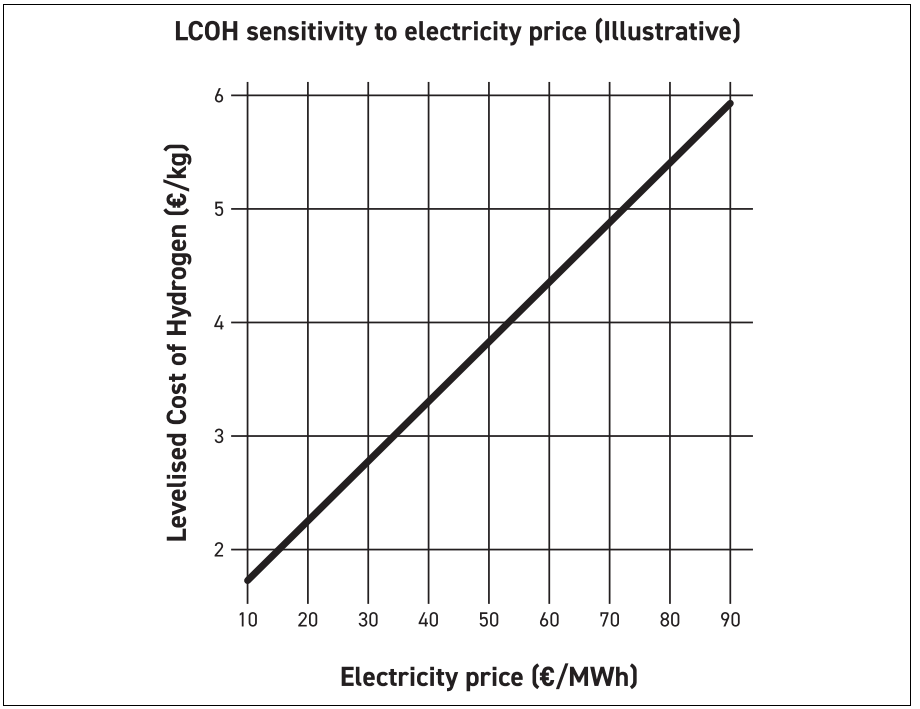
The implication is clear. Meaningful reductions in hydrogen cost depend less on marginal efficiency improvements in individual components and more on securing low-cost electricity at high utilisation, supported by stable, well-integrated system design.

Electricity Price — The Primary Driver

For green hydrogen, electricity is not just another input cost; it is the foundation upon which everything else rests. In most projects, electricity accounts for between 50 and 70 per cent of total LCOH, making it the single largest determinant of whether a project is viable or not.

A simple rule of thumb helps to make this concrete. Producing one kilogram of hydrogen through electrolysis typically requires around 50–55 kilowatt-hours of electricity, depending on the efficiency of the system. Here, efficiency should be

understood at the system level, including both the electrolyser stack and its balance-of-plant requirements, rather than as stack efficiency alone. This distinction matters when comparing technologies or manufacturers, as published efficiency and power-consumption figures do not always use the same system boundaries. Once this relationship is understood, the sensitivity of hydrogen cost to power prices becomes immediately clear. At an electricity price of €20 per megawatt-hour, the cost of power alone contributes roughly €1.00–1.10 per kilogram of hydrogen. At €50 per megawatt-hour, that figure rises to around €2.50–2.75 per kilogram, before any capital, operating, or financing costs are added. At €80 per megawatt-hour, electricity costs overwhelm the economics, rendering most green hydrogen projects uncompetitive without substantial policy support, regardless of how inexpensive the electrolyser itself may be.



This steep sensitivity explains why debates about hydrogen economics so often collapse into debates about power prices. But it also explains why average electricity prices can be misleading. Hydrogen does not need cheap electricity in general; it needs cheap electricity for many hours of the year. A project supplied with low-cost

power for only a fraction of its operating time will still struggle, because fixed costs must be recovered whether the electrolyser runs or not. By contrast, a system that combines low-cost electricity with high utilisation can tolerate higher capital costs and still achieve a competitive LCOH.

Seen in this light, electricity price and utilisation are inseparable. Cheap power without sustained operation produces disappointing results, just as high utilisation with expensive power locks in high costs. The projects that succeed are those that align both — securing electricity that is not only low-cost, but reliably available across thousands of operating hours each year.

In practice, however, the electricity price that appears in a model is rarely the price a project experiences — because procurement structures and operating strategies reshape it.

Electricity Procurement and Operations — Where LCOH is Won or Lost

PPAs — The Contractual Bridge to a Renewable Future

Producing green hydrogen depends first and foremost on securing renewable electricity — reliably, affordably, and at the right times. Power Purchase Agreements, which began as long-term hedging tools for the electricity sector, have become the backbone of early hydrogen projects.

A PPA is not simply a contract for power; it is a contract for stability. It defines price profiles, production windows, availability risks, imbalance charges, and sometimes even grid behaviour. A hydrogen project may sign a fixed-volume PPA, forcing the electrolyser to buy electricity even during high-price periods. Or it may sign a pay-as-produced contract, accepting variability in exchange for lower cost. A shaped PPA offers smoothness but at a premium.

The challenge is that hydrogen sits at the intersection of two volatile systems: the power market and the physical grid. PPAs allow developers to navigate this volatility by anchoring one side of the equation. But PPAs are, at best, an approximation of the system hydrogen will ultimately need. They are a bridge — nothing more, nothing less — but without them, the industry cannot move.

Electrolyser Loading Philosophies — The Heartbeat of the Transition

Once electricity is contracted, the next question becomes: how should the electrolyser run?

This is one of the most decisive variables in hydrogen economics. Electrolysers can operate at baseload, optimising utilisation but exposing themselves to high-cost electricity. They can run flexibly, chasing low-price and high-renewable hours, but at the cost of lower annual output. Alternatively, they can adopt hybrid strategies that balance utilisation with price sensitivity, using storage to buffer the mismatch between production and demand.

There is no universal answer. In regions with abundant cheap renewables, flexibility wins. In industrial clusters with constant demand, baseload may be necessary. In places with seasonal patterns, a hybrid model paired with cavern storage makes the most sense.

Loading decisions are also influenced by the electrolyser's own efficiency curve. Efficiency is not constant across the operating range: depending on the technology and system design, an electrolyser may operate more efficiently at certain partial-load points than when driven at maximum output. Plant operating philosophies can therefore be designed around these load curves, balancing electricity price, hydrogen demand, utilisation, efficiency, and equipment life rather than simply maximising production whenever power is available.

This transition period is where these philosophies collide. Electrolysers must be designed for flexibility even if they initially operate as baseload assets. Developers must embrace variability while aspiring to stability. And digital optimisation is not optional — it is the only way to navigate the volatility of a system still coming into existence.

Curtailment and Negative Prices — The Hidden Catalyst of Early Hydrogen Economics

One of the most powerful — and least understood — forces shaping early hydrogen economics is the surge of renewable curtailment and negative electricity prices. Wind farms are shut down because transmission lines are saturated. Solar arrays are disconnected because midday supply overwhelms demand. Generators are sometimes paid to stop producing electricity.

For hydrogen, this is not a problem — it is an opportunity.

Curtailment represents stranded renewable potential. Electrolysers turn that stranded potential into value. Negative prices transform electricity from a cost into a revenue stream. In markets like Spain, Denmark, Germany, California, and Texas, these conditions are no longer rare anomalies; they are structural features of the transitioning grid.

Hydrogen becomes the buyer of last resort for surplus electricity, absorbing volatility that would otherwise be wasted. In doing so, it accelerates the economics of both electrolysers and renewables. Curtailment is what allows early hydrogen plants to experiment with flexible operation, build digital optimisation, and drive down cost long before dedicated renewables arrive.

It is no exaggeration to say that negative-price events are quietly underwriting the first generation of workable hydrogen projects.

Grid Fees and Network Tariffs — The Invisible Cost Barrier

Yet while negative-price events help hydrogen projects, another force — far less visible to the public — works against them: grid fees, levies, and network charges.

In many countries, grid electricity is expensive not because the electrons themselves are expensive, but because the infrastructure and regulatory surcharges layered on top of them are. Network tariffs can double or triple the cost of delivered electricity. Renewable levies, balancing charges, and capacity payments add further strain. Ironically, in some markets, clean electricity delivered through the grid is more expensive to a hydrogen producer than fossil-based alternatives.

This creates a fundamental tension. Electrolysers could stabilise the grid, reduce curtailment, and provide flexibility — but only if the cost of connecting to the grid is rational. In many regions, exemptions or reform are becoming necessary simply to make hydrogen viable at all.

The public debate often focuses on electrolyser cost curves or renewable LCOE. But developers know the truth: grid fees are frequently the difference between a project that moves forward and a project that dies.

Green Electricity and the Grid — The Myth and the Reality

Overlaying all of this is the persistent misconception that electrolyser developers can simply “pull green electrons off the grid.” The grid does not distinguish between sources. Unless directly tied to renewables, electrolysers must rely on

contractual green mechanisms (PPAs, certificates, 24/7 matching) that vary in rigor and cost.

True 24/7 carbon-free operation demands alignment between electrolyser runs and real-time renewable generation. It is a noble aspiration — and ultimately necessary — but during the transition, it will be difficult and expensive. The system must evolve towards it gradually.

This is why the transition period is characterised by compromises that feel imperfect but are economically sensible. Grid-based hydrogen will not be perfectly green at first. Direct connections will not always be available. Yet the trajectory matters: each year, the proportion of clean energy increases, and the system becomes greener in practice as infrastructure catches up.

LCOH — The Messy Middle

The final reality is that green hydrogen remains expensive today, not because the technology is immature, but because the system around it is incomplete. Without dedicated renewables, without optimised operating models, without mature storage and pipelines, and with volatile electricity markets, the cost spread between different projects can be enormous.

Some projects, especially those aligned with high-curtailment regions or cheap solar basins, already approach competitive costs. Others remain heavily dependent on subsidies and policy support. This unevenness is normal for a system in transition.

The decisive point is this: costs cannot be understood without understanding operations.

Electrolyser utilisation, loading strategy, grid fees, PPA structure, curtailment availability, storage integration — these variables matter more in the next five years than electrolyser CAPEX itself.

As the system matures, cost convergence will occur. But the early period will be messy — because all transitions are.

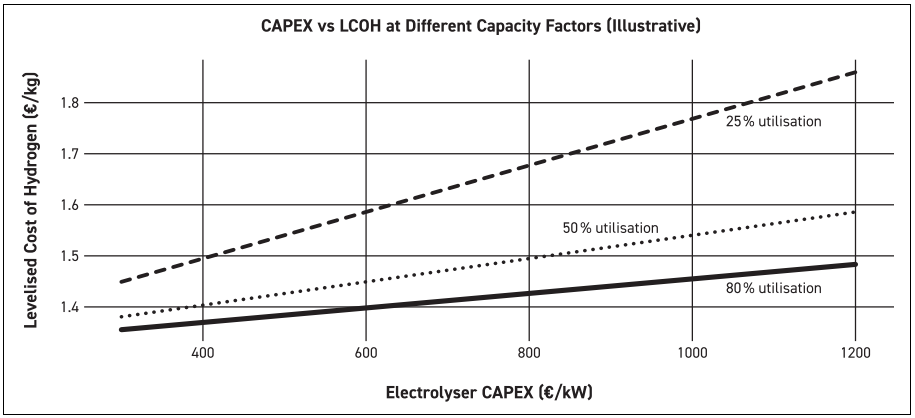
Once electricity procurement and operating strategy are understood, the next cost lever — and the one most often misunderstood — is capital.

CAPEX — The Electrolyser Cost Trap

Electrolyser capital expenditure is often the most visible number in hydrogen debates, and as a result it is frequently over-emphasised. There is good reason for the attention. Over the past decade, electrolyser costs have fallen sharply, from around €1,200–1,500 per kilowatt to €400–700 per kilowatt for large-scale systems today, with credible pathways pointing below €300 per kilowatt as manufacturing scales and learning effects accumulate.

These reductions matter. But they are also easy to misinterpret.

The uncomfortable truth is that lower CAPEX only reduces hydrogen cost if the asset is actually used. Electrolysers are capital-intensive machines. Their cost is largely fixed upfront and must be recovered over every kilogram of hydrogen produced during their lifetime. When an electrolyser operates at a high load factor — for example 80 per cent — that capital is spread across a large volume of output, making each kilogram of hydrogen cheaper. When the same electrolyser operates at 25 per cent load factor, the capital cost is spread across far fewer kilograms, dramatically inflating LCOH.



The figure above illustrates how electrolyser capital cost interacts with utilisation to determine LCOH. At high utilisation, reductions in CAPEX translate into modest but meaningful cost improvements. At low utilisation, however, the same reductions deliver far less benefit, as fixed capital costs are spread over a much smaller volume of hydrogen production. In this regime, even very low electrolyser prices struggle to compensate for underuse. The chart makes clear why CAPEX cannot be considered in isolation: its economic impact is governed by how often the asset is used, not simply by how cheaply it is purchased.

In this situation, even a “cheap” electrolyser becomes expensive in practice. The machine may be technically sound, but economically it behaves like an underused factory: most of the investment sits idle while financing costs continue to accrue. This is why headline CAPEX figures can be misleading when viewed in isolation. A low-cost electrolyser paired with poor utilisation can deliver worse economics than a more expensive unit operating consistently.

The key insight is simple but often overlooked. CAPEX is not a standalone lever. Its impact on LCOH is inseparable from utilisation. Reducing electrolyser costs helps, but only when the surrounding system — electricity supply, demand certainty, grid access, and operational strategy — allows the asset to run for thousands of hours each year. Without that, CAPEX reductions risk becoming a trap: highly visible progress that delivers far less economic benefit than expected.

Utilisation Rate — The Silent Killer

Utilisation, often referred to as load factor, is the most underestimated variable in hydrogen economics. While electricity price sets the floor for LCOH and capital costs determine how far above that floor a project lands, utilisation decides whether either of those advantages is ever realised in practice.

The arithmetic is unforgiving. A 10 MW electrolyser operating for 8,000 hours per year produces on the order of 1.5–1.6 kilotons of hydrogen per year (assuming ~50–55 kWh/kg), or roughly 29–32 kilotons over a 20-year lifetime. The same asset operating for 3,000 hours per year produces only around 0.55–0.60 kilotons per year — roughly 11–12 kilotons over twenty years — even though most of the capital cost is identical. The electrolyser does not become cheaper simply because it is used less. It becomes more expensive per kilogram.

This creates a brutal reality for project economics. Fixed costs — capital, financing, staffing, and maintenance — must be recovered regardless of how often the plant runs. When utilisation falls, those costs are spread across fewer kilograms of hydrogen, driving LCOH sharply upward. From a modelling perspective, this is one of the fastest ways to turn an apparently attractive project into an uneconomic one.

What makes this especially important today is that, in recent months, a growing number of proposed projects have begun to assume extremely low utilisation rates, sometimes well below what would historically be considered reasonable for industrial assets. This has been particularly visible in parts of southern Europe, including Spain, where projects are increasingly designed around curtailed renewable electricity rather than continuous supply. In discussions with

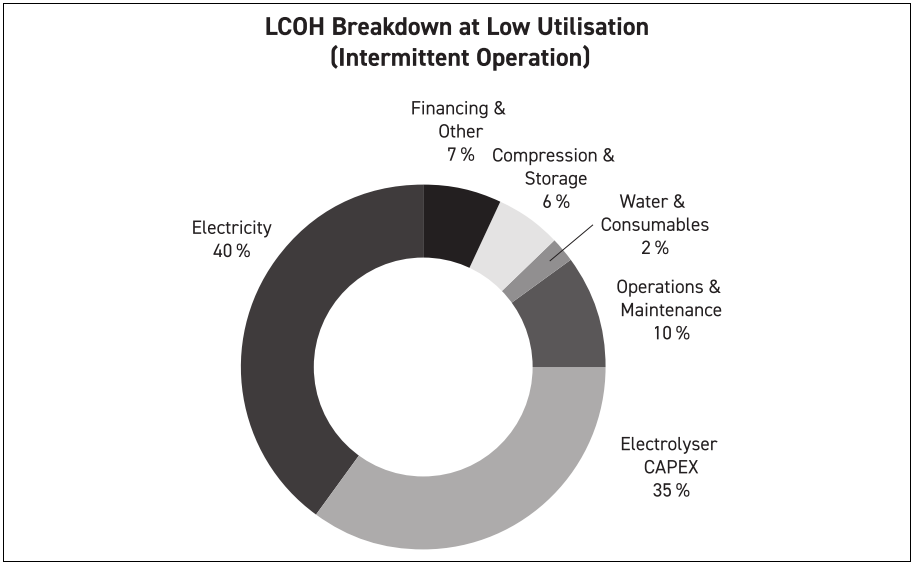
developers and electrolyser manufacturers, it is no longer unusual to hear utilisation figures that would have been unthinkable in other industrial contexts.

At first glance, the logic seems compelling. Excess solar and wind electricity can appear “cheap” or even free. If the power would otherwise be curtailed, why not use it to produce hydrogen? On paper, the electricity input looks attractive, and the headline power price in the LCOH calculation collapses.

The problem is that electrolysers are not marginal assets. They are capital-intensive industrial systems designed to operate continuously. Treating electrolysers simply as a way of absorbing surplus electricity whenever it is available changes the economics entirely. Curtailment-based hydrogen can look inexpensive when viewed through the narrow lens of electricity cost alone, but it often delivers very poor LCOH once capital recovery and financing are accounted for.

This is why hydrogen economics reward boring consistency, not heroic intermittency. A project supplied with moderately priced electricity for most of the year will often outperform one supplied with ultra-cheap electricity for a small fraction of the time. High utilisation does not make headlines, but it is the difference between a system that amortises its capital efficiently and one that quietly accumulates losses.

The contrast becomes clear when comparing a grid-connected electrolyser, able to draw power across thousands of hours, with an electrolyser dedicated exclusively to intermittent offshore wind or curtailed solar. The former may pay more for electricity, but it produces far more hydrogen over its lifetime. The latter may boast impressive power prices, yet struggle to escape structurally high LCOH. In hydrogen, the hours matter as much as the megawatts.



This figure illustrates how the composition of the levelised cost of hydrogen changes as utilisation falls. At low load factors, fixed costs dominate. The electrolyser is paid for whether it runs or not, and each kilogram of hydrogen must carry a larger share of that burden. In such systems, electricity is no longer the main problem — underutilised capital is.

This is precisely why projects like Stegra in northern Sweden are instructive. The hydrogen system at Boden is not designed to chase the cheapest possible electricity for a limited number of hours, but to support continuous, predictable operation aligned with industrial demand. By collocating hydrogen production with steelmaking, integrating with a stable low-carbon power system, and designing for steady offtake rather than opportunistic operation, the project prioritises utilisation over headline power prices. Electricity may not always be the cheapest available at every moment, but the electrolyser is built to run for thousands of hours each year, allowing capital costs to be amortised efficiently. In LCOH terms, this kind of design choice often matters more than securing the lowest possible price for a fraction of the year.

In hydrogen, the most dangerous projects are often not the most expensive ones, but the ones that look cheap for the wrong reasons.

Efficiency — Important, but Not a Silver Bullet

Electrolyser efficiency improvements matter. Every reduction in the electricity required to produce a kilogram of hydrogen translates directly into lower operating costs. These improvements are real, measurable, and worth pursuing. But their economic value cannot be considered independently of the price of the electricity being consumed.

Improving electrolyser efficiency from 55 kWh to 50 kWh per kilogram of hydrogen represents a reduction in electricity consumption of around 10 per cent. At an electricity price of €20 per megawatt-hour, that saves around €0.10 per kilogram of hydrogen. At €50 per megawatt-hour, the saving rises to €0.25; at €80, it becomes €0.40. Efficiency and electricity price are therefore fundamentally interlinked: the higher the cost of power, the greater the economic value of every kilowatt-hour avoided.

Efficiency alone, however, cannot compensate for structural weaknesses elsewhere in the system. An electrolyser with excellent efficiency but poor utilisation may still produce expensive hydrogen, just as favourable efficiency cannot overcome an inherently weak financing structure or poorly matched supply and demand. The economics emerge from the interaction between electricity price, electricity consumption, utilisation, capital cost, financing, and the way the overall system is designed and operated.

The key insight is therefore not that efficiency is secondary, but that it cannot be viewed in isolation. When efficient electrolysers are embedded in well-designed systems — with high utilisation, stable demand, and low-cost electricity — their benefits compound. When they are not, even the best-performing technology can struggle to deliver competitive LCOH.

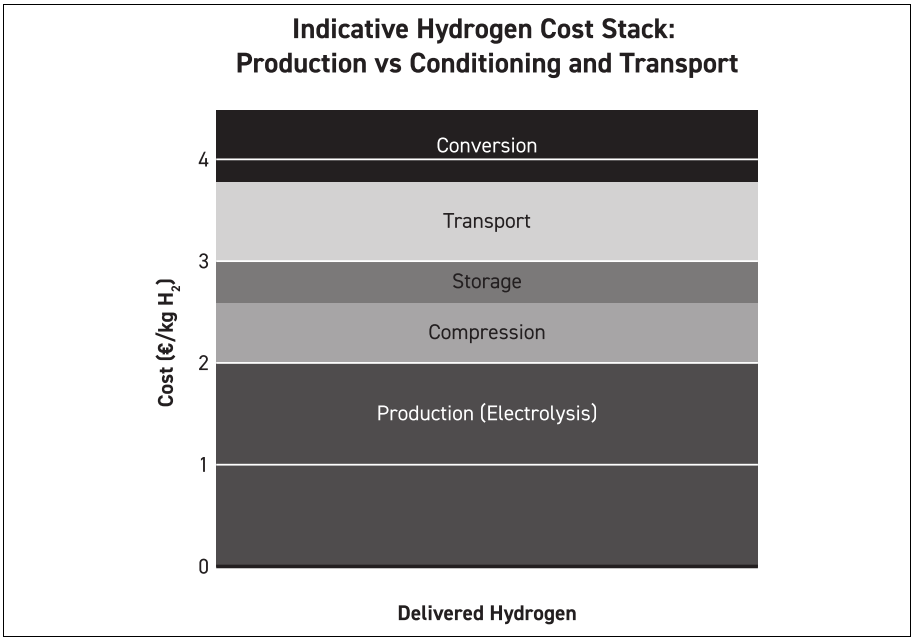
Transport, Storage, and Conditioning — The Forgotten Costs

Discussions of LCOH often stop at the electrolyser fence. Real hydrogen projects do not.

Once hydrogen is produced, it must be compressed, stored, transported, and often converted before it can be used. Each of these steps adds cost, energy losses, and complexity that are frequently underestimated in early-stage modelling. In practice, these downstream elements can reshape project economics as much as electricity price or electrolyser CAPEX.

Compression is the first hurdle. Hydrogen leaving an electrolyser must typically be compressed to 350–700 bar for storage or transport, consuming additional

energy and requiring specialised equipment. Liquefaction goes further still, imposing an energy penalty of 30–40 per cent of the hydrogen's energy content. Storage options — whether in pressurised tanks or underground salt caverns — add both capital cost and operational constraints. Transport introduces another layer entirely, with pipelines, trucking, shipping, or rail each carrying distinct cost structures, scale thresholds, and regulatory burdens. In many cases, hydrogen is also converted into another carrier — such as ammonia, methanol, or liquid organic hydrogen carriers — adding further processing steps and losses.



The cumulative effect of these stages is often surprising. In some projects, the cost of conditioning and delivery exceeds the cost of producing the hydrogen itself. A kilogram of hydrogen that appears competitive at the electrolyser outlet can become uncompetitive by the time it reaches the point of use.

This leads to a simple but uncomfortable conclusion. Hydrogen is often cheap to make where renewable energy is abundant. It is far more expensive to move across long distances. The further hydrogen travels, and the more times it is transformed along the way, the more its economics deteriorate.

This is why proximity matters. Projects that minimise transport and conditioning — by co-locating hydrogen production with demand, or by integrating directly into industrial clusters — consistently outperform those built around long-distance export. LCOH is not just a function of how hydrogen is produced, but of how far it must travel, how often it is handled, and how many times it is transformed before it is finally used.

Perhaps the most important lesson emerging from recent transport studies is that there is unlikely to be a single dominant method for moving hydrogen. Pipelines remain the most economical solution where sufficient volumes justify the investment—particularly where existing natural gas infrastructure can be repurposed. Over longer distances, however, hydrogen may be transported as ammonia, methanol, liquid organic hydrogen carriers (LOHCs), synthetic methane or liquid hydrogen, each offering different trade-offs between conversion losses, infrastructure requirements, transport costs and end-use suitability. Rather than converging on a single global solution, hydrogen logistics are likely to evolve as a portfolio of complementary pathways, each optimised for specific regions, distances and applications.

How LCOH Is Actually Modelled

Most calculations of the levelised cost of hydrogen are built on a discounted cash flow approach. At its core, the concept is straightforward: LCOH represents the total cost of producing hydrogen over the lifetime of an asset, discounted to reflect the time value of money, divided by the total quantity of hydrogen produced over that same period, also discounted. Written formally, it is simply total lifetime cost divided by total lifetime output.

The apparent simplicity of the equation is deceptive. In practice, the result is highly sensitive to a small number of assumptions, each of which can dramatically reshape the outcome. The cost of capital is one of the most powerful levers. A project modelled with a weighted average cost of capital of 6 per cent can appear competitive, while the same project modelled at 10 per cent may see its LCOH rise by more than 30 per cent, even though the physical system has not changed at all. Assumptions about asset lifetime play a similar role. Extending the operating life of an electrolyser from fifteen to twenty-five years spreads capital costs across a much larger volume of hydrogen, materially lowering the calculated cost per kilogram.

Other assumptions are less visible, but no less important. Electrolyser degradation reduces output over time, while stack replacements introduce additional

capital expenditure partway through the project's life. Small differences in how these factors are treated can accumulate into large divergences in reported LCOH figures.

As these parameters vary widely between studies, two projects quoting the same headline LCOH may, in reality, be describing fundamentally different systems under fundamentally different conditions. The number itself is therefore often less informative than the structure of the model that produced it. LCOH works best not as an absolute truth, but as a comparative tool — useful for testing options against one another, provided they are evaluated within the same modelling framework and under consistent assumptions.

For this reason, serious project analysis increasingly relies on integrated system models rather than isolated spreadsheets. Frameworks, such as the ETHOS model suite, place hydrogen production within a wider energy system, allowing interactions between electricity supply, utilisation patterns, storage, infrastructure constraints, and demand profiles to be explored together. Their value does not lie in predicting the future to the last decimal place, but in exposing which assumptions truly drive outcomes. By showing how quickly LCOH shifts as the surrounding system changes, these models reinforce a central lesson of hydrogen economics: cost is not a fixed property of the technology, but an emergent result of the system in which it operates.

How LCOH Has Evolved — and What That Tells Us

A decade ago, green hydrogen was widely regarded as an expensive and largely theoretical option. Cost estimates commonly ranged from five to eight euros per kilogram, placing it well beyond the reach of most industrial and energy applications. Grey hydrogen, produced from natural gas without carbon capture, typically cost between one and one and a half euros per kilogram, benefiting from mature fossil infrastructure, depreciated assets, and the absence of any meaningful price on carbon emissions.

Today, the picture is more nuanced — but also more honest. In most markets, green hydrogen remains expensive. When realistic assumptions are applied around utilisation, grid access, system integration, and financing, current production costs typically fall in the range of six to ten euros per kilogram. In that sense, hydrogen has not yet crossed the threshold into broad economic competitiveness, particularly in cost-sensitive industrial sectors.

Lower figures do exist, and they are often cited in the range of two to three euros per kilogram. These costs, however, are highly context-specific. They are primar-

ily associated with projects in China, where electricity prices, domestic electrolyser manufacturing costs, financing conditions, and regulatory requirements differ fundamentally from those in Europe or North America. Even in these cases, such figures generally apply at the production gate and under favourable assumptions, rather than to fully conditioned, transported, and certified hydrogen delivered to end users. They should be understood as signals of what is technically and systemically possible, not as a global benchmark.

What has changed over the past decade is therefore not a sudden collapse in hydrogen costs, but the emergence of clearer pathways toward lower LCOH under the right conditions. Renewable electricity has become cheaper and more widely available in many regions, reducing the single largest cost input for green hydrogen. Electrolyser systems have grown in scale and standardisation, lowering unit costs and improving reliability. Financing structures have matured as projects move from first-of-a-kind demonstrations toward repeatable designs. Perhaps most importantly, early examples of system integration — particularly in industrial clusters — have demonstrated how hydrogen production, power supply, and demand can be aligned to improve utilisation and reduce waste.

Equally revealing is what has not changed. The fundamental physics of electrolysis remain the same. Electrolysers still require large quantities of electricity, and they still depend on high utilisation to amortise capital efficiently. Infrastructure constraints — from grid connections to storage and transport — continue to shape what is achievable in practice. There has been no technological shortcut that bypasses these realities.

The lesson of LCOH's evolution is therefore not one of technological magic, but of systems learning. Cost reductions have come from designing better-integrated projects, scaling entire value chains, and reducing friction between components, rather than from optimising any single piece of hardware. That distinction matters because it points directly to what must change next — and to the conditions required for hydrogen to move from promising to competitive.

What It Takes to Reach Competitive Hydrogen

Across virtually all credible analyses, one conclusion emerges with remarkable consistency: competitive hydrogen does not result from a single breakthrough, but from the alignment of several conditions at once. No individual lever — cheaper electricity, lower electrolyser costs, or better efficiency — is sufficient on its own. Hydrogen becomes competitive only when multiple parts of the system reinforce one another.

The first and most obvious requirement is access to low-cost electricity. Since power remains the dominant input to green hydrogen, electricity prices must fall well below typical wholesale averages for LCOH to approach levels that industrial markets can tolerate. But price alone is not enough. Cheap electricity that is available only sporadically does little to improve economics if it leaves capital-intensive electrolysers idle for much of the year.

That leads directly to the second requirement: high utilisation. Competitive hydrogen systems must operate for thousands of hours per year, typically above four to five thousand, to amortise capital efficiently. This is why projects anchored to stable industrial demand or grid-connected supply consistently outperform those relying solely on curtailed or highly intermittent renewable power. In hydrogen, the number of operating hours often matters as much as the price paid for each megawatt-hour.

Capital expenditure is the third condition. Electrolyser costs must continue to decline toward the range of roughly five hundred euros per kilowatt or below. Yet, as earlier sections have shown, lower CAPEX delivers real value only when paired with high utilisation and reliable operation. Without those, even inexpensive equipment struggles to produce competitive hydrogen.

The fourth requirement is a low-risk operating environment. Hydrogen projects that face frequent interruptions, uncertain grid access, or unstable regulatory conditions tend to carry higher system costs, regardless of their technical design. Stability, predictability, and operational simplicity all translate into lower effective LCOH over the life of an asset.

The evolution of Levelized Cost of Electricity (LCOE) provides a useful parallel. Early discussions around renewable electricity focused overwhelmingly on the cost of wind turbines and solar panels. As these technologies matured, however, attention shifted towards financing, capacity factor, operational performance, system integration and market design. The hardware became increasingly standardised; competitive advantage came from how intelligently the wider system was designed and operated. Hydrogen appears to be following the same trajectory. While electrolyser costs remain important, long-term competitiveness will increasingly depend on how effectively projects are integrated, financed, utilised and operated rather than on hardware costs alone.

Taken together, these conditions explain a pattern that is now visible across the global hydrogen landscape. Industrial clusters outperform standalone projects because they combine demand, infrastructure, and power access in one place. Grid-connected systems often beat isolated renewable configurations because

they can operate more consistently. Policy frameworks matter not because they change physics, but because they shape whether these conditions can realistically be met at the same time.

Competitive hydrogen, in other words, is not the product of optimisation in isolation. It is the outcome of coordinated system design — and that insight leads directly to the broader lesson of LCOH itself.

The Bigger Lesson: LCOH is a System Metric

The central lesson of this chapter is simple but easily missed. The levelised cost of hydrogen is not an electrolyser metric. It is not a renewable energy metric. It is not even, in the narrow sense, an energy metric at all. LCOH is a system design metric.

Hydrogen costs do not emerge from a single component or decision. They emerge from the interaction between electricity supply, utilisation patterns, capital intensity, infrastructure, and demand. The same electrolyser, supplied by the same manufacturer and operating at the same efficiency, can deliver radically different LCOH outcomes depending on where it is placed, how often it runs, what surrounds it, and what problem it is trying to solve.

This is why so many debates about hydrogen cost become confused. They focus on isolated inputs — the price of electrolysers, the efficiency of stacks, the cost of renewable power — while overlooking the structure of the system those inputs are embedded in. Yet it is the system that determines whether capital is amortised efficiently, whether electricity is used productively, and whether hydrogen flows smoothly from production to use or leaks value at every interface.

The cheapest hydrogen is therefore not produced by the most efficient electrolyser, the lowest-cost solar farm, or the most aggressive cost assumption in a spreadsheet. It is produced by the system that wastes the least capital, the least energy, and the least time. High utilisation, proximity to demand, shared infrastructure, and operational simplicity matter as much as — and often more than — marginal improvements in component performance.

This perspective also explains why hydrogen economics often appear counterintuitive. Projects that look expensive on a narrow, component-level basis can perform well when integrated into a broader system, while projects that look cheap on paper can struggle once real-world constraints are applied. LCOH is not a property of hydrogen itself; it is a reflection of how intelligently the surrounding system has been designed.

Understanding LCOH, then, is not about identifying a single target number or waiting for a technological breakthrough. It is about recognising that hydrogen becomes competitive when systems are configured to support it — when electricity, infrastructure, and demand are aligned rather than treated as separate problems. That insight does not guarantee success, but without it, success is almost impossible.

With that understanding in place, the question shifts. It is no longer simply what hydrogen costs today, but what must be true — structurally and systemically — for those costs to fall in a way markets can accept.

The Push Pull effect

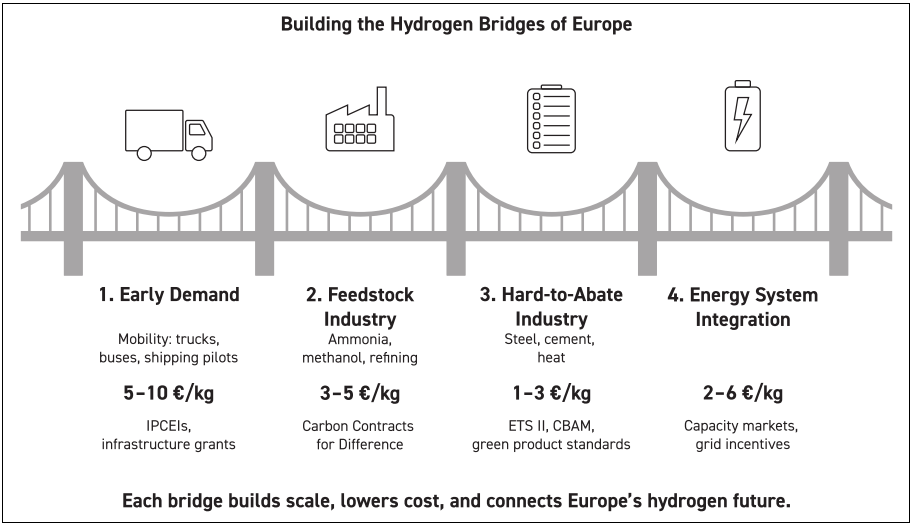
Much of the discussion around hydrogen economics focuses on a single direction of travel: pushing the levelised cost of hydrogen down. Lower electricity prices, cheaper electrolysers, higher utilisation, and better system integration all aim to compress LCOH until it fits the most cost-sensitive sectors of the economy. That pressure is real, and it is necessary. But it is only half of the equation.

The other half is demand.

Hydrogen does not need to be cheap everywhere on day one. It needs to be used where it already works and scaled in a way that allows costs to fall over time. In other words, alongside the push to reduce LCOH, there must be a pull from markets that can tolerate higher prices in the early years and help build volume, infrastructure, and operational experience.

This is a perspective I have argued for repeatedly, but it crystallised for me after being asked a deceptively simple question at a conference: what would you say if you had one minute in an elevator with Ursula von der Leyen? I did not answer it well at the time. But the question lingered because the answer is not about technology or ambition — it is about sequencing.

Europe's hydrogen economy will not be built in a single leap. It will be built bridge by bridge.



The hydrogen "bridge" sequence.
Early hydrogen demand emerges in sectors with higher willingness to pay, such as mobility and specialty industrial applications. These early markets help build scale, infrastructure, and operational experience, enabling hydrogen to move progressively into more cost-sensitive sectors such as steel, cement, and energy system integration. Each bridge lowers system cost and risk, making the next economically viable.

Some sectors can already absorb hydrogen at relatively high prices. High-value mobility applications — such as trucks, buses, material handling, and early shipping pilots — can often justify hydrogen in the range of five to ten euros per kilogram. These markets are small in volume, but they are powerful in what they enable: early infrastructure, operational learning, and committed offtake.

Beyond that sit industrial feedstocks. Ammonia, methanol, and refining already consume hydrogen at scale, and in many cases can begin transitioning in the three to five euro per kilogram range. These applications bring volume and continuity, anchoring utilisation and accelerating cost reductions across the system.

Only after those bridges are in place does hydrogen begin to make sense for the most demanding sectors: steel, cement, high-temperature heat, and other hard-to-abate industries. These sectors require hydrogen closer to one to three euros per kilogram, and they cannot be the first movers without imposing unsustainable costs. Asking them to lead is equivalent to asking the most price-sensitive parts of the economy to subsidise system learning for everyone else.

Each bridge builds the foundations for the next. Infrastructure becomes shared rather than bespoke. Skills move from novelty to routine. Supply chains stabilise. Most importantly, demand becomes real rather than hypothetical. Electrolysers do not scale because spreadsheets say they should; they scale because customers sign contracts and keep buying hydrogen through good years and bad.

This is where policy plays its most constructive role. The task is not simply to fund electrolysers or pipelines, but to provide certainty of demand through instruments that reduce risk and reward early commitment. Carbon contracts, offtake guarantees, and cross-border alignment are not distortions of the market; they are the scaffolding that allows it to be built in the right order.

Seen through this lens, LCOH is not just a number to be minimised. It is a signal that guides sequencing. Push costs down, yes — but pull demand forward where it already works. If those two forces are aligned, hydrogen does not need to wait for perfection. It can scale through use, learn through deployment, and arrive at competitiveness as a consequence, rather than a prerequisite.

9. Financing the Hydrogen Economy

'The problem at the moment is that as a society we see ecological need for Hydrogen, but we do not accept the economic reality. We need to make Hydrogen "Eco2". Both Ecological and Economical.' Anon at a H2 conference

If technology defines what is possible in the hydrogen economy, finance defines what is real. Hydrogen does not scale because it is promising; it scales because projects attract capital. And capital flows only when risk is understood, revenue is predictable, and policy frameworks give investors confidence that the economics will hold over decades. This chapter unpacks how that system works in practice, and what it takes to turn a hydrogen project from a concept note into a financial close.

One thing about marrying into a German family is that Christmas is very different, no trip to the pub on Christmas Eve, no sport events on Boxing Day, instead what you get is 3 days sitting around a table with family! If you tried doing that in the UK, I think most people who know would agree that the event would be catastrophic, and divorce rates would rise exponentially. In Germany, it is a little more civilised, but it is still a long time to sit around with people with very little left to do. In my wife's family the solution is 'Mau Mau'. It is basically Uno with a few extra rules. Of course the idea is to get out first 'Mau', when you do get out all the other players in the game get penalised by the value of the cards in their hand, the penalty is double if the person who gets out finishes on a Jack 'Mau Mau', It is also important to know that the Jack carries the highest number of points in the deck, so you don't really want to be caught holding one at the end of the round, especially not if your opponent does finish on the Jack, the penalty is high. There is an upper limit defined at the beginning of the game (say 300 or 500 points) and the game plays out till the first player crosses that barrier and the person with the least number of points win. Simple.

Now when I was first introduced to this game the first year, I was consistently getting annihilated. Losing high every round and the games did not last long, so I worked on a strategy of trying to 'lose small', I was basically using 'defensive decision making' to ensure that I didn't lose. However, by doing so, I also didn't win as I was initially just looking to ensure I wasn't last. As I got to understand the game better, I then developed tactics to know when to take on the risk of looking to win big, but also on when to dump cards.

Hydrogen financing works in much the same way: players—banks, developers, and investors—must learn when to play safe and when to take calculated risks.

Where we stand in hydrogen is a bit like where I was when I first learned Mau Mau, it was new to me the same way that hydrogen is nascent to a lot of banks and project developers, and I see that there is often the same defensive decision making in play when it comes to selecting projects. Decision makers are still attempting to work out their hand and where to 'bet big' and where to 'dump'. Nobody loses their job by being sensible, they do for taking risks. This is especially true for green hydrogen projects, where there is less knowledge on how to proceed, blue is better understood from the knowledge of the O&G and chemical industries so that helps them move forward smoother than green. From data from IHS from all of the green hydrogen projects out there, in the last years less than 10% of them were coming to Final Investment decision (FID). That number has slightly improved in 2025, and I believe it to now be slightly over 10%.

Hydrogen's Financial Reality: High Risk, High Capex

Hydrogen's financing challenge rests on three structural issues: high upfront capital requirements, uncertain offtake at competitive prices, and policy frameworks that are still evolving. Unlike wind, solar, or LNG, the commercial architecture around hydrogen is still immature.

If the introduction to this chapter illustrated anything, it is that hydrogen finance resembles a game where the players know only part of the rules, the cards are unfamiliar, and no one wants to be holding the Jack at the end. The German FFE's recent analysis for the Federal Ministry for Economic Affairs and Climate Action gives numerical weight to this intuition: investors currently view hydrogen—especially green hydrogen—as a high-risk asset class. In finance, high risk does not mean nervousness; it means higher capital costs, higher equity shares, and fewer projects reaching Final Investment Decision.

At the core lies a simple mismatch: in many applications, the willingness to pay for green hydrogen remains below its cost of production. Whether in mobility,

industry, heat, or power generation, most customers cannot absorb the premium. This is not a matter of ideology; it is arithmetic. If hydrogen costs €5–7/kg to produce but the off taker can only justify €2–3, a revenue gap emerges that banks will not underwrite.

Since cashflows depend on offtake, and offtake depends on economics, developers find themselves in a loop: everyone accepts the ecological case for hydrogen, but the economic case is still forming. FFE puts it bluntly: expected cashflows from electrolysers are highly uncertain and carry a high degree of risk. As a result, debt financing works only under very specific conditions—strong government support, an anchor industrial off taker, extremely cheap renewable power, or all three.

This perception of risk has measurable financial consequences.

You can see this dynamic in the many hydrogen projects that never make it past the glossy rendering stage. In more than one European case, developers assembled impressive consortia, signed MOUs with potential off takers, and secured land and grid connection—yet still failed to reach FID. The stumbling block was always the same: nobody was prepared to lock in long-term volumes at a price that covered the true cost of production. Lenders saw PowerPoint demand rather than contractual demand, and without that anchor the risk premium remained painfully high. The technology was not the problem. The economics were.

Final Investment Decision (FID)

The Final Investment Decision (FID) is a pivotal milestone in any major project's lifecycle. It marks the moment when project sponsors formally commit to full-scale investment and execution. In simple terms, FID is the point where planning ends and building begins.

Before FID, the project team completes extensive preparatory work—feasibility studies, front-end engineering design (FEED), financial modelling, and risk assessments—to ensure that technical, economic, and regulatory uncertainties are minimised. Only once this groundwork is complete can the project be considered ready for funding and implementation.

Approval of FID authorises the release of capital, the signing of major contracts, and the start of construction. It also signals to investors, lenders, and stakeholders that the project has met all critical conditions and is considered viable and ready to proceed.

The importance of FID lies in its finality and commitment. A well-supported FID reflects confidence in the project's design, cost estimates, and risk management. Conversely, a premature FID can lead to cost overruns, delays, and a loss of stakeholder trust. In essence, the FID transforms a project from concept to reality—turning detailed plans and forecasts into tangible results.

In hydrogen, FID is not simply a procedural milestone. It is the point at which policy promises, cost curves, offtake discussions, and technical assumptions must all survive the scrutiny of lenders simultaneously. This is why so few green hydrogen projects reach it—and why everything that follows in this chapter ultimately flows from whether FID is achieved or deferred.

FID is the Mau Mau moment: the point at which you either put the final card down and commit — or keep holding risk in your hand while the penalties quietly compound.

A Risk Premium Two to Three Times Higher Than Wind or Solar

Banks do not price risk emotionally; they price it mathematically. And today, the mathematics for hydrogen is punishing.

According to FFE's survey of 17 German banks and investors, hydrogen projects face a risk premium of 3.3% over EURIBOR:

- Onshore wind: 0.8–1.5%
- Hydrogen: ~3.3%

Hydrogen therefore carries a 2–3× higher risk premium than mature renewables. The resulting debt interest rate—around 5.7% as of December 2024—raises the cost of capital and helps explain why many hydrogen projects stall even when the engineering is sound and the decarbonisation case is compelling.

Equity Investors Want More Skin in the Game — And a Higher Return

Higher debt costs are only part of the picture. As financiers perceive hydrogen as risky, they require more equity to buffer potential volatility.

FFE reports that banks now expect an equity ratio of 35% in hydrogen projects, compared with roughly 20% in traditional renewable financing. This is a substan-

tial burden: higher equity means developers must raise—and lock in—much more of their own capital before lenders release a single euro.

Equity is always expensive. Unsurprisingly, equity investors typically demand a 12% return to compensate for the risk.

To put it plainly: hydrogen developers need more equity, and that equity must be compensated at a higher rate. The capital burden rises sharply in the early years, pushing up the threshold for project viability.

One phrase I heard frequently when I first entered the hydrogen sector—but thankfully far less today as discussions mature—is the mantra of "1 kg / 1 $ / 1 decade." It is catchy, but misleading. It treats hydrogen as a single global commodity with a single cost target, when in reality each sector has its own economics and willingness-to-pay thresholds. Steel has one price point, fertilisers another, mobility yet another.

Nevertheless, the mantra still casts a shadow over many FIDs. Investors may not expect $1/kg exactly, but they do expect a credible cost-reduction trajectory. If a project cannot demonstrate plausible movement toward the $1–2/kg zone within the next decade, institutional capital often hesitates. Investors follow cost curves, not slogans.

The Weighted Average Cost of Capital (WACC): The Number That Shapes Everything

Combine expensive debt with expensive equity and you reach the single number that governs investment decisions in infrastructure: the Weighted Average Cost of Capital (WACC).

FFE calculates the current WACC for hydrogen at 6.79%. At first glance, this may not look dramatic—until compared with wind and solar, which often achieve 3–5% in Europe.

Every percentage point matters. Higher WACC → higher levelised cost of hydrogen → lower willingness to pay → greater revenue uncertainty → higher perceived risk → higher WACC.

A circular trap that cannot be broken by market forces alone.

This dynamic mirrors the Mau Mau example earlier in the chapter; when the rules are unfamiliar and the stakes feel high, players default to defensive strategies—small, careful moves—until the incentives shift.

Why This Matters: Hydrogen Cannot Defy Economics

The FFE analysis makes one point unmistakably clear: the barrier to hydrogen's scale is not a lack of enthusiasm, nor an absence of government ambition. The real constraint is that capital flows toward certainty, not aspiration. Investors do not finance hope; they finance predictability. And in a sector where costs, regulations, and demand patterns are still evolving, that predictability must be engineered rather than assumed.

This is where policy mechanisms step in. Subsidies, Contracts for Difference (CfD's), tax credits, and long-term power purchase agreements are often described as incentives, but in the language of finance they are something far more fundamental: they are the mechanisms that turn volatile, ambiguous future revenues into something steady enough for a bank to lend against. They close the distance between what hydrogen currently costs to produce and what buyers are currently prepared to pay. They offer stability in place of uncertainty, replacing jagged price curves with something smooth enough to model over twenty years. They reduce the anxiety of lenders, shrink the amount of equity developers must commit, lower the overall cost of capital, and—crucially—allow projects to reach Final Investment Decision. Without these instruments, most business cases collapse under the weight of their own unpredictability.

Finance, at its core, does not respond to potential; it responds to numbers that stay still long enough to be trusted.

It is tempting to tell a triumphant story here, as if enthusiasm alone will pull hydrogen into maturity. Yet hydrogen finance today is still a story of caution—a careful, incremental learning curve in which every participant is trying to understand a market that does not yet behave like any of the familiar energy commodities. The Mau Mau analogy from earlier in the chapter fits uncomfortably well. Playing defensively keeps you from losing but rarely puts you in a position to win. Only when you start to understand the rhythm of the game—when to push, when to hold back, when to shed risk—do you influence the outcome rather than merely endure it.

Hydrogen finance sits precisely at that transitional moment. Banks, developers, policymakers, and investors are no longer total newcomers, but they are still reading the cards with caution. The FFE findings explain why: risk premiums remain elevated, offtake agreements remain fragile, and regulatory frameworks—though improving—still feel young. But the same analysis also hints at how quickly this could change. As markets mature, as PPAs and CfDs become standardised rather than bespoke, as electrolyser costs fall and industrial demand becomes

more predictable, the unfamiliar begins to feel routine. The risk premium starts to shrink. Equity requirements ease. The Weighted Average Cost of Capital edges downward. Projects that once felt experimental start to resemble the type of long-lived infrastructure investors understand instinctively.

Hydrogen, in other words, is still in its early rounds. The players are cautious—understandably so—but the shape of the deck is becoming clearer. Once the rules stabilise and the rhythm of the game becomes familiar, decision-makers will no longer play simply to avoid losing.

They will begin, finally, to play to win.

Policy Creates the Market: Pricing, Support, and Market Formation

Hydrogen does not scale on enthusiasm alone. It scales when the economics work, and in the early years of any new energy system that almost always requires policy support. Subsidies are not a reward for producers; they are a correction for markets that have not yet matured. Their purpose is simple: to close the gap between the cost of producing hydrogen and the price buyers can realistically pay.

That gap is structural. Green hydrogen competes against depreciated fossil assets, volatile electricity prices, and industrial customers operating on thin margins. Without policy intervention, many projects remain trapped between high production costs and limited willingness to pay.

Yet hydrogen's challenge goes deeper than a simple cost gap. Unlike electricity, oil, or gas, hydrogen is not entering a mature commodity market. Those systems evolved over decades—sometimes centuries—through massive public investment, geopolitical competition, and regulatory frameworks that gradually built supply chains, trading hubs, price benchmarks, and infrastructure networks. Hydrogen has none of that history. In most regions there is no liquid market, no universal reference price, and often not even the infrastructure needed to move molecules from where they can be produced to where they are required.

In such an environment, leaving development entirely to private actors leads to paralysis. Producers hesitate to invest without firm buyers. Buyers hesitate to commit without confidence in stable supply and predictable prices. Policymakers are cautious about imposing detailed rules before commercial activity emerges. Each group waits for the others to move first.

A hydrogen project is therefore a bit like trying to build a railway before the train schedules exist. The tracks only make sense if the trains will run, and the trains only make sense if the tracks are already there. No rational actor wants to move first. Policy is the mechanism that breaks this deadlock, creating the conditions under which supply, demand, and infrastructure can grow together.

Governments typically intervene through three complementary mechanisms: pricing, financial support, and market rules.

The first is carbon pricing, which corrects the long-standing distortion that has kept fossil hydrogen artificially cheap. Grey hydrogen has benefited for decades from mature infrastructure, depreciated assets, and the simple fact that its CO_2 emissions were not reflected in its price. Carbon pricing changes that equation. Whether delivered through an emissions trading system or a straightforward carbon tax, the result is the same: emitting CO_2 begins to carry a financial consequence.

As soon as that happens, the economics of hydrogen begin to shift. Grey hydrogen becomes more expensive because its emissions finally appear on the balance sheet. Blue hydrogen may gain competitiveness depending on capture rates and methane management. Green hydrogen benefits indirectly—not because it suddenly becomes cheap, but because the gap between clean and fossil-derived hydrogen begins to narrow.

The European Union provides a clear example. Under the EU Emissions Trading System (ETS), the cost of producing grey hydrogen has risen significantly, as producers have been required to pay for emissions that were previously externalised. As carbon prices rise and free allocations decline, the economics of unabated hydrogen production continue to deteriorate.

Where carbon pricing is absent, distortions remain. In those markets, low-carbon hydrogen must compete against fossil hydrogen that effectively enjoys a hidden subsidy—the ability to emit freely. This is one reason hydrogen economics vary so widely across regions: the difference is often not the cost of electrolysers or renewable power, but whether emissions carry a price.

Even so, carbon pricing alone rarely makes green hydrogen bankable. The cost gap between green and grey hydrogen remains too large in most markets, and carbon prices would need to be far higher—and more predictable—than they are today to close it completely.

This is where the second policy instrument enters: financial support.

Subsidies, tax credits, capital grants, and Contracts for Difference are designed to bridge the remaining gap between production costs and market prices. Their purpose is not generosity but risk reduction. By narrowing the difference between what hydrogen costs to produce and what customers can afford to pay, these mechanisms allow early projects to move forward while technology improves, and economies of scale begin to emerge.

In practice these instruments work in different ways. Operational subsidies reduce the effective cost of hydrogen production, capital grants lower the upfront investment required to build facilities, and Contracts for Difference stabilise revenues by guaranteeing a reference price over long periods. Together they provide the financial predictability investors need during the early stages of market development.

Despite their different designs, these instruments share a common objective: predictability. Investors do not need occasional injections of cash; they need confidence that revenue streams will remain stable across the twenty-year life of an asset. In a sector where the cost of capital is one of the dominant drivers of hydrogen price, stability is not a luxury—it is the foundation of bankability.

The third pillar of market creation lies in standards and definitions, which ensure that hydrogen can be traded, certified, and trusted across borders. Clear rules defining what qualifies as renewable or low-carbon hydrogen give investors confidence that projects built today will remain compliant tomorrow.

Together these three policy tools—pricing, financial support, and market rules—form the architecture of an emerging hydrogen economy. Carbon pricing corrects the distortion that made fossil hydrogen artificially cheap. Subsidies bridge the economic gap during the early years of market development. Standards create the shared framework that allows producers, buyers, regulators, and financiers to operate with confidence.

Only when these elements work together can hydrogen move from isolated pilot projects to a functioning marketplace.

Subsidies and Tax Credits: Bringing Costs Down

Subsidies work in the opposite direction to carbon pricing: they lower the cost of clean hydrogen during the early years of market formation. They are not permanent features of the market, but transitional tools designed to bridge the gap until technology improves, scale builds, and costs fall. Different governments have

approached this in different ways, but all are responding to the same basic problem: clean hydrogen still costs more than most customers are willing to pay.

In the United States, the original version of Inflation Reduction Act once stood as the most influential policy instrument shaping global hydrogen economics. In its initial form, the 45V production tax credit offered up to $3 per kilogram for the cleanest hydrogen, dramatically narrowing the gap between green and fossil-based alternatives. For a time, it effectively reset global expectations for hydrogen cost curves, and even today the industry still speaks about "pre-adjustment 45V economics" as a benchmark for what strong operational support can achieve. Later revisions reduced its impact, and the geopolitical consequences of those changes are explored in a later chapter, but historically the original 45V remains the clearest example of how policy can transform hydrogen from an expensive frontier technology into something immediately competitive.

Europe chose a different structure. The EU Hydrogen Bank offers fixed premiums designed to smooth revenue volatility and provide developers with predictable income streams in a market otherwise exposed to electricity price swings and uncertain demand. The United Kingdom adopted a similar approach through its Hydrogen Production Business Model, borrowing from the Contract-for-Difference framework that helped launch offshore wind.

Japan and South Korea, both heavily dependent on energy imports, have focused on subsidy systems designed to secure stable supplies of low-carbon hydrogen for domestic industry as they decarbonise.

Although these mechanisms differ in design, their purpose is similar: to provide the revenue certainty required for a Final Investment Decision. Subsidies cannot compensate for expensive renewable electricity or weak demand, but they do address the central obstacle in early hydrogen development—the fact that clean hydrogen still costs more than most customers are prepared to pay.

There is also a geopolitical dimension to these policies, which is explored later in the book. From a financing perspective, however, the principle is straightforward: structured support reduces uncertainty. When revenue expectations become more predictable, capital becomes less risk-averse—and projects begin to move from concept slides to construction.

Subsidies, in essence, are not about generosity. They are about reducing risk to a level the financial system can tolerate.

Standards, Certification & Guarantees of Origin: Making Hydrogen Real

If carbon pricing and subsidies shape the economic landscape of the hydrogen sector, certification shapes its credibility. It determines whether a kilogram of hydrogen can legitimately be called "green," whether it qualifies for public support, and whether it can be traded across borders without dispute. For investors, policymakers, and industrial buyers, certification turns hydrogen from a climate aspiration into a verifiable commodity with a measurable carbon footprint.

Achieving that credibility requires more than a simple label. A robust certification system must trace the emissions associated with hydrogen production—from the source of electricity used in electrolysis to the infrastructure and supply chains involved in producing and delivering the molecule. It must account for embedded emissions in equipment, upstream manufacturing footprints, and the interaction between electrolysers and the electricity grid. This level of transparency ensures that the declared carbon intensity of hydrogen can be independently verified rather than simply asserted.

This is why Guarantees of Origin and similar certification schemes matter. They create a recognised identity for each unit of hydrogen and allow buyers, regulators, and financiers to verify its environmental credentials. Increasingly, these systems go beyond simple claims about renewable electricity and require a full product carbon footprint—a cradle-to-gate accounting of emissions associated with hydrogen production.

The need for this level of rigour is hardly surprising. When governments allocate billions in public funding to support hydrogen projects, it is entirely reasonable that they demand evidence that the projects genuinely deliver the environmental outcomes they promise. Subsidies and grants are not simply financial support—they are public investments in decarbonisation. Certification provides the mechanism that allows those investments to be verified.

Meeting these requirements has created a new ecosystem of digital tools designed to track and verify carbon data. Siemens' SiGreen platform, for example, allows companies to calculate and share product carbon footprints using verified supplier information rather than broad assumptions. In hydrogen projects, such tools can account for emissions embedded in electrolysers, compressors, and other components while assembling a complete carbon profile for the facility. As hydrogen trade grows internationally—from Spain to Germany or from Australia to Japan—systems like these will become essential infrastructure for verifying the carbon credentials of traded molecules.

Data alone, however, is not enough. Producers also need systems that translate carbon accounting into recognised certificates and compliance documentation. This has led to the emergence of "compliance tech" platforms, such as Munich-based Atmen. As an accredited Compliance Technology Provider under the CertifHy RFNBO scheme, Atmen enables producers to collect operational data, calculate regulatory compliance, and generate the digital certificates required for Guarantees of Origin.

Together, carbon-footprint platforms, such as SiGreen and certification systems like Atmen, form the digital backbone of the emerging hydrogen market. They ensure that claims of "green hydrogen" are supported by traceable evidence rather than marketing language.

Without certification, hydrogen remains a promise. Buyers cannot prove decarbonisation, regulators cannot confidently allocate subsidies, and investors cannot underwrite long-term contracts. With credible certification, the opposite becomes possible. Hydrogen becomes a dependable, tradable commodity whose value rests on transparent and defensible carbon data.

In a market where carbon intensity increasingly determines value, certification becomes the currency of trust. It is the mechanism that allows governments to justify public support, companies to prove decarbonisation, and investors to believe the numbers in a project's financial model. And it is that trust—quietly embedded in certificates, data platforms, and verification systems—that ultimately turns prototypes into pipelines and pilot plants into full-scale industry.

How Hydrogen Projects Become Bankable

Once you strip away the jargon, a hydrogen project becomes bankable when it can tell a convincing story about its future cashflows. In hydrogen, investors fear three things above all: volatile electricity costs, uncertain offtake, and shifting regulations. Every tool discussed in this chapter—from carbon pricing and subsidies to certification, hubs, and blended finance—exists to neutralise some part of that fear. Not a story in the marketing sense, but in the sense of a spreadsheet that holds up when a lender's risk team tries to tear it apart. At that point the technology almost becomes secondary. Electrolysers, pipelines, storage and renewables all matter, but what ultimately matters is whether the project can reliably earn enough money, for long enough, to repay those who put capital at risk.

As hydrogen markets are still forming, that financial story rarely emerges on its own. It must be built. The policy tools and financial structures surrounding

hydrogen exist to bridge the gap between a risky, uncertain reality and the level of predictability that banks and investors require.

The first part of that bridge is the price. Left to itself, the price of hydrogen in early projects would fluctuate wildly with electricity costs, carbon prices, competing fuels, and uncertain demand. That kind of volatility is poison to project finance. The sector therefore borrows a concept from the power industry: long-term price stabilisation. Contracts for Difference (CfDs), introduced earlier in the policy discussion, are the clearest example. On paper they are simple mechanisms that top up revenue when market prices fall below a reference level and claw back excess when prices rise above it. For financiers, however, that simplicity is precisely the point. A CfD turns a volatile market revenue stream into something that behaves more like a regulated utility payment.

The difference can be dramatic. Consider a hypothetical 100 MW electrolyser connected to variable renewables in an industrial zone. Without support, its revenue depends on hydrogen prices linked to fluctuating electricity costs and competing fuels. Banks looking at that profile see a rollercoaster: profitable years, marginal years, and potentially years deep in the red. Add a 10–15-year CfD at a fixed reference price and the revenue profile stabilises. The underlying technology has not changed, but the cashflow has—and that is what allows lenders to consider financing the project.

Price stability alone, however, is not enough. Someone still must buy the hydrogen. This is where long-term offtake agreements enter the picture. These are contracts in which industrial customers—steel mills, refineries, chemical plants or mobility operators—commit to purchasing defined volumes of hydrogen over many years. From the outside, they are framed as decarbonisation commitments. From the financier's perspective, they serve a simpler purpose: they ensure the plant will not sit idle without revenue. A project with a credible off-taker, particularly one with an investment-grade balance sheet, immediately appears more bankable than one relying on speculative future demand.

Negotiating such agreements is rarely straightforward. Off-takers seek flexibility because their own markets remain uncertain, while producers require firm commitments to satisfy lenders. In practice, these negotiations often converge on long-term contracts with defined price formulas, volume commitments and some form of take-or-pay structure. Once signed, those agreements become the backbone of the financing model.

Even with price support and offtake secured, many hydrogen projects still carry risks that commercial lenders hesitate to shoulder alone. Technologies must be

integrated at scale, construction risks are substantial, and regulatory frameworks continue to evolve. In these early stages, public lenders and development banks frequently play a critical role. Institutions, such as the European Investment Bank, KfW, and multilateral development banks, can step in by accepting lower returns, offering longer loan tenors, or taking on portions of the risk that private banks find uncomfortable—for example through guarantees or first-loss positions. This form of blended finance does not transform a weak project into a viable one, but it can shift the risk–return balance enough to unlock private capital.

As more projects are built and operate successfully, something else begins to happen, specifically familiarity. Technologies move from "innovative" to "proven," contracts become standardised rather than bespoke, and regulatory frameworks stabilise. At that stage, a different class of investor begins to participate— infrastructure funds, pension funds and industrial asset managers who specialise in long-lived assets with predictable returns.

Offshore wind followed a similar trajectory. Early projects relied heavily on government support and corporate balance sheets. Over time, as revenues stabilised and costs fell, lenders grew comfortable and long-term infrastructure investors entered the sector. Hydrogen's financial evolution is likely to follow the same pattern.

From an investor's perspective, the evaluation ultimately comes down to a few simple questions. Is the revenue stable and sufficient? Does the technology perform reliably at scale? Are the regulatory rules predictable over the life of the asset? And does the project's cost trajectory suggest it will become competitive over time?

When those conditions are met—structured revenues, committed buyers, manageable technical risks, stable regulatory definitions, and a credible cost pathway—hydrogen projects begin to resemble the kinds of infrastructure investments financiers already understand.

Bankability, in other words, is not an abstract label. It is the outcome of deliberate choices about how risk is allocated and how future cashflows are structured. For hydrogen, getting that right will be just as important as getting the chemistry and engineering right.

Form Follows Function — and Financing Follows Use Case

Hydrogen finance is usually discussed from the production side — gigawatt-scale electrolysers, renewable PPAs, storage caverns, and industrial off-takers. But fin-

ancing challenges do not stop at the plant gate. Downstream use cases matter just as much because the economic logic of hydrogen is determined by how and where the molecule is used. In architecture, "form follows function" describes the way design emerges from purpose. In hydrogen, something similar happens; financing follows the use case. Capital flows most easily when the application creates genuine value, not simply when hydrogen is inserted into a process for the sake of substitution.

A striking example of this principle in action is the ATCO Hydrogen House in Canada — a project far removed from the usual industrial offtake models. On the surface, powering a single home entirely with hydrogen appears to violate the first rule of hydrogen economics: do not use hydrogen where electrons can do the job more efficiently. And yet the project makes financial sense for ATCO precisely because of its specific use case. The hydrogen produced on-site is not competing with electricity or heat pumps in a conventional residential market. Instead, the economics hinge on something far more interesting: the opportunity cost of natural gas.

ATCO realised that the natural gas that would ordinarily heat this house was more valuable if exported into the wider gas system at market prices than if consumed on-site. By producing hydrogen locally for the residence — even at a higher per-unit cost — ATCO can sell the freed-up natural gas into a market where its value is greater. The house becomes, in effect, an arbitrage machine: hydrogen is used not because it is the cheapest energy carrier, but because it enables a more profitable use of another one. The result flips the traditional cost hierarchy. Hydrogen is not the economical option in isolation; it becomes economical in context.

That is the core lesson for financiers. The bankability of downstream hydrogen applications depends not only on efficiency metrics or techno-economic comparisons, but on the strategic role hydrogen plays in a broader system. A use case that looks irrational on a levelised cost chart can become financially sound when placed in its real-world setting — when it unlocks value elsewhere, mitigates regulatory exposure, or enables revenue streams that the hydrogen itself does not carry. In other words, lenders don't finance hydrogen molecules; they finance the business logic around them.

Projects like the ATCO Hydrogen House remind us that downstream hydrogen economics are not monolithic. Industrial heat, mobility, grid balancing, synthetic fuels, and blending into networks all carry different risk profiles, different regulatory dependencies, and different revenue models. The right financing structure for each of them begins with understanding the underlying use case, not the hydrogen technology itself.

Hydrogen is only expensive when it is used in the wrong place.

In the right place, it can unlock value that no other energy carrier can.

The Evolving Investment Landscape

In the earliest stages of the hydrogen sector, the investors most comfortable with uncertainty are those backing technology rather than infrastructure. Venture and growth-equity capital appears here, but its role is limited: funding new electrolyser designs, materials innovations, digital platforms, and other components that might eventually shift the industry's cost curve.

Even at this stage, however, traditional venture capital is not the dominant force. Much of the early investment has come from strategic industrial players acting more like venture investors than venture funds. Companies, such as MAN Energy Solutions, through its investment in H-Tec Systems (now Quest One), have backed hydrogen technologies not primarily for short-term financial returns, but to secure technological footholds and influence the value chains they expect to depend on in the future. These investments matter, but they remain peripheral to the larger question of project bankability: they finance capabilities rather than finished assets.

As the sector matures, the centre of gravity shifts toward strategic corporates—energy majors, steelmakers, chemical producers, and heavy industries—whose participation signals that hydrogen has moved from speculative to strategic. These companies bring balance-sheet strength, engineering capability, and long-term commercial logic.

Only once revenues stabilise and risks become familiar do infrastructure funds and long-term asset owners begin to appear. These investors specialise in assets that resemble utilities: high upfront cost, modest but predictable returns, and limited surprises. Their arrival usually marks the point at which a sector moves from demonstration to scale, much as offshore wind did a decade ago.

Hydrogen Hubs: Concentrating Supply, Demand, and Confidence

One of the most important developments in the maturing hydrogen economy is the rise of the hydrogen "hub." The term is sometimes used loosely, but in its strictest sense a hub is far more than a cluster of projects on the same map. It is a deliberately designed ecosystem in which production, demand, storage, and infrastructure are developed together rather than in isolation. By aligning indus-

trial off takers, renewable power suppliers, infrastructure operators, and regulators in a single location, hubs create a level of certainty that individual projects struggle to achieve on their own.

A hydrogen hub typically begins with co-location. When production sits next to major industrial users—steel plants, refineries, chemical complexes, or ports—the cost and complexity of transporting hydrogen fall dramatically. More importantly, co-location resolves one of the central coordination problems in early hydrogen markets. Producers hesitate to invest without firm buyers, while potential customers are reluctant to commit until supply is guaranteed. Hubs break this stalemate by developing supply, demand, and infrastructure in parallel, giving both sides the confidence to move forward.

Shared infrastructure is another defining feature. Storage caverns, compressors, pipelines, high-voltage grid connections, and import or export terminals can serve multiple users rather than being built for a single project. By spreading these costs across many participants, hubs reduce the capital burden placed on individual projects and unlock economies of scale that isolated developments cannot achieve. For financiers, this shared infrastructure significantly improves the investment profile of each participating asset.

Hubs also diversify demand. A standalone hydrogen project may depend on a single industrial customer, leaving it vulnerable if that buyer delays or scales back its plans. Within a hub, hydrogen can be supplied to multiple users—industrial processes, mobility applications, power generation, or export routes. This broader demand base lowers counterparty risk and makes long-term offtake agreements easier for lenders to accept.

Governments frequently prioritise hubs for a simple reason: they concentrate industrial transformation, emissions reductions, and infrastructure investment in one place. As a result, hubs often receive coordinated support through permitting, infrastructure planning, subsidies, and long-term policy commitments. When a government formally designates a region as a hydrogen hub, it sends a strong signal that the regulatory environment will remain stable—something financiers value highly when evaluating assets designed to operate for decades.

Momentum then begins to build. The first anchor customer attracts additional users, shared infrastructure lowers costs for new entrants, and each successful project makes the next easier to finance. What begins as a small industrial cluster can gradually evolve into a self-reinforcing hydrogen ecosystem.

In essence, hydrogen hubs solve one of the sector's most fundamental challenges: synchronising supply, demand, and infrastructure. By bringing these elements

together, hubs replace the "wait-and-see" standoff that often paralyses early hydrogen development with coordinated progress. For an industry seeking to move from pilot projects to gigawatt-scale deployment, hubs are not simply a helpful concept—they are the financial architecture that makes large-scale hydrogen investment possible.

A Real-World Analogy: Ulsan, South Korea — A Hydrogen Hub Already in Motion

If the idea of a hydrogen hub still feels abstract, it becomes much clearer when you look at Ulsan on the southeast coast of South Korea. Ulsan is Asia's most advanced example of a hydrogen ecosystem that already functions as a coordinated hub rather than a set of isolated projects. What makes it compelling is not just the technology, but the way it binds together industry, infrastructure, policy, and finance into a single economic logic.

Ulsan did not become a hydrogen hub by planning a single project. It became one because the city hosts a cluster of heavy industries—petrochemicals, automotive manufacturing, shipbuilding, and refineries—all concentrated in one industrial zone. These are energy-hungry sectors with large carbon footprints and a clear strategic interest in securing a low-carbon fuel that can be produced at scale. Rather than waiting for a national hydrogen market to emerge, the city began building one internally.

The catalyst came from Hyundai, headquartered in Ulsan. With its deep interest in fuel-cell vehicles and long-term commitment to hydrogen mobility, Hyundai began consuming hydrogen locally and simultaneously pushed for the supporting infrastructure: fuelling stations, logistics networks, and small-scale production assets. This initial industrial demand gave financiers the confidence that hydrogen produced in Ulsan would have a guaranteed home. It was not a speculative market; it was anchored by one of Korea's largest companies.

But Ulsan's hub status accelerated when the city began integrating hydrogen into multiple sectors at once. Chemical producers signed early offtake agreements for low-carbon hydrogen blends. The city partnered with Korea Gas Corporation (KOGAS) to develop plans for new hydrogen pipelines. Heavy industries explored replacements for grey hydrogen used in refining. Mobility, industrial heat, and chemical feedstocks all began pulling hydrogen from the same emerging ecosystem.

For banks, this diversification was crucial. They were no longer assessing a single project with a single buyer; they were looking at a regional hydrogen economy

where demand came from a constellation of industries. The offtake risk—the Achilles heel of early hydrogen finance—was spread across many actors rather than concentrated in one. This is why Ulsan became one of the first regions in Asia where hydrogen infrastructure received substantial private financing: the projects were not isolated bets, but components of a stable industrial system.

The hub also benefited from strong government backing. The Korean government designated Ulsan as the country's first "Hydrogen Leading City," unlocking regulatory support, co-financing for refuelling stations, and streamlined permitting. National policy aligned with local industrial logic: build supply where demand is already dense and build demand where infrastructure can be shared. This policy alignment significantly reduced long-term project risk, particularly the fear that regulations or support schemes might shift midstream—one of the greatest concerns for financiers evaluating a 20-year asset.

Today Ulsan functions as a coherent hydrogen hub. Production, storage, transport infrastructure, and end-use applications feed into one another. A molecule of hydrogen produced there might fuel a Hyundai truck, power a chemical plant, or flow into an industrial pipeline. The hub works because the system works: the value chain is connected end to end, and risks are distributed rather than concentrated.

For anyone trying to understand how hydrogen hubs help unlock financing, Ulsan is a real-world lesson. It shows how co-location reduces infrastructure costs, how diversified demand reduces counterparty risk, how industrial anchors create confidence, and how public policy de-risks the entire ecosystem. It demonstrates that a hydrogen hub is not a theoretical organising principle—it is a practical financial architecture. It is what happens when supply, demand, and infrastructure are built together rather than waiting for each other.

Ulsan proves that hubs are not just blueprints. They are living markets—and where markets exist, capital follows.

Bringing It Together: Turning Policy into Megawatts

By this point in the chapter, one thing should be clear: hydrogen is not held back by physics or a lack of enthusiasm. It is held back by economics, risk, and the absence of structures that make long-term commitments possible. Technology shows what can be done; finance decides what actually gets built.

Policy, pricing, subsidies, certification, contracts, hubs, investors – they can all sound like separate topics. In reality, they are parts of a single machine whose job is to do one thing: turn a risky, uncertain idea into a bankable asset that can raise

capital at scale. A CfD on its own is not enough. A subsidy without clear standards is not enough. A beautiful electrolyser next to a reluctant off taker is not enough. Only when these pieces interlock do we move from pilot projects and press releases to steel in the ground and electrons flowing.

Financing is the bridge between ambition and delivery. Policymakers can set targets, engineers can optimise efficiency, consultants can write strategies, but none of it matters if investors do not believe the cashflows. Hydrogen will not scale because it is morally appealing or geopolitically fashionable. It will scale because the people who allocate capital – banks, infrastructure funds, industrial balance sheets, public lenders – see projects whose risks are understood, whose revenues are protected, and whose economics remain robust across market cycles.

That is why this chapter has treated hydrogen not just as a set of molecules, but as a financial system. In a mature hydrogen economy, you will not only find gigawatt-scale electrolysers and cavern storage; you will find standardised contracts, credible price signals, interoperable certificates, and hubs where multiple industries share infrastructure and demand. You will find governments that have accepted their role as market makers, corporates that have moved from "wait and see" to "build and learn", and investors who treat hydrogen assets with the same pragmatic familiarity they now reserve for wind farms or gas pipelines.

The success of that system will depend on collaboration rather than heroics. Governments must design policies that reduce risk rather than simply announcing targets. Corporates must be willing to sign long-term offtakes and invest in their own transition, even when the first projects are uncomfortable. Investors must move beyond defensive decision-making and learn to read hydrogen's risk profile with the same discipline they apply elsewhere, recognising that today's unfamiliarity will not last forever.

When those elements click into place – when pricing mechanisms are credible, support schemes are predictable, standards are trusted, and hubs align supply with demand – capital will move. And when capital moves, megawatts follow. The hydrogen transition will not arrive in a single breakthrough or a single piece of legislation. It will arrive the way all major energy shifts do; through a slow accumulation of bankable projects, each one slightly less exotic than the last, until one day the financial system treats hydrogen not as an experiment, but as part of the ordinary business of keeping the world running.

At that point, hydrogen will finally become what it needs to be: not a promise, not a gamble, but an "Eco²" solution – ecological and economical at the same time.

10. The Current Challenges in Delivering This Future

"The best thing about the future is that it comes one day at a time." — Abraham Lincoln

Régis and the Long Way Around

Before sunrise, Montmartre is still brushing sleep from its eyes. The cafés have not yet rolled up their shutters, and the streets around Sacré-Cœur lie quiet except for the soft hum of a single Toyota Mirai parked on Rue Joseph de Maistre. Inside sits Régis, a taxi driver who came to Paris from Champagnole in the Jura seven years ago and found himself pulled into the city's pulse. He likes the unpredictability of the job, the way every day takes him through a different rhythm of boulevards, shortcuts, and the shifting moods of the capital.

A few years before the Paris 2024 Olympics, he joined the city's hydrogen taxi initiative — a push funded partly through European programmes to prove that zero-emission vehicles could become a viable part of urban transport. Companies like Hype began rolling out fuel-cell taxis, and Air Liquide supported the construction of hydrogen stations around Île-de-France. For a moment, it felt as though Paris was trying to leap ahead.

For Régis, the transition was easy. He loved driving the hydrogen car almost immediately. The smooth acceleration, the quiet cabin, the absence of diesel fumes — it all made the city feel less frantic, less abrasive. Passengers noticed too, often remarking on the car's silence or the oddly calming sensation of gliding through traffic. In a city where noise is constant, a quiet taxi feels like a small revolution.

But the revolution only went so far.

The Mirai is ready when he starts his day, but refuelling is never simple. Montmartre has no hydrogen station, nor does most of central Paris. The nearest

options are on the edges of the city or farther out in the metropolitan belt — sites developed to support bus depots or dedicated fleets, but not the everyday rhythm of taxi drivers. Régis knows the stations by the quality of traffic he must fight to reach them: Issy-les-Moulineaux to the south-west, Porte de Bercy to the east, a handful of others scattered like afterthoughts around the périphérique. None are close enough to fold easily into a day's work.

So, each week begins with the same detour. Before he collects his first fare, he drives down through Pigalle, merges into morning congestion, and heads out of the city just to fuel. It can take an hour. Sometimes more. And every kilometre he drives to refuel feels like the underside of the hydrogen promise — a reminder of how elegant the technology is, and how clumsy the infrastructure still feels beneath it.

The irony is not lost on him. He loves the car. He loves what it represents. Yet to keep it running, he must leave the very streets it serves so well.

Paris itself is caught in a policy argument that Régis feels only in the consequences. The city wants cleaner air and quieter streets; it has set ambitious decarbonisation goals and publicly embraced zero-emission mobility. But it has also been hesitant to commit fully to one pathway. Debates at City Hall pit battery-electric advocates against hydrogen supporters, with each camp arguing for limited public space, limited investment, and limited political attention. In that tug-of-war, hydrogen infrastructure often becomes the compromise that keeps being postponed. Plans are drawn, revised, questioned, delayed. And drivers like Régis navigate the uncertainty every day.

The city's grand hydrogen vision has had its own turbulence. Hype, once the flagship operator of Paris's hydrogen taxi fleet, announced in 2025 that rising hydrogen costs and inconsistent supply had pushed it to shift back toward battery-electric vehicles. Air Liquide had previously partnered on hydrogen stations in the region, including a key one in central Paris, but local pressures and regulatory disagreements resulted in at least one early station being removed. The momentum that once felt unstoppable became uneven, almost cautious.

Régis feels this personally. Not politically, not scientifically — practically. Every detour chips away at the promise that clean mobility should be convenient as well as sustainable. He doesn't complain to his passengers, but sometimes they ask why he drives so far to refuel. He tells them the truth: the technology works beautifully, but the support system hasn't caught up.

Still, he has no intention of giving up the Mirai. He likes how it moves. He likes how people respond to it. He likes knowing he is part of something that could become ordinary one day. If Paris builds more stations — not dozens, just enough to weave into daily life — the whole system would feel different. A station in the north, maybe near La Chapelle or Saint-Ouen, would change everything for drivers like him. Even Montmartre could host one discreetly, tucked into a garage or service depot.

Sometimes he imagines finishing a long shift and refuelling close to home instead of navigating the périphérique at dusk. He imagines a network that supports the people who actually use it — not just demonstration fleets or temporary projects tied to global events.

The city has made efforts. Through European funding, several new stations are planned, and the region continues to expand capabilities for buses, trucks, and cars. But expansion takes time, and until it materialises in the right places, taxis like Régis' are the ones who feel the gap most sharply. Not because the hydrogen system fails, but because it succeeds — and now the city must match that success with infrastructure proportionate to the ambition.

As he drives down Boulevard Haussmann later that morning, picking up a fare near the Opéra, Régis notices the way Parisians move between impatience and elegance, improvisation and structure. The streets are crowded, but he manoeuvres through them with ease, quietly satisfied with the responsiveness of his car. He thinks again about the future: a Paris where hydrogen vehicles glide past cafés and monuments as effortlessly as petrol cars once did, where refuelling is as ordinary as stopping for bread.

His shift ends in Montmartre, the sun dipping behind the rooftops. He parks the Mirai in its usual spot, listens to the gentle ping of the cooling battery, and locks the door.

He loves this taxi.

He only wishes the city loved it just a little more.

Régis' daily detour is not an anomaly. Variations of the same story surface repeatedly across the hydrogen landscape — in ports struggling to align terminals with supply, in industrial clusters waiting for infrastructure that has been approved but not yet built, and in mobility projects caught between technical readiness and political hesitation. The details differ, but the pattern is remarkably consistent: the technology is often ahead of the system meant to support it.

I was reminded of this recently while presenting at an industry conference. One slide — notably less technical than most — lingered on the screen longer than planned. As it listed the challenges facing hydrogen today, phones came out across the room to capture it. There was no visible disagreement, no rush to debate the framing. The reaction was quieter than that: recognition. The list reflected what many in the room were already living within their own projects — high costs, infrastructure gaps, uncertain demand, policy friction, financing risk, immature markets, coordination failures, and public scepticism.

It is that shared, largely uncontroversial reality that structures the rest of this chapter. The sections that follow break these challenges down, not to argue that hydrogen's future is unworkable, but to understand why progress feels uneven — and what must align for scale to become something more than a promise.

Taken together, these experiences point to a broader tension shaping hydrogen's progress today. The core technologies are proven, pilot projects are operating, and political ambition is visible across regions and sectors. Yet delivery remains uneven. Hydrogen is not struggling because the science is unclear, but because it is being asked to scale inside systems — planning, permitting, markets, and infrastructure — that were designed for a different energy era. The friction lies less in the molecule itself than in the institutions meant to support it.

It is against this backdrop that the current challenges must be understood. The sections that follow break them down not as isolated obstacles, but as interlocking constraints that collectively determine whether hydrogen remains confined to demonstrations or evolves into a mature, reliable part of the energy system.

Economics Before Scale

No issue shapes the hydrogen conversation more than economics. Cost is the first objection raised, the primary filter applied by investors, and the metric by which projects are most often judged. As explored in earlier chapters, the levelised cost of hydrogen and the difficulty of securing finance are not peripheral concerns; they are central to whether projects proceed at all. Any honest assessment of hydrogen's current position must begin by acknowledging that, today, it remains expensive.

Yet high cost, in this context, is as much a description of timing as of technology. Hydrogen's economics look weak because the system around it is incomplete. Electrolysers are built before networks exist, infrastructure is planned ahead of demand, and renewable generation must be secured long before stable offtake is guaranteed. In such an environment, investors are not pricing inefficiency; they

are pricing uncertainty. Capital is expensive not because the chemistry is flawed, but because volumes, utilisation, and long-term certainty are still emerging.

The absence of a liquid hydrogen market compounds this challenge. Without transparent pricing, standardised contracts, or dependable anchor demand, projects struggle to move beyond bilateral agreements and bespoke financing structures. Demand exists in principle across industry, mobility, and energy, but it remains fragmented and conditional, often waiting for lower prices that themselves depend on scale. This circularity is not unique to hydrogen, but it is particularly pronounced in a sector attempting to establish both supply and demand simultaneously.

Comparisons with competing technologies further complicate the picture. Batteries and direct electrification frequently appear cheaper and more mature, and in many applications, they are the correct solution. Hydrogen's economic weakness is most visible where it is applied prematurely or inappropriately. Where it addresses use cases that electricity struggles to serve — high-temperature industry, long-duration storage, or certain forms of heavy transport — the question is not whether hydrogen is viable, but when the surrounding system will allow its economics to mature. In that sense, hydrogen today is economically incomplete, not economically impossible.

Infrastructure Built for a Different Century

If hydrogen's economics reflect an incomplete system, its infrastructure reveals where that incompleteness becomes tangible. Unlike electricity, hydrogen cannot rely on invisible networks that already reach every building and vehicle. It must be produced, stored, moved, and delivered as a molecule — and the physical pathways required to do so are still sparse, fragmented, and often misaligned with where demand is emerging.

Most existing energy infrastructure was designed around two assumptions: that fuels would be centralised and dispatchable, and that demand would follow predictable, well-established routes. Hydrogen challenges both. Production is increasingly decentralised, tied to renewable resources that are geographically uneven, while demand spans sectors that rarely share infrastructure — industry, transport, power, and heat. The result is a system in which supply, transport, and end use struggle to evolve in step.

The consequences are felt long before any technical limit is reached. Storage options remain limited and expensive, transport routes are incomplete and refuelling or delivery points are often located where it is administratively con-

venient rather than operationally useful. As with Régis' detour across Paris, infrastructure exists — but not yet in the places or configurations that allow hydrogen to blend naturally into daily operations.

Until these physical gaps are addressed, hydrogen will continue to appear more difficult than it needs to be. Not because the molecule resists scale, but because the networks required to support it are still being imagined, permitted, and built — often one project at a time, without the coherence that mature energy systems take for granted.

Policy, Permitting, and Time

As hydrogen projects move from concept to construction, the most stubborn barriers increasingly emerge outside the engineering scope. Permitting processes are slow, regulatory frameworks remain incomplete or inconsistent, and standards are often still under development. In many regions, it now takes longer to secure approvals than it does to build the assets themselves. This is not unique to hydrogen, but its cross-sector nature amplifies the friction.

Policy ambition has grown faster than institutional capacity. Governments signal support through strategies, targets, and funding programmes, yet implementation frequently stalls at the local and administrative level. Hydrogen projects must navigate environmental permitting, safety regulation, grid access, land use planning, and industrial approvals — often across multiple agencies that were never designed to coordinate closely. The result is uncertainty that is difficult to price and timelines that are hard to trust.

Examples of this kind of regulatory friction increasingly appear as hydrogen projects move from concept to construction. In one project I was involved in, engineers working on a relatively small hydrogen installation were required to implement a control architecture typically used for much larger industrial plants. From an engineering standpoint, the system already had the capability needed to operate safely and reliably, and those responsible for the design were confident it met the technical requirements of the facility. Yet existing regulations mandated the more complex structure because the rules had been written with conventional large-scale process plants in mind. The result was not a meaningful improvement in safety or performance, but additional cost and complexity imposed on a project that did not technically require it.

Situations like this illustrate a broader challenge. Much of the regulatory landscape governing energy and industrial infrastructure was developed for large, centralised facilities built around long-established technologies. While some

hydrogen projects will indeed resemble traditional industrial plants, many others are likely to be modular, distributed, and smaller in scale. When rules designed for a different industrial model are applied without adaptation, they can unintentionally slow deployment and increase costs — not because the technology demands it, but because the institutional frameworks around it have yet to evolve.

This disconnect is not due to a lack of engagement or intent. I was reminded of this recently at a dinner hosted at the Baden-Württemberg state parliament in Stuttgart, where the conversation at the table — which included senior political leadership, industrial executives, and representatives of regional commerce — centred almost entirely on the challenges of delivering a single hydrogen refuelling station in Ulm. What stood out was not resistance, but commitment. There was genuine passion around the table, a shared recognition of hydrogen's importance, and a clear desire to make progress. Yet even in that environment, the discussion returned repeatedly to permitting pathways, regulatory interfaces, and the difficulty of aligning responsibilities across institutions.

Time itself, in this context, becomes a risk factor. Hydrogen infrastructure requires long development cycles and patient capital, while political incentives favour short-term results and visible wins. Projects are announced quickly, but progress unfolds slowly, creating a growing gap between expectation and delivery. When delays accumulate, scepticism follows — not because the technology has failed, but because governance has struggled to keep pace with ambition.

Coordination, Confidence, and Credibility

Hydrogen occupies an uncomfortable space between sectors that have rarely needed to cooperate closely. Power generation, industry, transport, infrastructure operators, regulators, and financiers all play critical roles, yet no single actor controls the system end to end. In the absence of shared standards, interoperable designs, and clear responsibility for coordination, progress often fragments into isolated projects rather than cumulative systems.

This fragmentation affects confidence as much as efficiency. Safety standards evolve at different speeds across jurisdictions, technical specifications vary between suppliers, and contractual structures are frequently bespoke. None of this is unusual for an emerging sector, but together it creates friction that slows replication and undermines trust. Projects work, but they do not yet feel routine — and routine is what markets, regulators, and the public ultimately rely on.

Public perception amplifies these challenges. Hydrogen carries historical baggage and has suffered from cycles of inflated expectation followed by disappointment.

Safety concerns, while generally manageable, are easily magnified in the absence of familiarity, and scepticism grows when ambitious announcements outpace visible delivery. For some, hydrogen still carries a cultural association shaped less by modern engineering than by historical imagery — the Hindenburg remains a reference point in the public imagination, despite having little relevance to how hydrogen is produced, stored, or used today. The result is a credibility gap: not rooted in technical failure, but in inconsistency, overpromising, and the slow translation of pilots into everyday infrastructure.

Until hydrogen feels coordinated, standardised, and unremarkable, it will continue to face resistance that has little to do with performance. Confidence is built not through novelty, but through repetition — through systems that behave predictably, quietly, and safely enough to fade into the background of daily life.

Sequencing and Timing

Many of hydrogen's challenges ultimately converge on a single issue: sequencing. Supply is developed ahead of demand, infrastructure ahead of users, and policy frameworks ahead of practical delivery mechanisms. Each step makes sense in isolation, yet misalignment between them creates a sense of perpetual incompleteness — projects that work technically but struggle to mature commercially.

This sequencing problem traps much of the sector in what might be described as a permanent demonstration phase. Pilots succeed, lessons are learned, and confidence grows, yet scale remains just out of reach. Investors wait for certainty, users wait for price reductions, and policymakers wait for proof of demand. In the absence of coordination, progress advances incrementally rather than cumulatively.

The challenge, then, is not identifying what needs to be built, but deciding what must come first — and accepting that some elements must move ahead of clear market signals. Until timelines across policy, infrastructure, finance, and end use are better aligned, hydrogen will continue to feel slower than its technical readiness would suggest. The issue is not momentum, but synchronisation.

A Solvable Set of Problems

Taken together, the challenges outlined in this chapter are neither surprising nor exceptional. Each time a new energy system has been brought online, its early progress has been shaped less by technical uncertainty than by the difficulty of integrating it into structures built for what came before. Hydrogen's current friction fits squarely within that pattern. It reflects not doubt about whether the

technology works, but the challenge of adding a new energy vector into systems never designed to accommodate it.

What is striking is how few of these barriers are fundamentally insoluble. Costs remain high because volumes are low, infrastructure lags because development is uneven, and confidence wavers because policy, finance, and delivery move at different speeds. These are not mysteries of physics or chemistry. They are coordination problems — complex, politically sensitive, and slow to resolve, but ultimately responsive to clear signals, sustained commitment, and capable leadership.

History suggests that progress in energy systems tends to be cumulative rather than abrupt. New technologies gain ground through repetition: projects that operate reliably, standards that stabilise, institutions that learn, and people who carry experience from one deployment to the next. Where alignment begins to form — between policymakers willing to accept longer timelines, investors prepared to tolerate early uncertainty, and operators focused on reliability rather than novelty — momentum becomes structural rather than fragile.

From the outside, this process can look frustratingly slow, as though the sector is moving from one pilot project to the next without ever quite breaking through. Yet these pilots are not detours; they are the groundwork through which trust is built, operational knowledge accumulates, and unfamiliar technologies become routine. Everyone involved wants hydrogen to move faster, but speed without foundations would only recreate the very failures that sceptics fear.

Seen in that light, the barriers facing hydrogen are not signs of failure, but markers of a system in transition. They indicate how much alignment is required to add a new energy system at scale, and how deliberate that process must be. The question now is less whether hydrogen can function, and more whether the will exists to address these challenges with the patience, coordination, and seriousness they demand.

11. The Social and Environmental Dimensions

"We do not inherit the Earth from our ancestors; we borrow it from our children." — Indigenous proverb

Up to this point, much of this book has been written from an engineer's perspective. It has been grounded in data, systems, cost curves, and physical constraints. The arguments have followed a largely rational path: if the technology works, if the economics close, and if the policy framework is aligned, then progress should follow. That way of thinking is not wrong — but it rests on an assumption that human systems respond to logic in the same way technical systems do.

As discussed earlier in the context of battery electric vehicles and hydrogen, this assumption often breaks down. People do not adopt technologies solely because they are efficient or optimal; they adopt them because they feel familiar, intuitive, and compatible with everyday life. As Rory Sutherland has argued, what appears irrational when viewed through engineering or economic models is often entirely coherent when viewed through the lens of human perception, habit, and emotional cost. In *Alchemy: The Dark Art and Curious Science of Creating Magic in Brands, Business and Life*, he makes the case that behaviour frequently follows narratives and heuristics rather than linear logic — not as a flaw, but as a feature of how humans navigate complexity.

Energy transitions therefore do not unfold on spreadsheets alone. They unfold in towns, landscapes, workplaces, and communities. They are shaped not just by what is technically possible or economically efficient, but by how change is perceived, experienced, and absorbed into daily life. There is always a silent variable in any transition model — one that cannot be optimised or engineered away — and that is society itself.

This becomes visible in small, everyday conversations. When I speak to my mother about developments taking place in her part of the world, her response is often frustration: projects are delayed, infrastructure is misplaced, change feels

intrusive or poorly thought through. In those moments, it is tempting to invoke the familiar accusation of "Not In My Back Yard" (NIMBY) — to suggest that opposition reflects inconsistency or resistance to progress. And yet, that framing quickly breaks down. Nearing eighty, she is not weighing system efficiency or long-term optimisation. She is responding to loss of familiarity, diminished control, and the disruption of a world she has spent a lifetime navigating.

She once summed up this perspective far more succinctly than I ever could. When discussing the scale of change taking place around us, she remarked: "I'm happy that I'm in departures and not in arrivals."

There was humour in the statement, but also an uncomfortable truth. From her perspective, many of the benefits of the energy transition will be realised by future generations, while much of the disruption will be experienced by those living through it today. Whether one agrees with that view or not, it illustrates an important point: people often evaluate change through the lens of their own lived experience rather than through abstract discussions of long-term societal benefit.

I am not immune to this dynamic myself. Where I live, I am close to the proposed Jura-Leitung — a high-voltage direct current transmission line intended to carry renewable electricity from northern to southern Germany. I understand exactly why it is being built. I agree with the logic behind it. I can articulate the system need, the decarbonisation benefits, and the role such infrastructure plays in stabilising a future energy system.

Yet, I must admit that I am not particularly enthused by the idea of having it routed within a few kilometres of my home. That reaction sits uncomfortably with me. I am acutely aware of the irony: someone who spends his professional life arguing for system-level change, confronted with a tangible example of that change arriving close to his own front door. I find myself having to actively remind myself that opposing infrastructure simply because it is nearby — even when one agrees with its purpose — is precisely the behaviour so often criticised in others.

The instinctive question is not whether the project should exist, but why here?

That question is not irrational. It reflects the same human calculus seen everywhere in energy transitions: a tension between abstract benefit and personal impact, between collective necessity and local disruption. Understanding the system does not eliminate that tension — it merely makes it more visible.

There is even a kind of dark humour that develops around such projects. In discussions about whether Germans have a sense of humour, I sometimes answer "yes" — and then offer a single word: "Frankenschnellweg".

For readers outside Bavaria, this may sound meaningless. For many in Franconia, it prompts an immediate, knowing smile. The term has become shorthand for a piece of infrastructure whose development has been debated, redesigned, opposed, delayed, and litigated for so long that it has slipped from planning process into local folklore.

What began as a transport project has, over decades, turned into a symbol of how development can be hijacked — not by a single objection or decision, but by the cumulative weight of procedural friction, local resistance, shifting priorities, and institutional caution.

The humour matters, because it reveals something important: when delays become normalised to the point of satire, they stop being perceived as failures of the system. They become expectations. And once t hat happens, even projects that are broadly supported in principle struggle to move from agreement to execution.

That reaction is often labelled irrational. In reality, it reflects a different set of priorities.

This tension is not a personal quirk — it is the essence of the challenge facing the energy transition. Technologies that make perfect sense on paper can provoke resistance precisely because they ignore how people assess risk, value stability, and protect identity. Logical arguments may explain why change is necessary; they do not, on their own, make change acceptable.

The shift to a low-carbon energy system will not succeed through technology, finance, or policy alone. It will succeed only if society is willing to accommodate new infrastructure, accept new trade-offs, and trust that the costs and benefits of change are being shared fairly. Where that trust breaks down, projects stall, opposition hardens, and political support weakens — regardless of how compelling the technical case may be.

This chapter therefore moves away from electrolysers, balance-of-plant, and levelised costs, and towards the human dimension of the hydrogen economy. It examines water use, land requirements, lifecycle impacts, jobs, materials, and social licence — not as peripheral concerns, but as central constraints. This is because an energy transition that is logically sound, but psychologically mis-

aligned, will not be delivered — no matter how elegant the engineering behind it may be.

Social Licence: When Energy Moves into Everyday Life

Much of the discussion around the hydrogen economy focuses on systems, costs, and technical feasibility. These elements are necessary, but they are insufficient. The concept often used to describe what sits beyond them — social licence to operate — is frequently treated as secondary. In practice, it is a binding constraint.

Social licence is informal, fragile, and difficult to recover once lost. It is not granted by a regulator, nor secured by compliance alone. Instead, it emerges through trust: trust that risks are understood and managed, that benefits are shared fairly, and that local voices matter in decisions that reshape the spaces in which people live and work. Where that trust is absent, opposition hardens, projects slow, and political support weakens — regardless of how strong the technical or economic case may be.

Energy infrastructure is uniquely sensitive to this dynamic because it is inherently spatial. Pipelines, storage facilities, electrolysers, ports, and transmission lines must exist somewhere, and their impacts are experienced most acutely by those nearby. While the benefits of decarbonisation are diffuse and long-term, the disruptions associated with new infrastructure are immediate and local. Noise, land use, safety concerns, and construction impacts are not abstract risks; they are lived realities.

This imbalance between localised risk and distributed benefit is often reduced to the shorthand of NIMBYism. Yet this framing obscures more than it explains. In most cases, opposition is not rooted in ignorance or ideological resistance to change. It reflects a rational response to asymmetry: communities are asked to absorb disruption in exchange for benefits that may accrue elsewhere, arrive later, or remain uncertain.

Hydrogen infrastructure intensifies these tensions because it challenges long-standing assumptions about where energy is produced. For much of the last century, energy generation has been distant from daily life. Power stations, refineries, and heavy industrial assets were placed on the edges of cities or far beyond them, reinforcing a psychological separation between energy production and energy use. Risk, such as it was, remained largely out of sight.

In many applications, hydrogen is often most practical — and sometimes most efficient — when produced close to the point of demand. As a result, electrolysers are increasingly proposed within environments that are part of everyday experience: industrial facilities, transport depots, ports, hospitals, and municipal infrastructure. For many people, this will be the first time energy production is encountered not as a remote system, but as a visible, continuous process occurring nearby.

This proximity matters. Familiar risks are often tolerated with little reflection — a diesel generator powering a temporary event, a fuel tank behind a building — while unfamiliar technologies attract heightened scrutiny. The issue is not that hydrogen introduces unprecedented danger, but that it is new, visible, and poorly embedded in public memory. Where energy once arrived invisibly through cables and pipes, it now appears as equipment, process, and presence.

Over time, such infrastructure may become normalised, just as previous forms of energy production have been. But during the transition phase, visibility amplifies concern. Questions of safety, ownership, accountability, and control become central, particularly when infrastructure is introduced into spaces associated with care, work, or public service. Communities are no longer being asked simply to support an energy transition in principle; they are being asked to live alongside it.

For project developers and investors, the implications are concrete. Social opposition translates into planning delays, legal challenges, permitting uncertainty, and rising costs. These are not abstract social risks, but material factors that shape project viability. A hydrogen project that appears robust on a spreadsheet can quickly become unfinanceable if its social foundations are weak.

Social licence, therefore, cannot be treated as an afterthought or managed through communication alone. It must be built into project design from the outset, influencing choices around location, scale, ownership, and benefit sharing. The question is not how to persuade communities to accept hydrogen infrastructure, but how to design systems that give them good reason to do so.

The sections that follow examine the points at which this licence is most often tested — water use, land and place, employment, lifecycle impacts, and material supply chains. Each represents a moment where the abstract logic of the energy transition collides with lived experience. Understanding these collisions is not about slowing progress, but about ensuring that the hydrogen economy can be delivered at all.

Water: Quantity, Quality, and Competition

Water is often discussed in hydrogen debates in terms of volume. At a global level, the quantities required for electrolysis appear modest when set against agriculture or thermal power generation. This observation is broadly correct — but it is also incomplete. For hydrogen production, the critical issue is not only how much water is available, but what condition it is in, and what must be done to make it usable.

Electrolysers require water of very high purity. Impurities, dissolved salts, and organic matter degrade performance, reduce efficiency, and shorten equipment lifetimes. As a result, most hydrogen production systems cannot rely directly on the water sources that are locally available to communities. Surface water, groundwater, or even municipal supplies typically require significant treatment before they are suitable for electrolysis. This introduces an additional layer of infrastructure, energy use, and cost — and, critically, an additional point of social sensitivity.

From the perspective of local communities, this distinction matters. A project that claims modest water consumption may still place pressure on shared treatment facilities, compete for potable water, or require dedicated purification systems that alter local water flows. In regions where water access is already contested, the perception that high-quality water is being diverted — or further processed — for industrial energy use can quickly undermine trust, regardless of the absolute quantities involved.

The experience of other global industries offers a cautionary parallel. Beverage producers, most notably companies such as Coca-Cola, have faced sustained criticism and resistance in regions where their operations draw heavily on local water resources. In these cases, conflict has rarely been driven by total water use alone. Instead, it has centred on quality, priority, and control: who has access to clean water, who decides how it is allocated, and who bears the consequences when availability declines. Hydrogen projects risk triggering similar dynamics if water treatment and sourcing are not handled with transparency and local accountability.

This challenge becomes even more pronounced in regions frequently highlighted as ideal hydrogen exporters. Large-scale visions of solar-powered electrolysis in arid or semi-arid regions often overlook a basic constraint: sunlight is abundant, but water is not. Producing hydrogen in desert environments requires either long-distance water transport or desalination, both of which carry economic, environmental, and social costs. Without reliable and socially accepted water

solutions, such projects remain theoretical, regardless of the quality of the renewable resource.

Desalination is frequently presented as the answer to this problem, particularly for coastal developments. While it can provide a technically viable source of feedwater, it does not eliminate the water challenge — it transforms it. Desalination increases energy demand, raises questions about brine disposal, and adds infrastructure that must itself be accepted by local communities. Moreover, even desalinated water must often undergo further purification to meet electrolyser specifications, reinforcing the point that water for hydrogen is not simply abstracted, but processed.

These layers of treatment and conversion are largely invisible in high-level energy models, yet they are highly visible at the local scale. Pipelines, treatment plants, discharge points, and storage facilities become part of the physical and social landscape. For communities already managing water stress, the arrival of energy infrastructure that depends on high-quality water can feel less like a climate solution and more like an intensification of existing pressures.

The risk for the hydrogen economy is not that water requirements are insurmountable, but that they are underestimated in social terms. Water is a uniquely sensitive resource, bound up with health, livelihoods, and identity. Where access to clean water is perceived to be compromised, opposition is swift and durable. Projects that treat water as a purely technical input are therefore likely to encounter resistance, even when their environmental credentials are strong.

For hydrogen to scale sustainably, water strategy must be explicit, local, and credible. This means recognising that water availability, treatment, and governance are inseparable, and that solutions which appear elegant at a system level may falter if they fail to account for local realities. Hydrogen does not exist outside water systems; it is embedded within them. A transition that ignores this fact risks repeating the mistakes of other resource-intensive industries — and losing the social licence it depends upon.

Land, Place, and Industrial Memory

Energy infrastructure does not arrive on blank land. It is inserted into places shaped by history, identity, and lived experience. Land is not merely a technical input to be optimised; it is a social asset, carrying memories of past use, past promises, and past consequences. For the hydrogen economy, this reality is unavoidabl because much of its infrastructure is large, visible, and long-lived.

Electrolysers, compressors, storage facilities, pipelines, and export terminals occupy space in ways that cannot be abstracted away. Even when carefully designed, they alter how land is used and how it is perceived. In urban or peri-urban settings, this may trigger concerns about safety, noise, and visual intrusion. In rural areas, it can clash with agriculture, conservation, or landscape protection. In both cases, the physical presence of infrastructure becomes a focal point for wider anxieties about change, control, and fairness.

These reactions are rarely shaped by the proposed project alone. Many regions now identified as promising locations for hydrogen development carry deep industrial memory. Former coalfields, steelmaking centres, ports, and refining hubs often retain a complex legacy of employment, environmental degradation, and unfulfilled commitments. For communities that have experienced repeated cycles of industrial expansion and contraction, new infrastructure is rarely judged on its own merits. It is interpreted through what came before.

In such contexts, scepticism is not opposition to progress; it is learned caution. Promises of regeneration, jobs, and long-term benefit have often failed to materialise, while environmental and social costs remained local. Hydrogen projects proposed in these regions may therefore encounter resistance not because they are seen as dangerous, but because they resemble familiar patterns of extraction, disruption, and external control.

This dynamic complicates the common assumption that reusing existing industrial land is politically straightforward. From a system perspective, brownfield sites appear ideal: grid connections exist, transport corridors are established, and industrial activity is not new. Yet for local communities, these sites may represent unresolved harm. Re-industrialisation without acknowledgement of past impacts risks reopening old wounds, particularly where decision-making remains distant and opaque.

At the same time, industrial memory can also work in the opposite direction. In some regions, continuity rather than novelty offers a pathway to acceptance. The North Sea provides a compelling example of how existing energy landscapes can be repurposed rather than abandoned.

Decades of offshore oil and gas development have left behind an extensive network of platforms, pipelines, ports, service bases, and skilled workforces. As hydrocarbon production declines, much of this infrastructure faces decommissioning and removal — a process that is costly, environmentally disruptive, and socially sensitive. In parallel, the North Sea has become one of Europe's most

important regions for offshore wind, creating a convergence of renewable generation, industrial capability, and legacy energy assets.

From a technical and economic perspective, the case for repurposing is strong. Offshore wind can provide large volumes of electricity close to potential hydrogen production sites. Pipeline corridors already connect offshore fields to onshore industrial clusters. While many assets would require modification — including material assessment, re-lining, and new compression — the fundamental spatial logic is already in place. Ports that once served drilling and maintenance operations are increasingly supporting wind deployment and could play a similar role for hydrogen handling, storage, and export.

Crucially, this form of transition builds on familiarity rather than disruption. Offshore energy infrastructure is not new to North Sea communities. Ports, fabrication yards, and service bases have coexisted with energy production for generations. Repurposing assets in this context does not introduce industrial activity into previously untouched landscapes; it extends an existing energy identity into a new phase. For workers and communities accustomed to offshore energy, hydrogen and wind can be understood as evolution rather than replacement.

This continuity matters socially as much as technically. Where greenfield developments require social licence to be negotiated from scratch, repurposed infrastructure can inherit a degree of acceptance — provided it is handled transparently and fairly. Skills can be redeployed, assets reused, and regions remain central to the energy system rather than slipping into managed decline. In contrast to transitions framed around closure and loss, repurposing offers a narrative of adaptation.

There is also a broader, often overlooked, dimension to this approach. Much opposition to oil and gas has been shaped not only by emissions, but by what those industries have come to symbolise: environmental damage, extractive economics, and delayed action. Transition pathways that simply discard legacy assets risk reinforcing a sense of rupture and blame. By contrast, visibly transforming existing infrastructure to support low-carbon energy reframes that narrative. It suggests that past industrial systems can be cleaned, adapted, and redirected rather than erased.

For those concerned with the environmental legacy of fossil fuels, this matters. A transition that repurposes platforms, pipelines, and ports into components of a low-carbon system can feel more credible than one that displaces impacts elsewhere. It allows regions long associated with emissions to become demonstrably

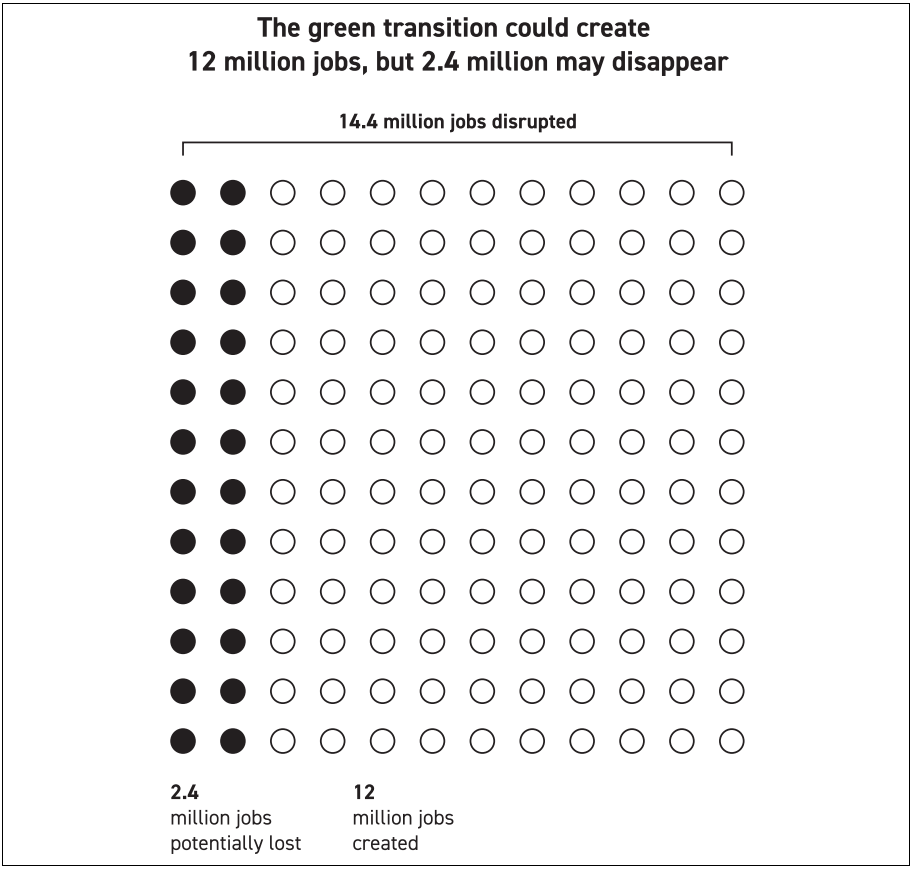
part of the solution, reducing cultural and political friction around new energy infrastructure.

None of this removes the need for careful design, local engagement, or fair distribution of costs and benefits. Repurposing that preserves private value while socialising risk will not build lasting trust. But where it is done well, it illustrates a wider principle: land use decisions are also narrative decisions. They shape how the transition is understood, remembered, and ultimately accepted.

Hydrogen will need space — offshore and onshore, coastal and industrial. Whether that space becomes a source of conflict or a foundation for continuity will depend less on engineering drawings than on how land, history, and identity are treated in practice.

Jobs, Work, and the Meaning of a "Just Transition"

Discussions of the energy transition often lean heavily on job numbers. Reports highlight millions of roles created, new sectors emerging, and long-term employment growth driven by clean energy investment. At a macroeconomic level, these projections may be accurate. But at the level where social licence is earned or lost, aggregate numbers matter far less than individual trajectories.



Adapted from World Economic Forum analysis; illustration redesigned for this book.

This is the view from 30,000 feet — useful, but rarely how transitions are experienced on the ground.

Global assessments of the energy transition often reflect this tension. For example, World Economic Forum analysis suggests that while millions of jobs may be lost in legacy energy sectors, a larger number will be created across clean energy, infrastructure, and associated industries. At an aggregate level, this points to net job creation. But such figures conceal as much as they reveal. Jobs are not lost and created in the same places, at the same time, or by the same people.

For most people, work is not simply a source of income. It provides identity, routine, status, and belonging. Communities form around industries; skills are passed between generations; local economies evolve to support particular forms of labour. When energy systems change, they do not merely reallocate capital — they disrupt lives. A transition that overlooks this reality risks being perceived as efficient but indifferent.

Hydrogen exemplifies this tension. Compared to the fossil energy systems it may replace, hydrogen production is typically more capital-intensive and less labour-intensive. Electrolysers, compressors, and storage facilities require fewer workers once operational, and those roles tend to be more specialised. This does not make hydrogen undesirable, but it does complicate claims that it will straightforwardly replace lost employment in legacy energy regions.

The concept of a "just transition" emerged to address this gap, but it is often invoked without precision. Too frequently, it is reduced to retraining programmes or future job promises detached from local context. Retraining can be effective for some workers, some of the time. It is far less effective when it assumes that skills, identities, and life circumstances are infinitely flexible.

A fifty-year-old offshore technician, a refinery operator, or a port worker is not simply a bundle of transferable competencies. They have mortgages, families, and social ties rooted in place. Asking such workers to relocate, requalify, or wait years for new opportunities is not a neutral request; it is a significant personal gamble. When transition narratives fail to acknowledge this, resistance is a rational response rather than an ideological one.

Where hydrogen offers a more credible pathway is in contexts where continuity is possible. Repurposing existing assets and redeploying existing skills — as seen in offshore regions, industrial ports, and logistics hubs — reduces the personal cost of change. Workers familiar with safety-critical environments, process industries, and energy systems are often well positioned to adapt to hydrogen-related roles, particularly when transitions are planned rather than reactive.

However, even in these cases, timing matters. Jobs lost today cannot be replaced by promises of employment tomorrow. A gap of several years between decline and opportunity is enough to erode trust, hollow out communities, and fuel political backlash. For regions that have already experienced industrial decline, patience is in short supply.

This is why the distribution of benefits matters as much as their existence. Projects that create value locally — through stable employment, local procure-

ment, and long-term operational presence — are more likely to earn acceptance than those that extract value while externalising risk. Conversely, hydrogen developments perceived as capital-intensive enclaves, owned and controlled from elsewhere, may struggle to gain support even when they align with climate objectives.

There is also a qualitative dimension to work that is often overlooked. Jobs associated with the energy transition are frequently framed as inherently virtuous, while legacy energy work is portrayed as something to be left behind. For many workers, this framing feels dismissive. It implies that decades of skilled labour were morally suspect, rather than socially necessary at the time. A transition that succeeds in economic terms but alienates those who built the existing system risks deepening cultural divides.

A more durable approach recognises that the skills, discipline, and experience embedded in legacy energy sectors are assets, not liabilities. Hydrogen, in particular, benefits from expertise in gas handling, process control, safety management, and industrial maintenance. Treating this continuity with respect helps shift the narrative from replacement to evolution — and from loss to contribution.

Ultimately, a just transition is not defined by job counts in distant futures, but by whether people can move from one system to another without falling out altogether. It is measured in continuity of livelihood, dignity of work, and confidence that change is being done with people rather than to them.

For the hydrogen economy, this is not a peripheral concern. Workforce acceptance shapes permitting, political support, and long-term stability. Where workers see themselves reflected in the future energy system, projects gain momentum. Where they do not, opposition hardens. Social licence, once again, is earned not through aspiration, but through lived experience.

Lifecycle Emissions: When “Green” Is Conditional

Few words are used as casually in the energy transition as “green”. It appears in strategies, press releases, and political speeches, often without qualification. Yet for technologies as complex and interconnected as hydrogen, greenness is not a fixed property. It is an outcome — dependent on assumptions, boundaries, and timing.

At the point of use, hydrogen appears deceptively simple. When consumed in a fuel cell or combusted under controlled conditions, it produces no carbon dioxide. This characteristic underpins much of its appeal and explains its prominence

in decarbonisation strategies. But social acceptance is not built on tailpipe emissions alone. It depends on whether the full system withstands scrutiny.

Lifecycle emissions expose that system.

Hydrogen's climate impact is shaped not by a single decision, but by a chain of choices: how electricity is generated, how electrolysers are manufactured, how water is treated, how hydrogen is stored and transported, and how end-use infrastructure is built and maintained. Each step carries embedded emissions, often invisible to the end user but highly relevant to public trust.

This complexity matters because hydrogen projects are frequently deployed in phases. Early projects often rely on grid electricity that is not yet fully decarbonised, imported equipment with carbon-intensive supply chains, and construction materials such as steel and concrete whose emissions are front-loaded. In accounting terms, this can mean that initial hydrogen production delivers only partial emissions reductions — or, in some cases, very little at all.

From an engineering perspective, this is neither surprising nor disqualifying. Transitions unfold over time. Systems improve as grids decarbonise, manufacturing scales, and supply chains adapt. But from a societal perspective, this temporal logic is far less forgiving. Communities are asked to accept new infrastructure now, while promised climate benefits arrive later. When that gap becomes visible, trust erodes quickly.

The risk, therefore, is not that lifecycle emissions invalidate hydrogen as a pathway. The risk is that claims outpace delivery. Public narratives that present early hydrogen projects as fully green invite scrutiny that immature systems struggle to withstand. Once credibility is lost, it is difficult to recover — even when performance improves.

This is where transparency mechanisms begin to matter. As discussed earlier in the context of financing, emerging certification and traceability systems offer a way to anchor claims in verifiable data rather than aspiration. Digital product passports, lifecycle tracking tools, and guarantees of origin make it possible to link hydrogen production to specific inputs, locations, and timeframes. For investors, this reduces uncertainty. For society, it provides something more fragile: confidence that progress is real and measurable. The European Union's move toward battery passports and tighter disclosure requirements reflects this shift. Batteries, having scaled earlier and attracted scrutiny sooner, are several years ahead of hydrogen in this respect — but the direction is clear. As hydrogen scales, similar expectations of transparency and traceability should be anticipated.

There is also a broader narrative risk that lifecycle emissions expose. A familiar image often circulates in debates about electrification: a battery electric vehicle charging cleanly at the curb, while the cable disappears into a distant power station emitting carbon dioxide. The implication is that electrification offers a comforting illusion of cleanliness while obscuring upstream emissions. Although the image is rhetorically powerful, it is analytically incomplete. It is not an argument against electric vehicles; it is a reminder that system boundaries matter.

The same logic applies to hydrogen. A fuel that appears clean at the point of use can only claim legitimacy if its upstream system is visible, accountable, and improving over time. When upstream impacts are hidden — whether in power generation, manufacturing, or material extraction — scepticism grows. This is not cynicism; it is a rational response to past experience.

There is also a spatial dimension to lifecycle emissions that intersects directly with social licence. Many of the emissions associated with hydrogen occur far from the point of use: in manufacturing plants, mining regions, or power systems in other jurisdictions. For local communities, this creates a disconnect. Infrastructure arrives nearby, while environmental impacts appear to be displaced elsewhere. When decarbonisation feels like relocation rather than reduction, accusations of greenwashing become harder to counter.

This sensitivity is particularly acute in regions with a history of heavy industry. Communities that have lived with pollution, degradation, and broken assurances are often acutely aware of lifecycle trade-offs. They are less likely to accept abstract arguments about future decarbonisation if present-day impacts remain tangible.

A more resilient approach treats lifecycle emissions as a design constraint rather than a communications problem. This means being explicit about starting conditions, improvement pathways, and trade-offs. It means acknowledging when hydrogen is lower-carbon rather than zero-carbon and explaining how and when further reductions are expected to occur. Such honesty may feel politically risky, but it is socially stabilising.

Over time, hydrogen systems can and should become genuinely low-carbon across their full lifecycle. Electrolyser efficiency will improve. Renewable generation will expand. Manufacturing processes will decarbonise. None of these developments benefit from exaggerated claims in the present. Social licence is strengthened not by perfection, but by transparency.

Lifecycle honesty, in this sense, is not a weakness. It is a prerequisite for scale.

Materials, Mining, and Displaced Environmental Impact

Lifecycle emissions broaden the frame of the hydrogen economy beyond the point of use. Materials widen it further — and often uncomfortably so. Every energy system rests on a physical foundation of metals, minerals, and processed materials. The energy transition does not escape this reality; it reorganises it.

Public debate around clean energy tends to focus on visible infrastructure: vehicles, turbines, electrolysers, charging points. Far less attention is paid to the material systems that make these technologies possible. Yet for many communities around the world, the most tangible impacts of the transition are not felt where energy is consumed, but where materials are extracted, refined, and transported.

This distinction matters because the environmental and social costs of mining are rarely borne by the societies that benefit most directly from decarbonisation. Water stress, land disruption, waste tailings, and labour conditions are concentrated in specific regions — often far removed from end users and subject to weaker regulatory oversight. As a result, transitions that appear clean at the point of use can feel extractive at the point of supply.

Hydrogen is often described as materially light, and in relative terms this is accurate. Its primary inputs are electricity and water. Compared with battery-electric pathways, hydrogen systems typically rely on significantly smaller volumes of critical materials per unit of energy delivered, with demand concentrated in a narrow set of components. Fuel cells and electrolysers depend on platinum-group metals, specialised steels, and copper, but these materials are used in small quantities, embedded in durable equipment, and recycled at high rates. The challenge here is not scale of extraction, but concentration of supply.

Battery-electric systems follow a different material logic. They rely on much larger volumes of a broader range of critical minerals — including lithium, nickel, cobalt, manganese, copper, and graphite — deployed at scale across millions of vehicles and storage systems. Most are not rare earth elements in the strict sense, but they create a similar challenge: sustained pressure on global mining and refining systems, with cumulative environmental and social consequences.

This contrast can be understood as a trade-off between volume and intensity. Battery pathways scale by volume, expanding extraction across multiple materials and regions. Hydrogen pathways scale by infrastructure, concentrating material demand in fewer, more specialised supply chains. Neither approach is impact-free. They simply distribute burden differently.

This distinction is rarely visible in public discourse, and that invisibility matters. Battery-electric vehicles are typically experienced as benign and familiar at the point of use, while the mining systems that support them remain distant and abstract. Hydrogen, by contrast, is highly visible: electrolysers, storage tanks, pipelines, and industrial facilities are difficult to ignore, even when their upstream material footprint may be smaller on a per-unit basis. Visibility attracts scrutiny; displacement often escapes it.

This asymmetry helps explain why debates around clean technologies can become polarised. Technologies that hide their material intensity behind consumer familiarity are often perceived as unambiguously green. Technologies that expose their infrastructure — even when materially efficient — are more readily contested. The result is not a rational assessment of impacts, but a selective one shaped by proximity and perception.

Over time, this imbalance is unlikely to persist. As battery deployment accelerates, attention is already turning toward lithium brine extraction, cobalt mining conditions, nickel processing emissions, and the geopolitical concentration of refining capacity. Policy responses such as enhanced recycling requirements and material traceability are emerging precisely because volume-driven material systems create cumulative pressure. Hydrogen, arriving later, benefits from this growing awareness — but will also be subject to similar expectations as it scales.

Importantly, this is not an argument against batteries, nor a claim that hydrogen is inherently cleaner. It is an argument against technological absolutism. Different pathways impose different stresses on the global system. Judging them solely on efficiency or point-of-use emissions obscures the deeper trade-offs involved.

From a systems perspective, diversity is not a weakness — it is a form of risk management. A transition dominated by a single material pathway concentrates geological, geopolitical, and social exposure. Technologies that rely on different material profiles — even if they appear less efficient when viewed narrowly — play a critical role in reducing systemic fragility. Hydrogen's role in heavy transport, industry, and long-duration storage is therefore not only about physics or economics, but about balancing material demand across the global economy.

There is also a moral dimension to how material burdens are distributed. Communities asked to accept new infrastructure locally may reasonably question why other communities are expected to absorb the environmental and social consequences of extraction elsewhere. Social licence depends not on whether impacts exist — they always do — but on whether those impacts are acknowledged, shared, and governed transparently.

Material dependence is not evenly distributed across the world, and this has consequences beyond environmental impact. Many of the critical minerals underpinning the energy transition are geographically concentrated, with extraction, processing, or refining dominated by a small number of regions. In several cases, the true bottleneck lies not in mining itself, but in downstream processing and refining capacity, which ultimately determines who controls supply.

This concentration introduces a different form of vulnerability. Supply chains that rely on large volumes of material are more exposed to disruption, price volatility, and geopolitical friction than those that depend on smaller quantities of specialised inputs. As demand accelerates, competition for access intensifies — not only between companies, but between states seeking to secure their own transition pathways.

Hydrogen and battery-electric systems encounter this challenge in different ways. Battery pathways depend on sustained access to high-volume mineral supply chains spread across multiple materials, often sourced far from end users. Hydrogen systems concentrate exposure in narrower material streams, such as platinum-group metals, where supply is limited, but absolute volumes remain comparatively low and recycling rates are high. In both cases, dependency exists; the difference lies in how that dependency manifests and how manageable it may prove to be over time.

These material geographies shape the transition in subtle but important ways. Technologies whose supply chains are perceived as vulnerable, externally controlled, or socially contentious may face resistance regardless of their technical merits. Conversely, pathways that align more closely with existing industrial capabilities, recycling potential, and transparent sourcing are likely to enjoy greater political and social durability.

As the hydrogen economy develops, its own material footprint will attract scrutiny. Platinum-group metals, copper, and specialised alloys will raise questions around sourcing, recycling, and concentration of supply. Anticipating this scrutiny — rather than reacting to it — is essential to maintaining trust.

Ultimately, the question is not whether the energy transition will consume resources. It will. The question is how those demands are distributed, where their impacts are felt, and how honestly trade-offs are discussed. Hydrogen does not eliminate the material challenges of decarbonisation — but it reshapes them in ways that may prove more manageable at global scale.

Recognising this does not weaken the case for hydrogen. It strengthens it — by situating it honestly within the physical, social, and strategic limits of the world it seeks to change.

Who Benefits, Who Pays?

Every energy system distributes costs and benefits unevenly. Some regions host infrastructure. Some supply resources. Some capture value. Others absorb disruption. The energy transition does not change this reality — but it does make it visible in new ways.

Much of the public debate around decarbonisation assumes alignment: that reducing emissions, creating jobs, and delivering economic growth naturally move together. In practice, these outcomes often diverge. The question of social licence ultimately turns on whether those divergences are acknowledged and addressed or quietly ignored.

Hydrogen makes this tension particularly clear. The benefits of hydrogen — reduced emissions, system flexibility, energy security — are frequently framed at national or global scales. The costs, however, are often local and immediate. Water use, land occupation, construction disruption, safety concerns, and visual impact are borne by specific communities. When the beneficiaries of a project feel distant, abstract, or external, acceptance becomes fragile.

This dynamic is not unique to hydrogen, but hydrogen amplifies it because of its role as an enabling technology. Many hydrogen projects exist not to serve local demand, but to support industrial decarbonisation, exports, or system balancing elsewhere. In such cases, communities are asked to host infrastructure whose primary value accrues beyond their borders. Without deliberate mechanisms to share that value, resentment is predictable.

Jobs are often presented as the answer — but employment alone rarely resolves the question. As discussed earlier, hydrogen systems are typically capital-intensive, with relatively modest operational staffing requirements. Construction phases create temporary work; long-term roles are fewer and more specialised. When communities are asked to accept permanent infrastructure in exchange for transient benefit, trust erodes quickly.

The same logic applies internationally. Many of the materials required for the transition are extracted in regions far removed from consumption centres. Water stress, land disruption, and labour impacts are concentrated in specific geographies, while climate benefits are global. If decarbonisation is perceived as relo-

cating environmental burdens rather than reducing them, its moral authority weakens.

This is where the idea of a "just transition" either becomes meaningful — or collapses into rhetoric. Justice in this context is not about perfect equality. It is about proportionality, transparency, and voice. Who decides where infrastructure is built? Who has influence over how risks are managed? Who participates in value creation, and who merely accommodates it?

Projects that answer these questions implicitly often fail. Projects that answer them explicitly are more resilient.

In practice, this means recognising that social licence is not granted once, but maintained over time. Early engagement, local ownership stakes, community benefit mechanisms, and credible long-term commitments all matter more than public consultation exercises conducted after decisions have effectively been made. Fairness is experienced, not declared.

There is also a temporal dimension to this question. Many of the costs of transition are immediate, while many of the benefits are deferred. Asking communities to accept disruption now in exchange for improvements later requires trust in institutions and delivery. Where that trust is weak, even well-designed projects struggle to proceed.

This places a burden on policymakers, developers, and system planners. It is not enough to demonstrate that a hydrogen project reduces emissions or improves system efficiency. It must also make sense in the lived reality of the people affected by it. Efficiency without legitimacy is brittle. Decarbonisation without consent is slow.

None of this implies that difficult decisions can be avoided. The energy transition will require trade-offs, and not all interests can be perfectly aligned. But acknowledging who pays and who benefits — openly and honestly — is the difference between contested progress and durable change.

In the end, hydrogen's success will not be determined solely by electrolyser efficiency, cost curves, or policy incentives. It will be determined by whether it can be embedded into society in a way that feels fair. Where benefits are visible, shared, and sustained, acceptance follows. Where costs are concentrated and value is distant, resistance is rational.

These questions rarely surface in technical language. They emerge instead in everyday conversations, in frustration with projects that feel imposed rather than

understood, and in a desire for continuity in a world changing faster than many people are comfortable with. When I speak to my mother about the energy transition, her concerns are not about optimisation or system design. They are about whether change feels intelligible, respectful, and proportionate to the lives already being lived.

The social and environmental dimensions of the hydrogen economy are therefore not secondary considerations. They are system constraints in their own right. A transition that acknowledges them stands a far better chance of proceeding not just efficiently, but legitimately — and of lasting.

12. The Geopolitics within the Energy Transition

'The trouble at the moment is I read the IRA (*Inflation Reduction* Act) *and I can get a general understanding of what is going on, with the European regulations I need to hire in a team of consultants for 3 months and I am still nowhere clearer as to what we actually need to do*' ----

Statement from the plenum session at the HyFCell in Stuttgart in 2024

A Note on Geopolitics in a Moving World

Of all the chapters in this book, this is the one that has undergone the greatest revision between the first draft and the final manuscript. Not because the underlying principles have changed, but because the geopolitical landscape itself has continued to evolve throughout the writing process. New elections have taken place, conflicts have continued, industrial policies have shifted, and alliances have adapted. Few aspects of the energy transition move as quickly—or change as frequently—as the geopolitical environment that surrounds it.

Writing about geopolitics is always a challenge, but in the energy transition it becomes particularly difficult. Infrastructure evolves over decades, yet the political environment surrounding it can shift almost overnight. Policies change, governments change, and strategic priorities move with them.

The pages that follow therefore represent a snapshot in time. They reflect the geopolitical landscape as it appeared during the writing and final revision of this book, principally across 2025 and early 2026. In a few years some of the specifics may look outdated—and hopefully they will. That would mean the hydrogen economy has continued to evolve, new partnerships have formed, and progress has

been made. In that sense, any future reader who finds parts of this chapter dated will also be witnessing how far the transition has come.

If that proves to be the case, then this chapter will have served a second purpose: not only explaining where we are, but documenting how far we have travelled.

The Genius Principle

Now this principle is something that you have never heard of before, and the reason for that is that it is something that I have had bouncing around my brain for the last while, so I ask for your patience and to bear with me.

The theory goes like this: 'The success of any given nation is directly proportional to its ability to create the geniuses it requires to drive industry forward.'

It is worth noting that 'genius' here is not necessarily a lone prodigy, but can also refer to exceptional innovators, problem-solvers, and highly capable teams whose combined efforts drive meaningful breakthroughs. Exceptional individuals are like seeds. On their own, they may have potential, but without fertile soil, sunlight, and a supportive network, that potential cannot be realized. Renaissance Italy illustrates this perfectly: the city-states did not merely have talented artists and thinkers—they cultivated an environment of patrons, schools, and collaboration that allowed these individuals to reshape art, science, and commerce. Talent alone was not enough; it was the ecosystem that amplified it.

Clever people exist everywhere, but certain conditions must be in place for this principle to work in practice. Even the most talented individuals cannot reach their potential in isolation. Networks, collaboration, and access to knowledge amplify individual brilliance, allowing talent to translate into transformative innovation. For a genius to be created, you require a stable environment, good health care, a nurturing system for recognising talent and being able to give them the opportunity to develop the skills that they require to be able to fulfil their potential. Top institutions not only provide knowledge but also foster exposure to diverse ideas, mentorship, and collaborative opportunities—elements that are crucial for turning potential into tangible impact. There must be the availability of world class higher learning institutes to be able to put the final touches to these people and introduce them to all of the potential options available to them so that they can find 'their' field, the area that interests them.

There then needs to be the companies available to them to offer them the appropriate salary and advancement possibilities. Also of great importance is the ability of any country to make venture capital available to enable geniuses to follow their

dreams. Without an ecosystem that supports risk-taking, experimentation, and the scaling of ideas, even the most brilliant minds may be unable to make a meaningful contribution to industry or society.

Once you get that, you get your geniuses, and they will drive your economy forward.

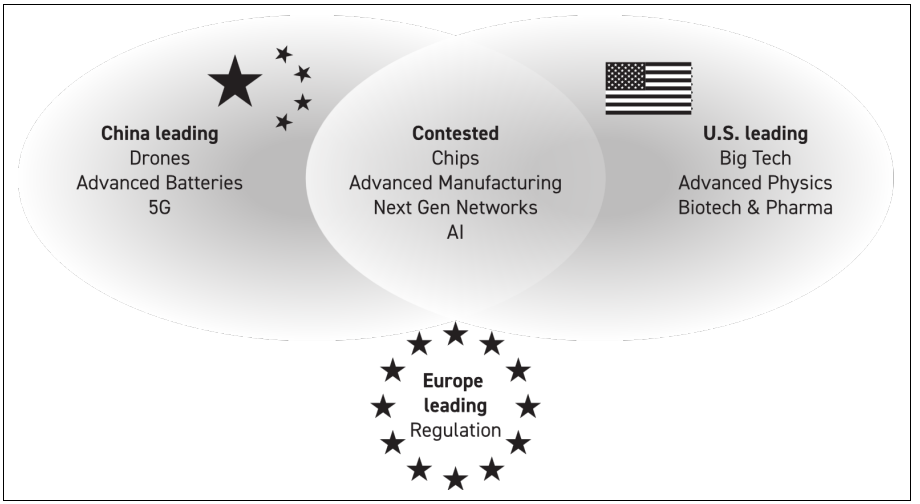
While population size can increase the absolute number of exceptional individuals, it is the combination of talent, resources, and the ability to collaborate effectively that ultimately determines whether a nation's potential is realised. A recent visit to China brought this idea into sharper focus for me. China has a long history of invention and scientific achievement, although for much of the past century the conditions needed to translate that potential into sustained industrial leadership were not always present. Whatever the reasons for that interruption, many of the conditions described above are now firmly in place: strong educational institutions, industrial capability, investment, infrastructure, and increasingly sophisticated innovation networks.

I was recently at a Hydrogen Hub event hosted by Extreme H racing and a comment made there was, 'no matter what the rest of the world does China will win due to weight of numbers'. That is a point to think about now that China has the conditions laid out above, it now has increasingly more genius minds to call upon to drive its industries forward. But creating talent is only one piece of the puzzle; retaining and attracting it is equally critical.

Now what about India you might say, they also have a large population and the higher learning etc. and you would be right, but then comes into play a caveat, it isn't only the ability to create geniuses, but the ability to keep them, or rather the ability to attract them to your own country, and here my belief is that India suffers due to a 'brain drain' of capable genius' leaving to go to primarily the United States or Europe.

This is probably America's biggest selling point, its ability to attract the best talent in the world to emigrate to them and drive their industries forward. China's biggest success currently is the volume of geniuses it can create and keep hold of, whilst Europe may have some to the best conditions, but isn't always able to keep hold of its talent when compared to what American venture capital and opportunities has to offer them. It is telling that if you have an idea in the US, you can get five times more venture capital to bring it to life than you would traditionally get in Europe.

In short, national success depends not only on the creation of exceptional individuals but on the environment, institutions, and networks that allow their brilliance to manifest at scale.



The European Union: The Gordian Knot of Regulation

Europe… where to even begin?

The EU, for me, has been one of the greatest political experiments the world has ever undertaken. There are many who argue that it has not met its goals or delivered fully on its promise, but it has undeniably given the continent its longest uninterrupted period of peace. The institutions built after 1945 helped drive prosperity, stability, and cooperation in a region that had spent centuries oscillating between conflict and uneasy rivalry. I say this not just as an observer of European history, but as someone who has personally benefited from the European project. Through the Erasmus programme I lived, studied, and moved freely across borders—and it was in that period of exchange that I met my wife. That kind of cross-pollination of cultures, languages, ideas, and relationships is not a minor feature of the EU; it is one of its greatest sources of strength. It has created a generation for whom Europe is not an abstraction, but a lived reality, a shared space. In ways that statistics never capture, this human integration has made the continent more resilient, more cohesive, and ultimately stronger.

To understand why Europe approaches hydrogen the way it does—slowly, cautiously, through layers of regulation—you need to understand the long arc of how Europe rebuilt itself after the war. Ian Kershaw's 'Roller-Coaster', his panoramic history of Europe from 1950 to 2017, captures this transformation brilliantly. He shows that Europe's success was not born from autonomy but from interdependence. Energy, steel, coal, trade—everything was pooled, shared, negotiated. Europe grew strong by weaving its nations together, not by standing apart.

The legacy of that system is still with us.

Europe's prosperity was built on a delicate web of shared markets, cross-border supply chains, and political compromises. But the price of interdependence is complexity. Integration happened economically, while political power stubbornly remained national. Kershaw describes this as the continent's defining tension: integration without unity.

Hydrogen is now colliding with this historical fault line.

From Germany's engineering sector to France's nuclear programme to the Netherlands' port infrastructure, Europe possesses vast industrial genius. Its universities, research institutions, and innovation clusters produce world-class talent. Yet converting that intellectual wealth into industrial dominance in hydrogen has been painfully slow. Projects emerge—hydrogen valleys, electrolyser factories, cross-border pipelines—but progress is uneven, fragmented, and often paralysed by the regulatory labyrinth that results when 27 nations try to innovate at once.

This is not a new failing; it is a historical pattern.

Europe learned after the war to advance through consensus rather than decree, through negotiation rather than direction. That made sense when rebuilding after 1945, during the Cold War, and through the integration of Eastern Europe after 1989. Nowadays however, as the energy transition becomes a geopolitical race, the very structures that made Europe resilient now make it slow.

Kershaw also emphasises how Europe has always been shaped by global forces—first U.S. security guarantees, then Soviet pressure, then the rise of globalisation and China. The energy transition fits squarely into this tradition. Europe watched solar, batteries, and clean-tech manufacturing slip abroad. It watched the Inflation Reduction Act unleash a brief period of extraordinary American clarity and speed—a rare moment when Washington moved with industrial purpose. But it is now watching that momentum stall, as political reversals and strategic drift begin to unpick what had looked like a decisive U.S. lead. Meanwhile, China con-

tinues scaling hydrogen and electrolysers with a level of coordination Brussels cannot hope to replicate.

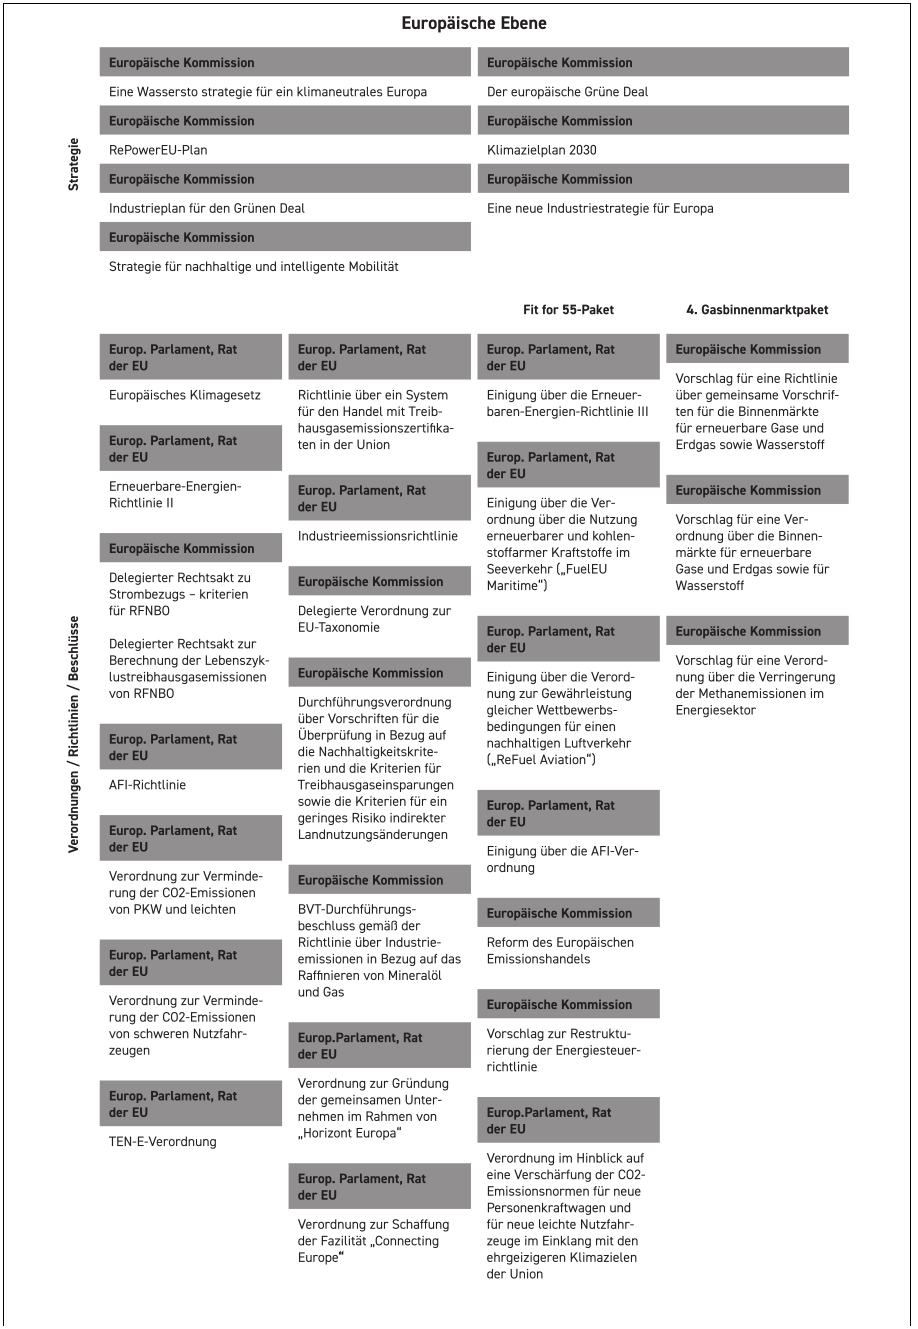
Hydrogen is not just an energy vector for Europe.

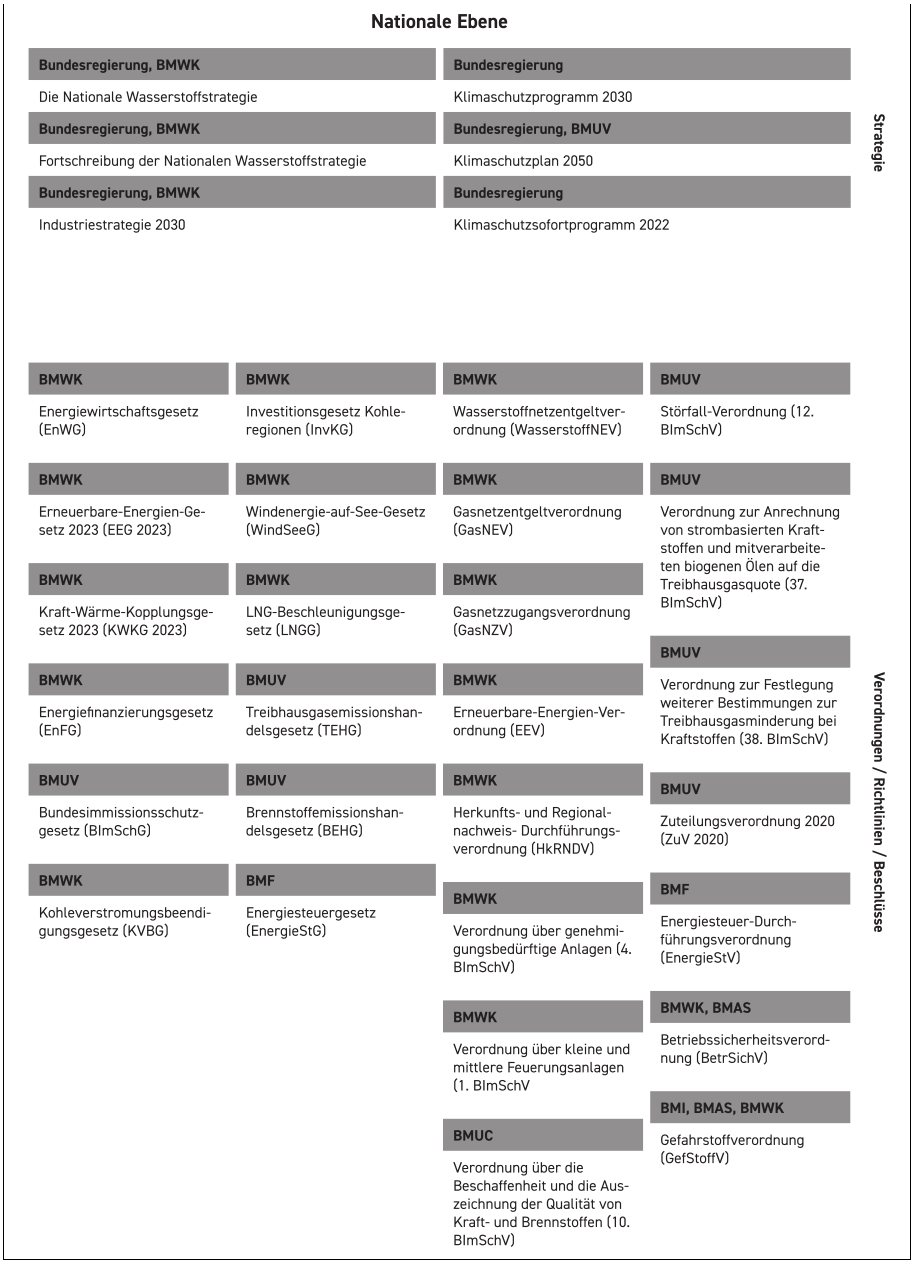
It is a last, urgent attempt to avoid repeating these losses.

Europe's genius is abundant but diffused. It has talent, expertise, and industrial heritage—but the ability to align these strengths into a cohesive strategic force is hampered by the historical DNA Kershaw describes; institutions built for stability, not speed; interdependence without sovereignty; ambition without executive authority.

And this is precisely why the image at the start of this chapter hits so hard. It is only a sketch—a provocation more than an analysis—but it reflects a truth Europeans recognise instinctively. While the United States and China dominate the commanding heights of technology and industrial capacity, Europe's comparative strength has settled into regulation. Not invention. Not scale. Not strategic manufacturing. Regulation. It is the legacy of a system designed to manage interdependence rather than to win industrial races. The picture is not destiny, but it is a warning: unless Europe can turn its genius into coordinated power, it risks becoming the world's rule-maker rather than a world-shaper.

Let us have a short look at Germany, I was recently at the General Assembly of the Hydrogen Alliance Bavaria, where an overview was shown that shocked me.


Europäische Ebene
Strategie
Verordnungen / Richtlinien / Beschlüsse
Europäische Kommission
Eine Wassersto strategie für ein klimaneutrales Europa
Europäische Kommission
RePowerEU-Plan
Europäische Kommission
Industrieplan für den Grünen Deal
Europäische Kommission
Strategie für nachhaltige und intelligente Mobilität
Europäische Kommission
Der europäische Grüne Deal
Europäische Kommission
Klimazielplan 2030
Europäische Kommission
Eine neue Industriestrategie für Europa
Fit for 55-Paket
4. Gasbinnenmarktpaket
Europ. Parlament, Rat der EU
Europäisches Klimagesetz
Europ. Parlament, Rat der EU
Erneuerbare-Energien-Richtlinie II
Europäische Kommission
Delegierter Rechtsakt zu Strombezugs – kriterien für RFNBO
Delegierter Rechtsakt zur Berechnung der Lebenszyklustreibhausgasemissionen von RFNBO
Europ. Parlament, Rat der EU
AFI-Richtlinie
Europ. Parlament, Rat der EU
Verordnung zur Verminderung der CO2-Emissionen von PKW und leichten
Europ. Parlament, Rat der EU
Verordnung zur Verminderung der CO2-Emissionen von schweren Nutzfahrzeugen
Europ. Parlament, Rat der EU
TEN-E-Verordnung
Europ. Parlament, Rat der EU
Richtlinie über ein System für den Handel mit Treibhausgasemissionszertifikaten in der Union
Europ. Parlament, Rat der EU
Industrieemissionsrichtlinie
Europäische Kommission
Delegierte Verordnung zur EU-Taxonomie
Europäische Kommission
Durchführungsverordnung über Vorschriften für die Überprüfung in Bezug auf die Nachhaltigkeitskriterien und die Kriterien für Treibhausgaseinsparungen sowie die Kriterien für ein geringes Risiko indirekter Landnutzungsänderungen
Europäische Kommission
BVT-Durchführungsbeschluss gemäß der Richtlinie über Industrieemissionen in Bezug auf das Raffinieren von Mineralöl und Gas
Europ.Parlament, Rat der EU
Verordnung zur Gründung der gemeinsamen Unternehmen im Rahmen von „Horizont Europa“
Europ. Parlament, Rat der EU
Verordnung zur Schaffung der Fazilität „Connecting Europe“
Europ. Parlament, Rat der EU
Einigung über die Erneuerbaren-Energien-Richtlinie III
Europ. Parlament, Rat der EU
Einigung über die Verordnung über die Nutzung erneuerbarer und kohlenstoffarmer Kraftstoffe im Seeverkehr („FuelEU Maritime“)
Europ. Parlament, Rat der EU
Einigung über die Verordnung zur Gewährleistung gleicher Wettbewerbsbedingungen für einen nachhaltigen Luftverkehr („ReFuel Aviation“)
Europ. Parlament, Rat der EU
Einigung über die AFI-Verordnung
Europäische Kommission
Reform des Europäischen Emissionshandels
Europäische Kommission
Vorschlag zur Restrukturierung der Energiesteuerrichtlinie
Europ.Parlament, Rat der EU
Verordnung im Hinblick auf eine Verschärfung der CO2-Emissionsnormen für neue Personenkraftwagen und für neue leichte Nutzfahrzeuge im Einklang mit den ehrgeizigeren Klimazielen der Union
Europäische Kommission
Vorschlag für eine Richtlinie über gemeinsame Vorschriften für die Binnenmärkte für erneuerbare Gase und Erdgas sowie Wasserstoff
Europäische Kommission
Vorschlag für eine Verordnung über die Binnenmärkte für erneuerbare Gase und Erdgas sowie für Wasserstoff
Europäische Kommission
Vorschlag für eine Verordnung über die Verringerung der Methanemissionen im Energiesektor



Adapted from NOW GmbH; *illustration redesigned for this book.*

What you can see here came from NOW GmbH, a German state-owned agency driving the country's shift to sustainable energy and mobility. It manages funding programmes, coordinates industry and research, and supports the rollout of technologies like hydrogen, fuel cells, and electric mobility. More than a think tank, it is a practical "do tank" turning policy goals into real-world projects.

What it shows is the amount of regulation you need to go through to bring a hydrogen project to life in the German market: from the EU, 7 Strategy guidelines and 24 regulatory documents, from the German government, 6 strategy guidelines and 28 regulatory documents. It can make you understand the frustrations stated in the quote at the beginning of the chapter, 13 strategic guidelines and over 50 pieces of documentation that need to be both understood and that is before requesting any potential funding to help make the economics of a project work. It simply is too much.

I think that the point has been made about the hurdles there, but maybe lets now try to understand what the EU is trying to do to push hydrogen forward and that means understanding both the Renewable energy Directive III (RED III) and the European Hydrogen Bank (EHB).

RED III and the Push for Renewable Hydrogen

The European Union's Renewable Energy Directive III (RED III) strengthened the legal framework for renewable energy and explicitly expanded the role of renewable hydrogen — and related fuels, known as RFNBOs (Renewable Fuels of Non-Biological Origin) — beyond transport.

RED III sets binding quotas for renewable hydrogen use: by 2030, at least 42% of hydrogen in industry must be renewable, rising to 60% by 2035. For transport, the directive also establishes sub-targets, requiring a share of energy supplied to come from advanced biofuels and RFNBOs. Member states have some flexibility in how they implement these requirements, but compliance is non-negotiable.

To prevent "greenwashing," RED III relies on detailed delegated acts. These acts define the rules for renewable electricity used in hydrogen production, including requirements for additionality, temporal and geographic correlation, and a life-cycle greenhouse-gas savings threshold compared to fossil alternatives. In practice, any hydrogen project seeking public support or market access must plan for green electricity procurement, certification, documentation, and early investment in renewable-hydrogen infrastructure.

These requirements do more than complicate electricity procurement and certification; they can influence how the electrolyser itself is operated. Temporal correlation rules, for example, link qualifying hydrogen production increasingly closely to the availability of qualifying renewable electricity. A developer may therefore be unable simply to run an electrolyser whenever electricity is cheapest or whenever hydrogen demand is highest, while still treating all of the resulting hydrogen as renewable under the RFNBO framework. Regulation consequently becomes another variable in the operating philosophy, alongside electricity price, utilisation, efficiency, storage, and offtake demand.

Even before funding comes into play, RED III adds a layer of EU-level compliance, certification, and quotas to the already complex regulatory landscape — as your NOW GmbH diagram vividly illustrates.

The European Hydrogen Bank: Turning Quotas into Reality

Since renewable hydrogen is more expensive than fossil-based alternatives, especially in early phases, the EU created the European Hydrogen Bank (EHB) — not a physical bank, but a financing instrument designed to close the investment gap and make green hydrogen economically viable.

The EHB is built on four "pillars":

- Domestic financing, using an auction system where producers bid for a premium (€/kg of hydrogen produced). Winning bidders receive subsidies for up to 10 years, bridging the cost gap between renewable and conventional hydrogen.
- International financing, enabling imports of renewable hydrogen from non-EU countries under defined sustainability standards.
- Transparency and coordination, ensuring clarity and efficiency in project support.
- Coordination with other EU and national instruments, blending grants, loans, and state aid to maximise leverage.

The first auction, launched in late 2023, had a budget of around €800 million, supporting several renewable-hydrogen projects. The second call increased the budget to roughly €1–1.2 billion, with a third call expected by the end of 2025.

Together, RED III and the EHB are designed to create a market for renewable hydrogen: not through direct industrial direction, but by incentivising economically viable production, enabling long-term contracts, reducing the cost gap, and signalling to investors that hydrogen is a strategic pillar of decarbonisation. The

EU steers markets, sets standards, and structures incentives — but does not build factories itself.

CBAM: Protecting Industrial Competitiveness

Hydrogen policy in Europe does not exist in isolation. One of the underlying concerns driving European decarbonisation policy is the risk of carbon leakage—the relocation of energy-intensive industries to jurisdictions with weaker environmental requirements and lower carbon costs.

This concern gave rise to the Carbon Border Adjustment Mechanism (CBAM), one of the most significant industrial policy tools introduced by the European Union in recent decades.

In simple terms, CBAM seeks to ensure that imported products face a carbon cost comparable to that paid by European producers operating under the EU Emissions Trading System (ETS). During its initial phase, the mechanism covers sectors such as steel, cement, fertilisers, aluminium, electricity, and hydrogen—all industries that are both carbon-intensive and strategically important to Europe's industrial base.

The logic is straightforward. If a European steel producer invests in renewable electricity, low-carbon hydrogen, carbon capture, or other decarbonisation technologies, it should not be placed at a competitive disadvantage against imported steel produced using higher-emission processes. Without some form of border adjustment, Europe risks exporting emissions rather than reducing them, while simultaneously undermining its own industrial competitiveness.

For hydrogen, CBAM is particularly significant because it helps create confidence that investments in low-carbon production pathways will not be immediately undercut by imports from regions operating under less stringent environmental frameworks. In this sense, RED III creates demand, the European Hydrogen Bank helps support supply, and CBAM seeks to preserve a level competitive playing field.

Yet CBAM also highlights one of the broader dilemmas facing Europe. While it can encourage decarbonisation and protect domestic industry, it is viewed by some trading partners as a form of protectionism. Supporters see it as a necessary correction that reflects the true environmental cost of production. Critics argue that it risks creating new trade barriers and could provoke retaliatory measures.

Whether viewed as climate policy, industrial policy, or trade policy, CBAM illustrates a broader reality: the hydrogen transition is no longer simply about technology. It is increasingly about competitiveness, economic sovereignty, and the rules that will govern global markets in a decarbonising world.

Challenges: Where Regulation Meets Urgency

RED III and the EHB also highlight Europe's regulatory complexities. Some challenges are structural, others practical:

- National transposition: Binding EU quotas only matter if Member States implement RED III in national law — and do so on time. As of 2025, many countries are behind schedule, delaying investments and discouraging long-term contracts.
- Compliance complexity: Meeting RFNBO requirements (additionality, temporal/geographic correlation, lifecycle emissions) is technically demanding, increasing administrative costs and barriers for smaller players or early-stage projects.
- Cost gap persists: Even with EHB support, renewable hydrogen remains more expensive than conventional hydrogen or fossil fuels. Long-term competitiveness depends on scale, infrastructure, and consistent demand — all of which remain uncertain.
- Coordination and harmonization: Across 27 member states, regulatory harmonisation, certification, monitoring, market transparency, and infrastructure rollout must all align, making hydrogen one of the most complex energy vectors to scale in the EU.

Europe's regulatory DNA creates a strong and principled framework — but one that is inherently slow, fragmented, and risk-averse. For hydrogen to become transformative rather than symbolic, regulation and incentives must be matched by political will, industrial coordination, and structural clarity. In a 27-member-state union, that alignment is far easier to design than to execute.

The Reality of Transposition

RED III was meant to be fully implemented across the European Union by May 2025. Yet as of early 2026, the reality remains uneven. Denmark was the only member state to notify full transposition by the legal deadline. In the months that followed, the European Commission initiated infringement procedures against the remaining member states for failing to notify full implementation of key pro-

visions — and those procedures continue into 2026, signalling that delayed transposition remains an active enforcement priority rather than a closed chapter.

Since then, progress has continued — but largely in partial or staged form. Several countries — including Germany, France, and other larger economies — have adopted implementing legislation covering selected elements of RED III, such as transport-fuel mandates, permitting acceleration, renewable-electricity integration, and preparatory hydrogen compliance frameworks. However, full transposition — including comprehensive incorporation of RFNBO quotas, certification systems, and hydrogen-specific compliance mechanisms — remains incomplete in most member states.

Germany provides a useful example of national divergence. It has advanced legislation implementing core RED III elements, particularly around renewable permitting and transport-sector mandates, but has chosen not to impose company-specific RFNBO quotas. Instead, it relies on subsidy mechanisms, climate-protection contracts, and funding instruments to reach the directive's targets. Other large member states are pursuing similarly differentiated approaches, reflecting domestic industrial-policy priorities and administrative constraints.

This patchwork illustrates the structural tension at the heart of Europe's regulatory model. Ambitious pan-European frameworks collide with national political calendars, administrative capacity constraints, and differing industrial strategies. The Commission's continued enforcement actions demonstrate seriousness of intent, yet they also underscore how difficult synchronised implementation is across sovereign jurisdictions.

For hydrogen developers, the consequences are tangible. High production costs, technical compliance requirements, and rigorous certification already make early-stage projects capital-intensive. Add uncertainty around national implementation timelines, and the path toward a harmonised European hydrogen market becomes even more complex. RED III provides the scaffolding of a future hydrogen economy — but as of early 2026, much of the building work remains ongoing.

The deeper question, however, is not whether RED III will eventually be transposed — it will be — but whether Europe can move quickly enough to convert regulatory ambition into industrial momentum. In a global race shaped increasingly by Chinese scale and American volatility, delay carries strategic cost. Europe's genius is not in question; its capacity for coordinated acceleration is. The test of RED III is therefore not legal completeness, but whether it allows

European talent and industry to scale before others define the rules of the hydrogen economy.

RED III Transposition Status — EU Member States (Snapshot: Early 2026)

Member State / Group	Status (Early 2026)	Notes
Denmark (DK)	✓ Fully transposed by deadline	The only member state to notify full transposition by the May 2025 legal deadline.
Majority of EU Member States (incl. DE, FR, IT, ES, NL, PL)	⚠ Partial or staged transposition	Key elements implemented (e.g., transport mandates, renewable permitting, hydrogen framework preparation), but full RFNBO quota and certification integration remains incomplete.
Subset of Member States	✗ Not yet notified full transposition — infringement procedures ongoing	European Commission infringement actions initiated in July 2025; enforcement continues into 2026.
Germany (DE) — Example of Divergent Implementation	⚠ Advanced legal implementation, alternative compliance model	Has implemented core RED III elements but relies on subsidies and climate-protection contracts rather than company-specific RFNBO quotas.

Legend:

✓ = fully transposed
⚠ = partial / staged / divergent implementation
✗ = not fully transposed; infringement procedures active

Note to reader: This table provides a high-level structural snapshot as of early 2026. Detailed transposition of RFNBO quotas, hydrogen certification systems, transport-fuel mandates, and related sub-elements varies by Member State and continues to evolve.

A Future for European Hydrogen: Talent, Tenacity, and Opportunity

Despite the regulatory maze, fragmented transposition, and slow pace of implementation, there is still much to be optimistic about. Europe may not move as fast as the United States or China, but it is far from stagnant. Across the continent, the foundations for a vibrant, renewable-hydrogen economy are being laid — and the most important ingredient is already here: human genius.

From universities to research institutes, innovation clusters to industrial hubs, Europe possesses a deep reservoir of talent. Engineers, chemists, policy designers, and strategists are already working on the cutting edge of electrolysers, fuel cells, and hydrogen transport. This is not abstract potential; it is concrete capab-

ility, cultivated over decades, now ready to meet the urgent challenge of decarbonisation.

Programs like the European Hydrogen Bank, RED III quotas, and national support mechanisms are imperfect and uneven, but they provide direction, incentives, and structure. They signal to investors, entrepreneurs, and policymakers that renewable hydrogen is a serious, long-term priority. Every auction, every support programme, and every partially implemented law is a stepping stone — evidence that Europe is learning to translate genius into coordinated action, even if imperfectly.

History shows that Europe excels when it brings its collective talent together. After the Second World War, the pooling of resources, ideas, and expertise created the world's most integrated and resilient continent. The same principle applies to hydrogen today: the genius is there, the industrial and intellectual foundations exist, and the regulatory scaffolding, though complex, is being built. What Europe needs now is patience, persistence, and continued alignment — turning scattered brilliance into systemic, strategic capability.

In other words, while the path may be slow, frustrating, and at times convoluted, it is not without hope. Europe's strength has always been its people: a generation of innovators, thinkers, and problem-solvers capable of turning complex challenges into transformative opportunities. If regulatory frameworks, financial incentives, and industrial coordination continue to evolve — and if European genius is allowed to meet strategic urgency — the continent can yet become a world-shaper, not just a rule-maker.

The hydrogen transition is not merely a test of technology or policy; it is a test of Europe's ability to harness its greatest resource: its talent. It is also increasingly a question of energy security. For a continent that imports much of its energy, the transition is not only about decarbonisation but about resilience—reducing exposure to volatile global fuel markets and geopolitical disruption. In that sense, hydrogen infrastructure is not simply climate policy; it is part of Europe's long-term strategy for energy sovereignty. And in that talent lies the promise of a sustainable, competitive, and resilient future.

Europe may be slow, regulated, and complex, but with its unmatched talent and industrial genius, it still holds the power to cut through its own Gordian Knot and turn promise into global leadership in hydrogen.

The United Kingdom: Regulation Done Right

Hailing originally from the UK but living on the continent for over 20 years, I have had my fair share of frustrations with UK governmental policy. So, it is no small thing for me to highlight an area where the UK has truly done well. Looking back across the channel, and acknowledging that not everything is perfect, I can nonetheless see signs of progress. From the outside, the UK's approach to hydrogen appears increasingly coherent. While progress has not always been as rapid as many in industry would like, the combination of targeted regulation, innovative funding mechanisms and a growing willingness to support domestic manufacturing suggests a policy framework that continues to mature.

Historically a hub of scientific and industrial innovation—from the Industrial Revolution to cutting-edge technology today—the UK demonstrates how smart regulation can accelerate industrial progress. Through a combination of clear policy frameworks, targeted incentives, and streamlined approval processes, the UK enables projects to move forward efficiently. Hydrogen clusters in Teesside, Scotland, and the Humber show that when regulation is designed to reduce friction rather than create bottlenecks, industrial ambition can be realised quickly and effectively.

The UK's experience highlights two key lessons. First, regulation can be a powerful accelerator when structured with clarity and purpose. Second, aligning incentives with strategic industries allows government policy to translate into real infrastructure and industrial growth. Together, these elements allow the UK to punch above its weight, showing that thoughtful regulation can achieve outsized results.

HS2 — When Smart Policy Drives Innovation

One of the clearest examples of the UK getting regulation right is the high-speed rail project, HS2. From the outset, this was never simply a transport project — it was conceived as a national platform for sustainable infrastructure. The government set out clear, measurable expectations: construction must reduce carbon emissions, adopt renewable energy wherever possible, and minimise the environmental impact of operations. Unlike many projects that falter under ambiguous rules or excessive bureaucracy, HS2 provided a precise regulatory framework that encouraged innovation while keeping progress accountable.

This clarity created opportunities for new entrants and green innovators. Among them was GeoPura, a company producing hydrogen-powered generators capable of replacing diesel units on construction sites. While I explore GeoPura's story in

depth later in the book, it is worth noting here how the company's emergence is directly linked to HS2's regulatory environment. By mandating "green energy" use on site, the project offered a clear market signal, allowing innovative technologies to move from concept to deployment. Trials of GeoPura's hydrogen units demonstrated that they could reliably supply power across construction sites, eliminating significant CO_2 and other emissions, while reducing noise and particulates compared to diesel generators.

HS2's regulatory approach extended beyond power sources. Low-carbon construction materials, hybrid machinery, and other sustainability measures were embedded into contracts and enforced through oversight programs, such as the industry "Plant Charter." Contractors who implemented these requirements successfully and were recognised with awards for their sustainability performance. These measures ensured that environmental goals translated into concrete actions across every stage of the project, showing that regulation can be both ambitious and operationally effective.

The project has not been without complexity. Large-scale infrastructure inevitably involves trade-offs — from land use and habitat disruption to embodied emissions in concrete and steel. Yet, the HS2 framework demonstrates how regulation can steer an industry toward innovation and lower-carbon solutions, even in projects of enormous scale and complexity.

The lesson is clear: regulation, when designed with precision, clarity, and measurable targets, can accelerate industrial innovation and deployment. HS2 demonstrates that policy is not merely a set of rules to comply with — it can be a lever for change, creating opportunities for new technologies, shaping behaviour across an entire sector, and translating ambition into tangible results. For the UK, this approach has enabled projects like HS2 to move forward efficiently, while simultaneously fostering the growth of companies like GeoPura. As noted, I will explore GeoPura in much greater depth later in this book, but their story here illustrates the principle: smart regulation can turn policy goals into real-world outcomes.

The same philosophy is now visible beyond infrastructure. Through the Hydrogen Allocation Rounds (HAR), the UK has developed one of Europe's most thoughtful hydrogen support mechanisms. Rather than relying solely on upfront capital grants, HAR provides long-term revenue certainty through Contracts for Difference-style support, with successful projects able to receive revenue support over a 15-year period, reducing commercial risk and giving developers the confidence to reach final investment decisions. Companies such as GeoPura, whose hydrogen-powered generators are now being deployed commercially, are

examples of the wider ecosystem that benefits from this approach. The mechanism recognises an important reality: emerging industries rarely fail because technology is unavailable—they fail because investors lack confidence that future revenues will justify today's capital investment.

There is, however, one area where the framework could be strengthened further. While HAR successfully reduces financial uncertainty, it places relatively little pressure on successful applicants to progress projects within clearly defined timeframes. A stronger emphasis on delivery milestones or time-limited commitments could help ensure that awarded projects move more rapidly from approval to construction, accelerating both industrial learning and market development without undermining the stability that makes the scheme attractive in the first place.

Teesside and the East Coast Cluster – From Carbon-Intensive Past to Low-Carbon Future

In Teesside, an industrial region long associated with heavy industry and high emissions, the UK is deploying regulation not as a brake, but as an engine for transformation. Under the government-backed East Coast Cluster (ECC), a carefully structured regulatory and policy framework is enabling hydrogen production, carbon capture and storage (CCUS), and low-carbon power generation at a scale that could reshape both the region and the nation's industrial emissions profile.

The key projects speak to ambition: H2Teesside promises to become one of the UK's largest hydrogen plants, while Net Zero Teesside Power points to a gas-fired power station with integrated carbon capture — producing dispatchable low-carbon power for hundreds of thousands of homes. Crucially, both tie into a newly licensed CO_2 transport-and-storage system under the Northern Endurance Partnership (NEP), backed by government regulation and long-term permits.

This regulatory certainty — from CCUS licensing to hydrogen and power permitting — transforms what used to be high-risk, carbon-heavy industry into a viable path toward net-zero. It shows that with well-designed regulation, governments can steer entire industrial clusters toward sustainability, while offering investors the confidence they need to commit major capital.

For Teesside, the promise is not only environmental. The ECC scenario anticipates thousands of new jobs, supply-chain renewal, and regional regeneration — a modern "just transition." This example proves that regulation, far from constrain-

ing industry, can become the backbone of industrial renewal, aligning national climate goals with economic revitalisation at a regional scale.

Regulation as an Enabler: Harnessing Ambition and Innovation

What the HS2 and Teesside examples demonstrate is that, when regulation is smartly designed, it does not simply constrain—it enables. Clear, measurable, and enforceable rules provide certainty to innovators, investors, and industrial actors. They allow ambition to meet opportunity, and ideas to become tangible solutions. Companies like GeoPura thrive because the regulatory environment gives them a stage: a framework where hydrogen-powered technologies are not an abstract concept but a real-world requirement, backed by policy and enforceable standards. It is precisely this clarity that transforms a promising idea into deployed technology, helping innovators move from lab to site.

The UK's approach is multidimensional. HS2 highlights green hydrogen in action, showing how regulation can directly shape deployment in infrastructure. Meanwhile, Teesside underscores how blue hydrogen, coupled with CCUS, is being brought to scale under government-backed cluster policy. The combination of regulatory certainty, financial incentives, and long-term industrial planning ensures that hydrogen — in all its forms — can become a practical, scalable solution rather than a theoretical aspiration.

In both cases, the lesson is the same: regulation can amplify genius. It does not create talent, but it creates the conditions in which talent and ambition can flourish. Where bureaucratic uncertainty or fragmented rules might stifle innovation, the UK has demonstrated that clarity, consistency, and purpose can unlock it.

By fostering an ecosystem where green innovators like GeoPura can grow, and where large-scale industrial clusters can transition to low-carbon energy through hydrogen and CCUS, the UK is showing that regulation can be a powerful lever — one that turns vision into action, potential into performance, and ambition into concrete industrial achievement.

More recently, the UK's approach has continued to evolve. Having established regulatory frameworks that encourage deployment, government policy is now becoming more directly involved in building domestic industrial capability. The creation of Great British Energy and its investment in companies, such as ITM Power, signals a subtle but important shift. Rather than acting solely as a regulator and market designer, the state is increasingly positioning itself as a strategic investor, seeking to preserve manufacturing capability, strengthen domestic sup-

ply chains and retain technological expertise within the UK. This reflects a wider geopolitical trend seen across many developed economies, where energy security, industrial competitiveness and technological sovereignty are encouraging governments to play a more active role in shaping strategic industries.

In short, the United Kingdom reminds us that when regulation is precise, targeted and forward-looking, it does more than enforce standards: it creates confidence. Confidence encourages entrepreneurs to innovate, investors to commit capital and industries to modernise. Equally, as the hydrogen economy matures, regulation alone is no longer sufficient. Long-term revenue mechanisms, such as the Hydrogen Allocation Rounds, provide developers with the certainty needed to reach investment decisions, while more recent initiatives, including, Great British Energy demonstrate a willingness to support strategically important manufacturing capability directly. There remains scope for improvement. Although HAR has proven to be one of Europe's most effective funding mechanisms for reducing investment risk, stronger delivery milestones could encourage successful projects to progress more rapidly from approval to construction. Nevertheless, the UK's willingness to refine its approach as the market evolves illustrates an important lesson: successful industrial strategy is not static. It adapts alongside technology, markets and geopolitical priorities. In doing so, the UK provides a valuable example of how regulation, investment and industrial policy can work together to accelerate the energy transition.

China: The Great Enigma

It was while thinking about China that the idea for my “Genius Principle” truly crystallised. China is, unfortunately, a place I haven’t managed to visit as often as I would like. My first trip was in 2008, followed by two more visits in 2025.

Before that first visit, the world was consumed with talk of the BRIC nations—Brazil, Russia, India, and China—as the next engines of global growth. Having travelled to India multiple times by then, I expected something similar. But nothing could have been further from the truth.

In 2008, I encountered a country that felt—at least in my own interpretation—like an adolescent testing its strength, trying to understand its capabilities, and preparing to shoulder more responsibility on the world stage. Beijing was gearing up for the Olympics and China was ready to declare its “coming of age.”

The China I returned to in 2025 was unrecognisable from that earlier version. This was a nation firmly into early adulthood—confident, disciplined, and keenly aware of what it could deliver. It had a clarity of purpose that few countries possess, and

like a top athlete, it carried a calm swagger born of inner belief: not only would it compete, but it fully expected to take gold.

And yet, this is where China becomes the great enigma in the context of the Genius Principle. Its strength lies in extraordinary organisational capacity and an unmatched ability to execute at scale—qualities that have produced massive electrolyser factories, sprawling fuel-cell deployment programs, and industrial clusters that appear almost overnight. But genius, in the disruptive, paradigm-shifting sense, requires something else: creative freedom, independent risk-taking, and the space for eccentricity to flourish. China's model raises the central question of this section: can a system built on discipline and coordination also nurture the kind of free-form innovation that leads to true breakthroughs?

I think that question has already been answered with a resounding, yes!

In 'China Shakes the World', James Kynge illustrates China's extraordinary ability to mobilise entire industries at astonishing speed. He shows how long-term governmental planning enables China to exert a unique "weight of scale" — the power to create, shift, or even absorb global markets simply by choosing to prioritise a sector.

The world often finds China both awe-inspiring and unsettling — not because of unpredictable flashes of creativity, but because of its relentless capacity to organise, coordinate, and out-scale everyone else. When I returned in 2025, I witnessed this even more vividly: a nation that turns ambition into reality with industrial precision and at a pace that borders on the surreal.

This is something that we in the Western world have seen time and time again. Over the past 20 years, innovations in wind power, photovoltaics, electric vehicles, and the vast ecosystem of associated materials and components have often begun in Western labs and companies — yet the industries themselves have ended up dominated by China. Not because China was first, but because its strengths lie in refinement, mass-production, cost leadership, and overwhelming industrial scale rather than first-mover invention.

What James Kynge described as China's "industrial tidal wave" is now clearly visible in hydrogen: electrolyser gigafactories appearing almost overnight, state-backed fuel-cell clusters expanding across provinces, and supply chains assembled at a scale no other nation is structurally capable of matching.

It is also wrong to assume that China is built solely on "old technology." This was another misconception I held for far too long. I once believed that the West — with its conditions for creative freedom, disruptive thinking, and the 'genius prin-

ciple' — would naturally maintain a lead in the race to commercialise cutting-edge technologies.

Then, earlier this year, I joined a trade delegation to North America. We visited a promising hydrogen company operating in a regional innovation hub, staffed with engineers who had previously worked for one of the world's leading PEM manufacturers. During our discussions, it emerged that they had recently entered a partnership with a Chinese firm. What was happening there felt, frankly, like "Alabama Fried Chicken" trying to reverse-engineer the Colonel's secret recipe: transferring world-class insights back to China in real time.

The world does not sleep. And if one party chooses to, the other will take advantage. China did not have advanced PEM technology a few years ago. It does now.

It would, however, be a mistake to assume that China's progress is driven solely by acquiring technology developed elsewhere. That narrative is becoming increasingly outdated. Alongside technology transfer and industrial scaling, China is also producing its own innovators and technology leaders.

A good example is Shanghai H-Ray. Rather than simply manufacturing technology developed abroad, the company has developed its own PEM electrolyser technology and has recently begun exporting systems beyond China's borders. Companies such as Shanghai H-Ray demonstrate that China is no longer only learning from the rest of the world; increasingly, it is contributing its own ideas, designs, and innovations to the global hydrogen industry.

This distinction matters. Catching up by adopting existing technology is one challenge. Competing with a country that combines industrial scale, manufacturing capability, and original innovation is something altogether different.

Given China's track record — in solar, wind, batteries, EVs — the pattern is unmistakable. Once a sector becomes globally strategic, China does not wait for homegrown genius. It absorbs technology through partnerships, licensing, co-development, and, where necessary, reverse-engineering. It then unleashes industrial scale and state coordination to turn that knowledge into market dominance.

And if this is what China can do with imported technology, the more unsettling question follows naturally: what happens when China begins applying its own talent to developing original hydrogen breakthroughs, as it is now doing in the EV sector? How far ahead could it leap when scaling and innovation finally converge?

If the rest of the world sleeps much longer, the catch-up will be complete — and you must ask the uncomfortable question: has it already happened?

Let's look at what is already happening. For China, hydrogen isn't a fringe policy; it has been elevated to the level of national strategic planning. In 2022, the country adopted a dedicated medium- and long-term plan for hydrogen industry development, covering the period 2021–2035. Under this roadmap, China aims to produce between 100,000 and 200,000 tons of renewable hydrogen per year by 2025, while deploying approximately 50,000 hydrogen fuel-cell vehicles on the road.

On the production side, as of 2024 China has already established around 125,000 tons per year of renewable hydrogen capacity — roughly at the lower bound of its 2025 target. The country has also expanded hydrogen refuelling infrastructure, with over 380 stations in operation, and the industrial base is ramping up rapidly. However, most of China's total hydrogen supply still comes from fossil sources, and fuel-cell vehicles remain a small fraction of the broader new energy vehicle fleet, which is dominated by battery-electric vehicles.

Even with these caveats, the plan demonstrates a systemic approach: from renewable-powered electrolysis through storage, transport, refuelling infrastructure, and application in transport, industry, and power. China is treating hydrogen not as a niche experiment, but as a strategically essential sector, similar to the approach it previously took with solar and electric vehicles. The message is clear: scale, coordination, and execution are real — the question now is whether the rest of the world can keep pace.

China is transforming its hydrogen plans from paper targets into tangible industrial capacity. Central to this effort is the massive ramp-up of electrolyser manufacturing. By 2023, Chinese firms and regional industrial clusters accounted for roughly 60% of global electrolyser manufacturing capacity. This dominance is not accidental; it reflects a strategic combination of state coordination, industrial policy, aggressive investment in infrastructure, and partnerships with Western electrolyser OEMs. Many global companies are collaborating with Chinese firms to jointly manufacture, adapt, and localise electrolyser technology, while domestic companies continue to innovate and scale production independently.

The impact of this scale is already visible in cost reductions. Chinese electrolyser factories have driven prices down dramatically, in some cases producing units significantly cheaper than their Western counterparts. This reduction in equipment costs feeds directly into the broader economics of green hydrogen.

Estimates indicate that production costs have fallen from approximately ¥42 per kilogram (~€5.5/kg) in 2019 to around ¥22 per kilogram (~€2.9/kg) in 2024.

Analysts project that costs could drop even further: potentially below ¥20/kg (~€2.6/kg) by 2025, and even as low as ¥15/kg (~€2/kg) by 2030, provided China continues to expand electrolyser production, improve efficiency, and leverage its growing supply of renewable electricity from wind and solar. The formula is clear: large-scale manufacturing, integrated supply chains, and cheap renewable power combine to create a structural cost advantage. Importantly, these reductions are not purely the result of subsidies — while government support helped jump-start the sector, the scale and efficiency gains themselves are driving real competitiveness.

What emerges is a clear pattern: China is not simply adopting hydrogen technology, it is reshaping the economics of the entire sector. By driving down costs through industrial scale, partnerships, and home-grown innovation, the country is positioning itself to dominate global green hydrogen markets in a way that mirrors its earlier successes in solar, wind, and battery production. The message is unmistakable — execution at scale, more than first-mover innovation, is proving to be China's greatest strength.

The hydrogen sector in China is being underpinned by a carefully orchestrated network of investment and coordination across the value chain. Major state-owned giants like Sinopec have committed serious capital to the sector. In 2025, Sinopec launched a hydrogen-energy-focused venture fund worth ¥5 billion (≈ €4.8 billion) to back early-stage startups, core equipment and materials development, and broader hydrogen value-chain projects. This level of funding signals a strategic commitment: the goal is not only to produce hydrogen, but to develop an integrated industrial ecosystem capable of supporting the country's long-term energy ambitions.

Traditional automakers are also aligning with this vision. Companies such as SAIC Motor and BAIC Group are incorporating hydrogen fuel-cell technology into their "new energy vehicle" strategies, combining battery-electric platforms with fuel-cell options. By doing so, they are hedging their bets while simultaneously creating a pipeline of vehicles capable of scaling as hydrogen supply and infrastructure expand. This mirrors China's broader approach to new energy: coordinated development across multiple stakeholders, ensuring that production, application, and market adoption move in tandem.

Infrastructure development further reinforces this coordinated effort. Provinces and municipalities are actively building hydrogen refuelling-station networks to

match growing demand. While subsidies for station construction and operation remain in place, many have been phased down or restructured. For example, stations meeting certain compressor-capacity thresholds receive lump-sum grants, and hydrogen producers benefit from preferential tax treatments or electricity-price advantages. This targeted support incentivises efficient expansion without relying on open-ended government handouts.

Taken together, these investments and policies reveal a strategically orchestrated ecosystem, spanning central ministries, state-owned enterprises, local governments, and private industry players. This is not a fragmented set of pilot projects — it is a coordinated, system-level approach designed to create sustainable industrial and technological leadership in hydrogen, from production to vehicle deployment.

Hydrogen in China is no longer just about fuel-cell vehicles. While fuel-cell vehicles are often the most visible application, national and regional planning documents make clear that hydrogen is intended to play a role across industry, power generation, energy storage, and chemical production. This reflects a deliberate strategy to integrate hydrogen into multiple pillars of the economy, rather than treating it as a niche or experimental technology.

In particular, China is targeting heavy-industry sectors that are difficult to decarbonise through electrification alone. Steel, chemicals, and refining processes are prime examples where hydrogen could replace coal or natural-gas-derived hydrogen, dramatically reducing carbon emissions. Some provinces and leading industrial companies are already piloting hydrogen-based solutions in these sectors, positioning themselves to scale once production and supply chains reach maturity.

This multi-sector approach demonstrates that hydrogen is being embedded systemically, not as a series of isolated pilot projects. By using hydrogen to decarbonise transportation, industrial processes, and power systems simultaneously, China is creating a feedback loop: scaling production lowers costs, which encourages adoption in more sectors, which in turn drives further investment and innovation.

Taken together, these initiatives suggest that China is positioning hydrogen as a core pillar of a future green-industrial economy. This is not about incremental or isolated experimentation; it is about building an integrated, economy-wide ecosystem capable of leveraging the country's industrial scale, policy coordination, and engineering talent. In other words, it is yet another example of China's

execution-focused genius: identifying strategic sectors, mobilising resources, and systematically embedding hydrogen across the economy.

However, arguably the most important signal is where China has already pushed the price to. Recent estimates suggest that Chinese green hydrogen production costs are now around ¥28 per kg — roughly €3.5–4/kg (with some indications that projects may already be approaching €2/kg in favourable conditions). By comparison, European green hydrogen prices still often exceed €10/kg in many markets.

This gap is not theoretical — it is already emerging in practice. It suggests that China is not only scaling hydrogen industrial capacity, but doing so in a way that creates a tangible and potentially decisive cost advantage.

In short: for China, hydrogen is not developing as a premium fuel — it is rapidly becoming a commodity. If current trends continue, China may begin to set the global "floor price" for green hydrogen, with profound implications for industry, trade, and the wider energy transition.

Whilst visiting in summer 2025, at dinner with a local customer, he told me, "in the West you debate the future; in China, we build it." At the time, I thought it was bravado. Now, looking at the solar panels, wind turbines, batteries, EVs — and increasingly hydrogen — that flow from Chinese factories into global markets, I see it differently. The world may question China's methods, its politics, or its long-term trajectory. But no one can question its impact. For better or worse, this is a country that currently has massive influence on the world.

India: The Coming Man

India is a challenging country to fathom. It is, to me, a land of profound contradictions: immense wealth sitting alongside deep poverty; world-class education systems next to schools without electricity; a vibrant, ambitious demographic living amid enormous social pressures created by sheer scale. In my younger days as a commissioning engineer, I spent time in India working on pipelines, gas-treatment stations, and related infrastructure. That period allowed me to see parts of the country far removed from the usual visitor's experience. From Panipat to Mumbai, from Gandhidham to Visakhapatnam, I glimpsed more of the "real" India than most outsiders ever do.

One image has never left me: travelling to a site just outside Ajmer, passing small dwellings where children sat doing their homework by candlelight. It revealed something profound—an unwavering drive to learn, to improve, to move forward,

regardless of circumstance. The colleagues I worked with reflected the same spirit: forward-looking, ambitious, and confident in India's potential.

Viewed through the lens of the Genius Principle introduced at the start of this chapter, India has—by sheer volume of people—the potential to stand alongside China in latent talent. Its demographic dividend, its expanding STEM base, and its cultural familiarity with engineering and problem-solving give it a deep well of future "genius." Yet for decades, much of that talent has flowed outward, powering innovation in other parts of the world rather than at home.

What is changing now, however, is the arresting sense that more of that talent is staying—and that it has something to stay for. In hydrogen specifically, you can already see the outlines of "big plays" emerging: ambitious national targets, large-scale renewable-hydrogen projects, and an emerging ecosystem of domestic innovators positioning India not merely to participate in the global transition, but to shape it.

India's demographic advantage is often described in abstract terms, but its scale is unmatched: a young, educated, increasingly tech-literate population entering the workforce at a pace no other major economy can replicate. This is the raw material of national genius. India's longstanding strengths in software, pharmaceuticals, and engineering show that when talent is given even modest opportunity, it produces globally significant innovation. The question—particularly in hydrogen—is whether India can convert this latent capability into industrial leadership at home rather than brilliance exported abroad.

The challenges remain substantial. Infrastructure gaps can slow even the most promising projects. Regulatory hurdles—varying from state to state—create uncertainty for investors who need long-term clarity. Capital constraints, especially for early-stage hydrogen technologies, can limit experimentation and scale-up. These frictions are not trivial; they have, for decades, acted as a brake on India's ability to turn potential into realised industrial power.

Yet India also has a unique opportunity to leapfrog. Unlike mature energy economies with legacy assets to protect, India can build parts of its hydrogen economy from scratch, guided by the latest technology rather than constrained by entrenched systems. Strategic partnerships—whether with the Gulf, Europe, or East Asia—are already beginning to supply capital, know-how, and export markets. The private sector, too, is stepping forward: major conglomerates are setting out multi-billion-dollar hydrogen strategies, betting that India can become a low-cost producer and eventually a global supplier.

India's motivation extends beyond creating a low-cost export industry. Hydrogen also offers a route towards improved energy security, reduced dependence on imported natural gas, and greater resilience for India's fertiliser industry, which underpins the country's agricultural productivity and food security. For a country whose economic growth remains closely linked to secure and affordable energy supplies, hydrogen is becoming part of a broader national development strategy rather than simply a climate policy.

In recent years, India has begun laying down a regulatory foundation to support this emerging hydrogen economy. The centrepiece is the National Green Hydrogen Mission, which allocates over ₹17,000 crore to accelerate domestic hydrogen production, electrolyser manufacturing, and the development of dedicated hydrogen hubs. Through the Mission's SIGHT programme, India offers production-linked incentives for both green hydrogen output and the manufacturing of electrolysers—lowering costs and reducing risk for early investors. Rather than supporting isolated demonstration projects, these measures collectively aim to develop an integrated hydrogen ecosystem encompassing manufacturing, skills, certification, research, infrastructure and export capability.

A series of complementary measures further strengthens this push: hydrogen producers using renewable electricity now receive a 25-year waiver on interstate transmission charges; import duties on key electrolyser components have been reduced to ease domestic deployment; and the government is developing a formal certification system to define what qualifies as "green" hydrogen for both domestic use and export markets. Several states have also launched their own green hydrogen policies—most notably Maharashtra, Gujarat, and Rajasthan—offering land access, capital subsidies, and concessional power tariffs to attract large-scale projects and manufacturing clusters. Together, these policies represent a deliberate attempt to turn India's ambition into an investable, scalable industrial pathway.

If India's ambition in hydrogen once seemed aspirational, the reality emerging on the ground today suggests something more substantial. The country is not merely outlining targets; it is assembling the industrial architecture required to meet them. The National Green Hydrogen Mission—launched with multi-billion-dollar funding commitments—aims to produce millions of tonnes of green hydrogen annually, supported by large-scale electrolyser manufacturing, dedicated renewable-energy zones, and state-level incentives designed to attract global investment. Across the country, pilot projects, industrial clusters, and export corridors are taking shape with a speed that reflects both domestic urgency and international opportunity.

Yet ambition alone will not determine success. As in every emerging hydrogen market, the challenge now shifts from announcing projects to delivering them. Grid infrastructure, long-term offtake agreements, financing, water availability and skilled labour will determine how much of today's project pipeline ultimately reaches commercial operation. India's opportunity remains enormous, but execution will now become the defining test.

At the centre of this movement stand India's industrial heavyweights—companies with the scale, capital, and ambition to shift national trajectories. Among them, Reliance Industries is undoubtedly the most consequential. Under its New Energy strategy, Reliance is constructing a fully integrated hydrogen ecosystem: from gigawatt-scale solar and wind generation to electrolyser production, hydrogen storage, and end-use applications in refining, chemicals, and mobility. The company's stated goal is clear—to deliver green hydrogen at a cost low enough to make India one of the world's most competitive producers. If achieved, this would reshape not only India's energy landscape but global supply chains as well.

Adani Group, too, is pursuing an expansive hydrogen agenda, leveraging its ports, logistics infrastructure, and renewable-power portfolio to build vertically integrated hydrogen hubs. Its proposed green hydrogen and ammonia facilities—some among the largest announced globally—are explicitly designed for export markets in Europe and Asia. Meanwhile, state-owned entities such as NTPC are piloting hydrogen blending in gas networks, developing hydrogen mobility solutions, and exploring green hydrogen for industrial heat—applications that could create domestic demand at scale.

These headline players are matched by a growing ecosystem of innovators. Start-ups and mid-sized engineering firms are developing electrolyser technology, hydrogen storage materials, and fuel-cell systems tailored to India's cost-sensitive market. Ports such as Kandla and Paradip are planning hydrogen-ready export infrastructure, while states like Gujarat, Rajasthan, and Karnataka are competing to attract manufacturing clusters with land, policy stability, and renewable-energy access.

Rather than creating isolated demonstration projects, India is laying the foundations for its first hydrogen islands. These regional industrial ecosystems bring together production, manufacturing, infrastructure, research, skills and investment in strategic locations. As they mature, they become natural centres of innovation and growth, attracting new supply chains, technologies and investment before ultimately connecting into a broader national hydrogen economy.

Together, these initiatives signal something decisive: India is not simply preparing to join the hydrogen economy—it is positioning itself to shape it. For a nation with a young population, strong engineering culture, and the capacity to build at scale, hydrogen offers not only an industrial opportunity, but a chance to convert latent national talent into global influence.

If India succeeds in turning its human capital into domestic hydrogen capability, it will no longer be the world's source of talent, but a centre of industrial gravity in its own right. That is the moment when India stops reacting to the global energy transition and begins directing it. As India's first Prime Minister, Jawaharlal Nehru, once observed, the country is a palimpsest—a civilisation that rewrites itself without erasing its past. If India brings that same creative layering to hydrogen—absorbing global ideas, adapting them, and building upon them—it may move from merely surprising the world to shaping the next chapter of the global energy system.

The United States: The Sleeping Giant

America is a country I know far better than China or India. Through business and holiday travel, I've spent over a year of my life there, and I like to believe that gives me some insight into its character. Yet the more time I spend with American friends and colleagues, the more I'm reminded of how difficult America is to truly understand. There is always a gap — a lacuna — in how we think and approach issues. Not hostility, not misunderstanding, just an unmistakable divergence.

And almost inevitably, the stumbling block is politics. You can't ignore it; you can't sidestep it; it inserts itself into everything. It's not just that the UK and the US are two nations divided by a common language (a line that also holds for much of Europe compared to America). It's that our mental models for how problems should be framed — and solved — have drifted apart.

It's hard to say exactly when this divergence began. Growing up in Margaret Thatcher's UK, we looked to the United States as the "leader of the free world," and for a long time that sense of shared purpose felt real. But over the last 30 years something has shifted. The common ground has thinned. On many global challenges, we now find ourselves approaching the same problem from very different angles. At times, it feels like "them" and "us," even though we are all in the same boat and desperately need aligned goals.

If I had been writing this book two years ago, this section would have read very differently. Back then, I would have been praising America's return to global climate leadership. I would have described how the US was once again demonstrat-

ing its unmatched ability to turn genius into real-world solutions — how the Inflation Reduction Act seemed to embody the country's foundational strengths: ambition, innovation, scale. For a moment, it genuinely felt like the US was leading the world into a cleaner future.

But today, the narrative is more complicated, and the questions deeper. The promise of the IRA has largely faded, its direction overtaken and reshaped by the same broader political dynamics that define so much of American life. To appreciate where the US stands in the hydrogen economy, we need to go beyond the headline policies and look at the structure of American genius itself: decentralised, powerful, creative — and often frustratingly incoherent.

At the heart of that story lies a fundamental paradox—one that the Genius Principle captures quite neatly. The US is, by any reasonable measure, the most fertile environment for innovation in human history. No nation has produced more game-changing technologies, more scientific breakthroughs, more transformative companies. But at the same time, no advanced economy makes it so difficult to turn that genius into aligned, national-scale execution.

The country's innovation model is almost anarchic in its decentralisation. It does not rely on coordinated national planning or long-term industrial strategy. Instead, American genius bubbles up from thousands of independent nodes: universities, venture capital firms, tech startups, federal research labs, industry consortia, corporate R&D divisions, and a culture that tolerates risk, celebrates ambition, and forgives failure. This decentralised ecosystem creates extraordinary creativity — and extraordinary fragmentation.

When we look at the great American success stories, this pattern constantly repeats itself. Consider the birth of the personal computer: it didn't emerge from a national strategy or an industrial masterplan, but from a scattered network of hobbyists, rebellious engineers, university dropouts, defence-funded labs, and small electronics firms experimenting on the fringes. What became one of the defining technologies of the modern world started as a decentralised swirl of ideas rather than a coordinated mission. This is the American genius at work: breakthroughs born not from order, but from creative disorder. Even the shale revolution — perhaps the most significant energy innovation of the last fifty years — was not a state programme but a chaotic, high-risk bet by a handful of stubborn entrepreneurs drilling holes in Texas and Oklahoma long before anyone thought they would strike anything.

This is the American genius: it thrives in the margins, in the gaps, in the places where no one has decided what the rules should be yet.

Yet this same decentralisation is also why the United States struggles to align itself behind long-term industrial missions. Unlike Europe or China, the US does not have a tradition of national coordination around energy systems. It has no equivalent of an EU Energy Directorate. It has no five-year plans. It has fifty states, each with its own regulatory culture, political incentives, and industrial base — and at times they seem to exist in parallel universes. For decades this fragmentation has not necessarily been a strategic disadvantage: the global hydrocarbon system evolved in ways that largely aligned with American financial, industrial, and geopolitical influence.

What emerges from this is a country capable of inventing the future, but often incapable of systematising that invention. The US can produce the world's best electrolysers, catalysts, materials science breakthroughs, and hydrogen software platforms — but that doesn't mean it can build pipelines, coordinate markets, or design regulatory certainty. Not because it lacks capacity, but because its system isn't built for coherence; it's built for creative chaos.

And this is the central tension when we apply the Genius Principle to America: the abundance of individual brilliance does not automatically translate into collective capability. It requires alignment — political, regulatory, industrial — and this alignment is becoming harder to achieve.

For hydrogen, this matters enormously. Hydrogen is not a product you can invent in a garage or deploy through a smartphone app. It is an infrastructural technology — it requires pipelines, storage caverns, refineries, transport corridors, port facilities, long-term power contracts, and enormous capital. In other words: hydrogen demands coordination.

And coordination is the one thing America is not naturally good at.

This is where the paradox sharpens. If hydrogen were purely a technological race, the US would already be the global leader. If it were purely a capital race, the US would dominate. If it were purely a research challenge, the national labs would outperform everyone. But hydrogen is not just those things — it is a systems challenge. It demands integration across massive industries and jurisdictions, coordination that the United States finds notoriously difficult. And in recent years, that difficulty has taken on an extra layer: hydrogen, and clean energy more broadly, has become an ideological battleground. No longer simply a matter of economics or engineering, the debate is now entangled with partisan politics, cultural identity, and broader disputes over the role of government. Even technologies with immense promise can stall when they are cast as symbols in a political struggle, rather than tools for practical progress.

Yet the lesson from history is also clear: when America does manage to overcome this fragmentation — when it aligns its genius behind a common purpose — it does so with a force that reshapes the world. The Manhattan Project, Apollo, the interstate highway system, the biotech boom — these were moments when the US system, usually so decentralised, briefly cohered into something like a national mission.

Hydrogen will require something similar.

Whether the US can achieve that — whether it can transform decentralised brilliance into coordinated scale — is the question that will determine its place in the next energy era.

The IRA: A Vision for Awakening the Sleeping Giant

When the Inflation Reduction Act became law in 2022, it felt, for a moment, like the United States was stirring from a long slumber. Here was a government policy designed not merely to subsidise a technology, but to harness the full scale of American ingenuity and industry for a decisive shift toward clean energy. Hydrogen, often dismissed as exotic or futuristic, was suddenly cast at the centre of a national vision.

At the heart of the IRA's promise was the new federal tax credit for clean hydrogen production. Known as Section 45V, it offered producers up to three dollars per kilogram for hydrogen meeting strict emissions standards. This was no minor incentive — it had the potential to make hydrogen cost-competitive with fossil fuels, transforming the economic calculus for investors and industrial actors. Even hydrogen that was slightly less clean still qualified for smaller credits, creating a graded system that encouraged continuous improvement. Critically, the credits were transferable and eligible for direct payment, meaning that companies could access funding upfront rather than waiting for tax reconciliation — a structural innovation designed to overcome the risk that has so often slowed American infrastructure projects.

Yet the IRA did not stop at financial incentives. It envisioned a network of regional clean hydrogen hubs, distributed across the country, each acting as a nucleus for a new hydrogen economy. California, Texas, the Midwest, Appalachia, the Northeast — these hubs were to coordinate production, storage, distribution, and end-use. They were designed to connect innovators with industry, connect capital with infrastructure, and connect ambition with execution. It was an attempt to convert America's famed decentralised genius into coordinated, systemic change — to create coherence where the system is usually so delightfully chaotic.

The scale of ambition was extraordinary. Federal dollars were intended to catalyse massive private investment, multiplying their impact. These hubs were not about small pilot projects; they were about real, industrial-scale hydrogen production, capable of delivering millions of tonnes per year, and decarbonising the sectors that are hardest to clean: steel, fertiliser, ports, shipping, trucking, and utilities. For the first time, hydrogen was not an idea on the margins, it was a pillar of a national industrial and environmental strategy.

And the promise extended beyond energy. The hubs were designed to create jobs, revitalise regions that had seen industry decline, and secure American competitiveness on the global stage. The IRA envisioned hydrogen as both a climate solution and an economic engine, a rare example of policy marrying innovation, industrial power, and social impact.

For a moment, the IRA felt like the US had found a path to align its creative chaos — the universities, venture capital, startups, industrial labs — with the demands of infrastructure and scale. It was a vision of the "sleeping giant" waking: a country capable of leading not only in invention but in execution, using its extraordinary innovation ecosystem to deliver results that could reshape the global energy landscape.

From Promise to Uncertainty: The Trump Disruption

For a moment, the Inflation Reduction Act seemed to hold the promise of something extraordinary. It was an attempt to convert America's decentralised genius into national-scale action: a map for hydrogen infrastructure, a mechanism to de-risk investment, and a pathway to a new industrial renaissance. The hubs were meant to be the skeleton of a hydrogen economy, linking production, transport, and industrial use in ways that could finally awaken the sleeping giant.

Hydrogen sat right at the centre of that promise. The clean hydrogen production tax credit — the 45V provision — offered up to three dollars per kilogram for the cleanest hydrogen, a level of support capable of fundamentally reshaping project economics. By early 2025, detailed federal guidance had finally clarified how "clean" hydrogen would be defined and how lifecycle emissions would be assessed. After months of uncertainty, investors at last had a framework. For short while, policy, capital, and industrial ambition appeared aligned.

But that vision has run into an all-too-familiar American obstacle: politics. With Donald Trump's return to the presidency, the delicate architecture of the IRA came under immediate strain. Federal funding for the regional hydrogen hubs was paused, reviewed, and in several cases restructured or withdrawn as part of a

broader reassessment of clean-energy spending. Contracts that were expected to seed billion-dollar projects suddenly carried new uncertainty. What had been conceived as a coordinated network of seven regional hubs — stretching from California to the Midwest, Appalachia, and the Gulf Coast — no longer felt like a guaranteed national build-out. Some projects continue to move forward, particularly where state-level support remains strong, but the overarching federal momentum that once bound them together has weakened. Instead of a clear national rollout, the hub programme now reflects a more fragmented reality, dependent on shifting political priorities and local resilience rather than unified federal direction.

The instability did not stop there. A sweeping tax-and-budget overhaul in mid-2025 modified elements of the clean-energy incentive landscape created under the IRA. Even where credits technically remained in place, timelines and eligibility conditions shifted. For hydrogen projects that depend on multi-decade returns and stable financing assumptions, this was not a minor adjustment — it altered risk perception. Financing models tightened. Risk premiums rose. Some projects slowed; others quietly went back to the drawing board.

By late 2025, the Department of Energy had terminated or re-evaluated hundreds of federally backed projects across multiple programmes, citing economic viability and fiscal discipline. While not all were hydrogen-related, the broader message to the market was unmistakable: federal participation could not be assumed to be durable. For an infrastructure vector like hydrogen — which requires long-term coordination between public and private actors — durability is everything.

And markets, being rational, respond to uncertainty.

One of the clearest corporate signals came from Accelera, the hydrogen-focused division of Cummins. Hydrogenics — a pioneering electrolyser company — had been acquired by Cummins as part of a long-term bet on a hydrogen-led future. Like much of the sector, the business was operating at a loss, hardly unusual in an emerging infrastructure market. Yet in late 2025, Cummins took substantial charges connected to its electrolyser operations and signalled a pullback from new commercial deployment activity in this area.

This was not a verdict on hydrogen as a technology, nor even on electrolysis itself. It was a commercial judgement shaped by timing, policy certainty, and market visibility. In effect, a major U.S. industrial player looked at the near-term prospects for green hydrogen deployment at scale in the United States and concluded: not yet — and not here.

For a country whose strength lies in innovation, but whose weakness lies in coordination, this decision is less an anomaly than a symptom.

This is the American paradox writ large. Even when the country devises a policy capable of mobilising its genius at scale, the promise can be diluted — or reversed — within a single political cycle. Innovation and capital remain abundant. The laboratories are still world-class. The engineering depth is unquestioned. But hydrogen is not a software market. It is pipelines, storage caverns, port infrastructure, offtake agreements, and assets designed to last decades. It demands durable coordination.

When hydrogen becomes another ideological football, as it has under Trump, even the most advanced technologies struggle to escape the gridlock of partisanship.

For hydrogen, as for so many complex systems in the United States, this is the moment where potential meets political reality — and the sleeping giant risks nodding off once more.

Yet there is another interpretation of what is happening in the United States today. From a distance, it is easy to conclude that momentum has simply evaporated. The headlines certainly point in that direction. Funding has been delayed, projects have been reconsidered, and political attention has shifted elsewhere.

But the view from inside the industry is often more nuanced. Many companies understand that the hydrogen economy extends far beyond the United States. They operate globally, serve international markets, and recognise that the long-term drivers behind hydrogen — energy security, industrial competitiveness, and decarbonisation — have not disappeared simply because political priorities have changed.

As a result, some parts of the ecosystem appear to have become quieter rather than disappearing altogether. Research continues. Partnerships continue. Technologies continue to mature. Capabilities are maintained. Yet much of this activity now takes place with less fanfare and fewer headlines than during the peak enthusiasm surrounding the Inflation Reduction Act.

This is perhaps another manifestation of the American model. Innovation often continues beneath the surface, outside the spotlight, waiting for the right market conditions, policy environment, or commercial opportunity to emerge. The risk in assessing the United States is therefore not only overestimating its ability to coordinate, but also underestimating how much progress may still be occurring beyond public view.

Hydrogen Reimagined: The Conditional Rise of Blue

Despite the disruption of the IRA's original vision, hydrogen is far from dead in the United States. The promise of green hydrogen — electricity-driven, zero-emission, renewable — has stumbled under political and regulatory uncertainty. But where one path slows, another emerges. Today, the focus has shifted decisively toward blue hydrogen, produced from natural gas with carbon capture and storage. It is a pragmatic, if controversial, adaptation: a form of hydrogen that leverages the country's existing fossil infrastructure while still offering a pathway to emissions reduction — at least in theory.

The driver of this shift is clear: the oil and gas industry. Major energy firms — long accustomed to navigating political uncertainty and scaling massive projects — are seizing the moment. They recognise that, even if green hydrogen remains constrained by policy and cost, blue hydrogen offers a tangible, investable opportunity. Federal and state incentives may have shifted, but private capital is still flowing, and technology is advancing. The O&G sector is now positioning itself at the centre of a hydrogen economy that aligns with their existing assets: natural gas pipelines, storage facilities, industrial hubs, and technical expertise in large-scale gas processing.

What this means is a subtle but important redefinition of the U.S. hydrogen story. The sleeping giant is not asleep; it has simply changed posture. Hydrogen continues to attract investment, innovation, and industrial experimentation — but the terms have shifted. Where once the vision was a nationwide network of renewable-powered hubs, now the narrative centres on leveraging carbon capture, retrofitting industrial clusters, and creating markets for hydrogen that are compatible with existing energy infrastructure.

Blue hydrogen also demonstrates a uniquely American approach to energy transition: it is opportunistic, pragmatic, and heavily influenced by the players with the deepest pockets and operational experience.

There is, however, a critical condition embedded in this pathway: the carbon capture element must actually be delivered, operated, and maintained at scale. Blue hydrogen only fulfils its climate promise if the associated CO_2 is captured reliably, transported securely, and permanently stored. Building a reformer without building the full carbon management system is not transition — it is delay. History offers reasons for both optimism and caution. While carbon capture technologies are well understood and have been deployed in industrial contexts for decades, large-scale, high-capture-rate implementation has often lagged behind initial ambition.

And this is where the current U.S. climate introduces uncertainty. Announcing a blue hydrogen plant is one thing; financing, permitting, and constructing the full CCUS chain — pipelines, injection wells, long-term monitoring, and liability frameworks — is another. In a policy environment marked by regulatory volatility and shifting federal priorities, it remains to be seen how many blue projects will integrate carbon capture at the scale and capture rates assumed in their models. The risk is not technological impossibility, but incomplete execution.

In this sense, blue hydrogen remains a "show-me" pathway. Its success will not be determined by announcements or investment commitments alone, but by measured capture rates, verified storage performance, and long-term operational discipline. Whether the carbon capture components of these projects are built — and built to the standards implied in modelling assumptions — will ultimately decide whether blue hydrogen functions as a credible bridge or merely a rebranding of existing gas use.

While the ideology around green hydrogen has polarised policy debates, blue hydrogen has quietly advanced, driven less by political symbolism and more by industrial pragmatism. In a system where coordination is hard and policy swings can stall projects; this may be the most politically feasible path forward — but feasibility is not the same as integrity.

In short, hydrogen in the U.S. has not vanished — it has evolved. The sleeping giant is still stirring, just in a form that reflects both the country's political reality and the strengths of its industrial ecosystem. Blue hydrogen may not be the revolutionary vision once imagined, and its climate credibility will ultimately depend on whether carbon capture is executed fully rather than partially. If it is, blue hydrogen could lay the groundwork for a broader hydrogen economy. If it is not, it risks becoming a transitional narrative without a transformational outcome.

The Future: What Awakening the Sleeping Giant Would Look Like

The question now is what it would take for the United States to fully unleash its hydrogen potential. Hydrogen, as we have seen, is not merely a technological challenge or a matter of capital; it is a systems-level mission requiring coordination across federal and state governments, industry, and infrastructure. It demands policies that are long-term and bipartisan, designed with the same continuity and resilience as defence strategy or Medicare — commitments that survive administrations and political cycles.

Achieving this would require more than incentives and hubs. It would mean linking the hydrogen economy to national narratives of competitiveness and industrial leadership, framing it as essential to America's place in the world. There is a real risk that, without such alignment, leadership could pass to China, Germany, or the Gulf states — nations willing to commit for the long haul.

There is also a deeper structural dimension to the transition that is rarely discussed openly. For much of the last half-century, global energy markets — particularly oil — have been intertwined with the dollar-based financial system. As hydrocarbons are widely traded in U.S. dollars, the architecture of energy trade has also reinforced American monetary and geopolitical influence. A world in which countries increasingly produce their own energy — through renewables, hydrogen, or electrified systems — would not only transform the energy landscape; it would gradually reshape that financial architecture as well. That reality does not dictate American policy, but it does form part of the strategic backdrop against which energy decisions are made.

In that sense, the energy transition is not merely an industrial transformation; it is a quiet reconfiguration of the geopolitical and financial order built around the fossil-fuel age.

We can imagine three possible scenarios. In the first, the US becomes a global leader, marrying its decentralised genius with coordinated national strategy, mobilising private and public capital, and creating a hydrogen infrastructure that underpins industrial growth, decarbonisation, and geopolitical influence. In the second, the US remains a laggard, producing brilliant innovations in isolation but failing to scale them, leaving other nations to define the rules of the emerging hydrogen economy. In the third, America acts as a disruptor, advancing niche technologies and market models that force competitors to adapt, but without ever fully realising the full systemic potential.

However, there is a fourth variable in this equation — one that may ultimately matter more than federal policy alone: the behaviour of multinational corporations.

Global companies do not think in four-year cycles. They think in decades. Faced with U.S. volatility, they have choices. Some may take the Cummins route — trimming exposure, redeploying capital, and focusing on jurisdictions where regulatory frameworks are more predictable. Capital is mobile. Engineering talent is mobile. Manufacturing can be relocated. If the U.S. environment becomes persistently unstable, multinationals will simply follow certainty elsewhere.

If the United States becomes merely a market for imported hydrogen technology rather than its architect, its geopolitical leverage in the energy transition would shrink accordingly.

Others, however, may choose to stay the course. They may interpret current turbulence not as structural decline, but as cyclical interruption. They may continue investing selectively, maintaining optionality, watching developments in Europe, China, and the Middle East, and positioning themselves for a potential policy realignment after the next election cycle. For these firms, hydrogen in the United States is not dead — it is paused.

The strategic risk for America is not that hydrogen disappears, but that industrial ecosystems quietly relocate. If scale happens elsewhere — if standards are written elsewhere, if supply chains consolidate elsewhere — then even a future policy pivot may not fully restore lost ground. Leadership in infrastructure sectors compounds over time. Once clusters form, they are hard to unwind.

The path to awakening the sleeping giant is thus clear in theory, though fraught in practice. It will require bridging the ideological divides that now frame energy debates, overcoming the fragmentation of federalism, and converting the country's prodigious creative capacity into coordinated, long-term industrial execution. Only then can the US transform its promise in hydrogen into tangible, world-shaping outcomes — finally turning the sleeping giant into a force that leads once more.

The Rest of the World

Now, I apologise if you happen to be reading this in a country not highlighted above. You may be thinking, "but my country also has a hydrogen industry that deserves attention," and you would probably be right. The fact is, according to the IEA, 60 countries — 30 advanced and 30 emerging — already have hydrogen policies in place, with incentives and actions occurring on a global basis. For the planet to meet the goals laid out in earlier chapters — replacing 620 EJ of energy — action is required everywhere.

Before moving on to export strategies, industrial clusters, and national roadmaps, it is worth pausing to consider a different expression of hydrogen geopolitics — one that does not begin in ministries or ports, but at the edge of the grid itself. In some parts of the world, hydrogen is not primarily about trade or industrial dominance, but about resilience, access, and sovereignty. Few places illustrate this more clearly than Namibia.

Namibia — The First Light of a Hydrogen Continent

Just outside Walvis Bay, where the Namib Desert leans into the Atlantic in long, pale strokes of sand and salt, a small cluster of solar panels and steel tanks stands under a fierce blue sky. It looks almost delicate against the vastness around it, like a tiny outpost planted at the edge of the world. But this modest installation is poised to answer a question far larger than itself: how do you bring reliable, clean energy to places the grid cannot reach — and to places where the grid breaks?

The system emerging here is a hydrogen microgrid, developed under the guidance of Dr.-Ing. Ulrike Beyer and her team at Fraunhofer IWU in Chemnitz. It is as compact as it is ambitious. Sunlight hits the panels, electricity feeds a modular electrolyser, in this case from Enapter, and water is split into hydrogen and oxygen. The hydrogen is stored in quiet, pressurised tanks; the oxygen is channelled into local water treatment to improve irrigation quality. And as soon as the sun dips beneath the dunes, a fuel cell steps in, turning the stored hydrogen back into electricity with nothing but warm water vapour drifting into the cooling night.

From a distance, you wouldn't know it was anything unusual. There is no smoke, no noise, no vibration, none of the signatures of a diesel generator fighting its way through the day. The microgrid breathes in sunlight and breathes out electricity. That is all. And for the team building it, that simplicity is the point. Because power that works is only useful when it works everywhere — in cities, yes, but also far beyond them, in forgotten towns, along unpaved roads, and in the places where disaster strikes hardest and infrastructure comes last.

Namibia is an instructive place to begin this story. It is one of the world's sunniest countries, a land of extraordinary renewable potential, and one of the first African nations to articulate a national hydrogen vision. Its government has charted a course toward producing green hydrogen and ammonia at industrial scale, with the Hyphen Project already drawing international investment and attention. But alongside its grand export ambitions, Namibia also recognises a deeper truth: the energy transition must reach people where they are, not only where supply chains and shipping routes would like them to be.

That is why this microgrid matters. It is planned to enter full operation in 2026, built not in a laboratory or a controlled testbed but on Namibian soil, with Namibian partners including the Namibia University of Science and Technology. It will provide electricity to a local school, support water treatment, and serve as a training ground for technicians who will one day install and maintain systems like it across the region. It is a beginning — small, deliberate, and profoundly symbolic.

Standing on the sand, watching the systems come together, you can feel the weight of its implications. The microgrid is not designed only for development settings. It is also engineered to be rapidly deployable, its components ruggedised, modular, and transportable. The same containerised system that keeps a Namibian school running could just as easily be sent into a disaster zone. When storms flatten power lines, when earthquakes silence substations, when conflict or drought break supply routes, this technology offers something invaluable: energy independence, no fuel deliveries required. A pallet of solar panels, a small electrolyser, a fuel cell, and a pair of hydrogen tanks can keep a clinic open, a water pump running, a communication tower online. In humanitarian response, that can be the difference between a crisis and a catastrophe.

This is what excites the Fraunhofer team — the idea that hydrogen, often framed in terms of megaprojects and global supply chains, can also operate at human scale. A microgrid like this can support remote agricultural communities by powering cold storage for crops. It can electrify rural clinics where darkness means vulnerability. It can give teachers classrooms where computers stay on during the day and lights stay on at night. It can bring industry to villages that otherwise drift at the margins of the economy. And critically, it can do all of this without diesel generators, without the endless logistics of fuel trucking, and without the emissions that have weighed on African air and economies for decades.

Yet the narrative does not stop at a single microgrid. Namibia sees hydrogen as a route to industrialisation — a new economic pillar that can create jobs, attract manufacturing, and reduce dependence on imported fuels. The country's plans stretch from large export facilities to local mobility pilots, mining applications, and renewable-powered desalination. If these systems flourish, Namibia could become one of Africa's hydrogen leaders, exporting energy molecules while applying the same technology to lift its own communities.

And this pilot is the first real thread of that future. A thread that ties together global research networks, African educational institutions, local technicians, European science funding, and a technology that feels almost too elegant to be possible in such a stark and simple landscape.

As the afternoon light begins to fade across the desert, the panels continue to harvest the last energy of the day. A faint click from the electrolyser signals another kilogram of hydrogen stored. Soon, the fuel cell will take over, its quiet hum replacing the sun's intensity. If you were standing next to it, you might not notice the shift. The lights would stay on. The systems would continue as they always do. And that is the extraordinary part — the way the future arrives not in

a roar but in a whisper, working exactly as it should, without demanding attention.

Someday, systems like this may dot the edges of deserts, forests, mountains, coastlines, refugee camps, mining outposts, and towns struck by storms. Someday, the idea of sending diesel into disaster zones may seem as outdated as sending coal to power a hospital. Someday, a continent long defined by energy scarcity may find itself exporting clean molecules while electrifying its own communities with the same sunlight that pours down freely each day.

And when that happens, someone may remember this first installation — this small cluster of panels and tanks near Walvis Bay — and recognise it for what it was: not a demonstration, not a prototype, but a declaration. A message, written quietly in hydrogen and sun, that Africa's role in the energy transition will not be to wait but to lead.

Namibia's microgrid is not a substitute for national hydrogen strategies or export ambitions — it is a complement to them. Other countries are pursuing different paths, shaped by geography, scale, and strategic intent.

Even though we often talk about leaders and laggards in the global hydrogen race, no single country will dominate entirely. While China accounts for roughly 60% of global electrolyser manufacturing capacity today, the scale of the challenge means that we will need manufacturers and innovators from Norway to Chile.

Chile is an instructive example. It is positioning itself as a global leader in green hydrogen through its National Green Hydrogen Strategy, aiming for 25 GW of electrolysis capacity by 2030 and domestic production of 1 million tonnes per year. The strategy combines ambitious production targets with a regulatory framework, tax incentives, and the creation of "hydrogen districts" to attract investment and foster local industry. A cornerstone is local electrolyser manufacturing: by 2025, three companies — Hygreen Energy, a Chinese consortium, and a Spanish consortium — were selected to build assembly facilities ranging from small to industrial-scale units, supported by government co-financing and local content requirements. Through these efforts, Chile seeks to internalise value, reduce costs, and establish a fully integrated hydrogen ecosystem, laying the groundwork to become a major exporter of green hydrogen and derivatives.

Beyond Chile, a variety of smaller nations and regional clusters are carving out niches in the global hydrogen landscape. Japan, for example, has adopted its Clean Fuel Value Chain (CfV) strategy, which focuses on building an end-to-end hydrogen ecosystem — from production and distribution to end-use — and positions

the country as a leader in fuel cell deployment and hydrogen-import infrastructure. South Korea, meanwhile, is developing a network of hydrogen valleys — regional innovation hubs where industry, research, mobility, and infrastructure are co-located — and continues to advance one of the world's most mature fuel cell industries, particularly for stationary power. Australia seeks to leverage its abundant renewable resources to pursue large-scale green hydrogen production for export, while the Gulf States aim to integrate hydrogen into their existing energy infrastructure and become major exporters.

Even countries without large populations or industrial bases can exert outsized influence by identifying niches, fostering specialized expertise, and aligning policy, investment, and infrastructure with strategic objectives.

These examples underscore a broader point: the global hydrogen landscape is not a zero-sum game. Success does not hinge solely on size or existing industrial might. Instead, innovation thrives where countries can identify specific niches, leverage unique resources, and align policy and infrastructure to empower talent. Smaller players can excel by focusing on areas such as advanced fuel cells, hydrogen storage, or local electrolyser manufacturing, often outpacing larger nations in specialised fields. At the same time, collaboration — through technology partnerships, cross-border supply chains, or shared standards — allows these niches to scale globally, amplifying their impact.

In practice, this means the hydrogen ecosystem will be highly interconnected and complementary. For instance, countries like Chile can export green hydrogen to markets where Japan and South Korea supply fuel cells and storage technology, while Australia and the Gulf States provide large-scale production to meet industrial demand. This interplay of capabilities creates a landscape where multiple players, large and small, contribute to solving the enormous global energy challenge. Genius in hydrogen, therefore, is not simply about who produces the most or invests the largest sums, but about strategic focus, coordinated action, and the ability to turn regional strengths into global influence. The coming decade will reveal a mosaic of players, each leveraging its own combination of resources, talent, and vision to shape the hydrogen economy in distinctive and sometimes surprising ways.

The Protectionism Trap

One of the defining geopolitical questions of the energy transition is how nations balance competitiveness with protection. Across Europe, North America, India, and elsewhere, governments are increasingly intervening in markets through tar-

iffs, local-content requirements, subsidies, and industrial policies designed to support domestic industries.

The reasoning is understandable. No government wishes to see strategically important industries disappear, particularly when those industries may play a central role in future economic growth, employment, and energy security. Providing domestic companies with time to adapt, invest, and scale can therefore be entirely rational.

The key word, however, is time.

Protectionism does not create competitiveness. It creates an opportunity to become competitive.

The automotive sector provides a useful example. During recent visits to China, I was struck by the pace of development taking place within the electric vehicle market. Many of the vehicles now available there are simply not available in Europe. Some are technologically impressive, highly integrated, competitively priced, and advancing at a remarkable rate.

European and American manufacturers have responded, in part, through various forms of market protection. Whether justified or not, such measures can provide breathing space for domestic producers. Yet the rest of the world does not necessarily operate under the same conditions.

A customer in South America, Southeast Asia, Africa, or the Middle East is often less concerned with protecting European or American industry than with purchasing the product that best meets their needs. They will typically evaluate performance, reliability, features, productivity, service support, and cost. In other words, they operate on a far more open playing field.

This same debate is increasingly visible within the hydrogen sector. Numerous industry leaders have publicly argued that European electrolyser manufacturers require greater protection from international competition, particularly from lower-cost imports. The concern is understandable. Electrolysers are strategically important technologies, and Europe has invested heavily in developing domestic capabilities. Few policymakers would wish to see those capabilities disappear before the market has had an opportunity to mature.

The concern extends beyond preserving market share or protecting existing jobs. Industrial capability is not something that can simply be recreated once it has disappeared. Electrolyser manufacturing is supported by far more than production lines alone; it depends upon engineers, specialist suppliers, research institu-

tions, manufacturing expertise, testing facilities, certification bodies and practical experience accumulated over many years. If these ecosystems are allowed to erode, rebuilding them can take decades. In this sense, supporting domestic hydrogen manufacturing is not simply about today's commercial success, but about retaining the capability to innovate, compete and respond to future strategic needs.

This should not be interpreted as an argument for permanent protection. Competitive industries must continue to innovate, improve productivity and earn their place in global markets. Rather, it is an acknowledgement that industrial ecosystems are strategic assets in their own right. Once lost, they are often far more difficult—and far more expensive—to rebuild than many policymakers anticipate.

There is, however, another dimension to this problem. Emerging industries do not only face the risk of eventually becoming uncompetitive; they also face the risk of disappearing before the market they were created to serve has fully developed. Many companies across the hydrogen economy have invested in engineering capability, manufacturing capacity, people and products in anticipation of demand that governments, industry and energy strategies all expected to emerge. When that demand develops more slowly than anticipated, even technologically capable companies can find themselves caught between the investment already made and the revenues that have yet to arrive.

This creates an uncomfortable policy question. If governments regard technologies such as electrolysers, hydrogen compression, storage, refuelling equipment or fuel cells as strategically important, it is not enough simply to set deployment targets for 2030 or 2040. The industrial ecosystem required to deliver those targets must still exist when the market finally arrives. That does not mean keeping every company alive indefinitely, nor protecting businesses that cannot ultimately compete. But it does mean recognising the danger of allowing viable technologies, specialist knowledge and manufacturing capability to disappear during the very period in which the market is struggling to establish itself. The challenge is therefore not merely to create future demand, but to build the conditions that allow the strongest parts of the emerging industrial ecosystem to survive long enough to meet it.

Beyond the economics, another factor underpins many of these discussions: the perception of whether competition is taking place on a genuinely level playing field. Across industry there are frequent debates regarding differing levels of state support, industrial policy and financing available to manufacturers in different regions of the world. While the extent of such support is often difficult to quantify

objectively, the perception alone is sufficient to influence political debate and strengthen calls for protective measures. Ultimately, the objective should not be to eliminate competition, but to ensure that competition encourages innovation, investment and long-term competitiveness.

Yet the same question remains: what is done with the time that protection provides?

This is where the danger emerges.

A company can maintain market share within a protected environment while simultaneously losing competitiveness internationally. The domestic market may appear healthy, yet global competitors continue to improve. They innovate faster. They reduce costs. They scale production. They refine their products. Eventually the gap becomes difficult to ignore.

The lesson is not that protectionism is inherently wrong. In many circumstances it may be entirely justified. Emerging industries often require support. Strategic capabilities may need safeguarding. Critical technologies may warrant temporary intervention.

The question is what happens next.

If the period of protection is used to create genuinely world-leading products, improve productivity, strengthen supply chains, and drive innovation, then protection has served its purpose. It has bought the time necessary to compete.

If, however, protection becomes a substitute for improvement, it risks creating complacency. Companies may begin competing against regulations rather than competitors. Success becomes measured by preserving market position rather than by building products that customers around the world actively choose.

This challenge extends well beyond the automotive sector. The same pressures can be found in hydrogen, batteries, fuel cells, electrolysers, renewable energy technologies, and countless other industries that will shape the future energy economy.

Ultimately, global markets tend to be unforgiving. Customers rarely purchase products out of sympathy. They buy the solutions that offer the best combination of performance, reliability, productivity, and cost.

Protection may buy time.

What determines success is how that time is used.

Resilience Under Pressure

Recent disruptions to global oil markets during the first half of 2026 served as a timely reminder that the geopolitics of energy are shaped not only by resources, but by resilience. While the disruption itself centred on oil, the underlying lessons extend far beyond hydrocarbons. Countries with diversified supply chains, strategic storage, flexible infrastructure and multiple import options weathered the disruption far more effectively than those dependent upon a small number of suppliers or transport corridors. The impact was not evenly distributed. Some regions experienced genuine concerns over energy security, while others remained comparatively insulated. Geography, infrastructure and preparedness mattered just as much as the availability of energy itself.

For hydrogen, this offers an important insight. The energy transition does not eliminate geopolitical risk; it changes where that risk resides. Future hydrogen systems will still rely on ports, pipelines, storage facilities, electricity networks, critical materials and international supply chains. The objective therefore cannot simply be to produce the lowest-cost kilogram of hydrogen. It must also be to build systems capable of absorbing disruption. Strategic storage, diversified production, flexible infrastructure and multiple trading relationships are not inefficiencies to be minimised—they are investments in resilience.

This also challenges one of the assumptions that has underpinned the fossil-fuel era. For decades, energy security has often been viewed primarily through the lens of securing imports. The emerging hydrogen economy presents an opportunity to rethink that model. Wherever practical, countries will seek to maximise domestic production from renewable resources, supported by international trade rather than wholly dependent upon it. In that sense, hydrogen is not merely another energy commodity; it is an opportunity to redesign the architecture of energy security itself.

Energy Deficit and Energy Surplus Nations

One of the most common assumptions surrounding the future hydrogen economy is that energy-deficient nations will simply replace today's oil and gas imports with tomorrow's hydrogen imports. The logic appears straightforward. Countries with abundant renewable resources produce energy. Countries with limited resources buy it. The model has worked for hydrocarbons for decades, so why should hydrogen be any different?

The answer is that hydrogen changes some of the underlying economics.

For much of the fossil-fuel age, transporting energy was often easier than producing it locally. Oil could be shipped around the world. Natural gas could be moved through pipelines or liquefied and transported as LNG. The result was a global energy system in which many countries became accustomed to importing a significant proportion of their energy needs from elsewhere.

This has shaped how many policymakers continue to think about the future. Discussions frequently focus on importing hydrogen from Australia, the Middle East, North Africa, South America, or India. The assumption is that today's energy relationships will simply be replicated with different molecules.

Yet hydrogen does not behave like oil.

As discussed earlier in this book, hydrogen is difficult and expensive to transport over long distances. Compression requires energy. Liquefaction requires even more energy. Converting hydrogen into ammonia, methanol, synthetic fuels, or other carrier molecules introduces additional losses and costs. Upon arrival, many of these products must then be converted again before they can be used. Every additional step adds complexity, inefficiency, and expense.

This does not mean international hydrogen trade will not exist. It almost certainly will. Some regions possess renewable resources so abundant that exporting energy will remain economically attractive. Australia, Chile, parts of the Middle East, and certain African nations may ultimately become major exporters of hydrogen-derived products.

However, it would be a mistake to assume that every country currently being discussed as a future exporter will prioritise global markets over domestic needs.

Take India as an example. The country has ambitious hydrogen plans and enormous renewable-energy potential. Yet it is also home to a rapidly growing population, expanding industry, increasing electricity demand, and, in some regions, ongoing challenges around energy reliability and grid stability. It is difficult to imagine any government willingly exporting large quantities of energy if doing so comes at the expense of domestic economic development or the expectations of its own citizens.

This is not unique to India. Governments are ultimately accountable to their own populations. Energy security, economic growth, industrial competitiveness, and social stability will almost always take priority over export ambitions.

As a result, the future hydrogen economy may look very different from the fossil-fuel system it is often compared against.

Recent discussions surrounding international LNG trade provide a useful reminder that energy dependence can have consequences beyond price and availability. Where a country or region relies heavily on imported energy, it may also find itself under pressure to modify regulations or standards in order to secure supply. This is not unique to any one supplier or fuel; it is a structural feature of dependence itself. One of the less obvious benefits of increasing domestic renewable and hydrogen production is therefore not simply lower emissions, but greater strategic freedom. Imports will always remain part of the global energy system, but resilience comes from ensuring they remain a choice rather than a necessity.

Rather than relying predominantly on imported energy, many nations may find themselves pursuing a strategy of maximum practical domestic production, supplemented by imports where necessary. The objective shifts from import-first to domestic-first. Imports remain important, but they become a complement to local generation rather than the foundation of the system.

In some cases, this localisation may be surprisingly literal. Rooftop solar, local wind generation, community energy systems, distributed storage, and decentralised production all bring energy closer to the point of consumption. The energy transition is therefore not only changing what energy we use, but also where it is produced.

This represents a profound shift in thinking. For decades, energy policy has largely been shaped by access to distant resources. The emerging energy system increasingly rewards nations that can harness resources within their own borders.

The countries that thrive in this new world will not necessarily be those that secure the largest number of import agreements. They will be those that first understand how much of their future energy demand can realistically and economically be met close to home.

International energy trade will remain important.

But the age of assuming that someone else will provide most of your energy may gradually be coming to an end.

Risks and Fragility

Genius alone does not guarantee success. Even the most innovative nations are vulnerable to structural and systemic risks that can derail progress. Political instability, regulatory misalignment, supply chain vulnerabilities, and economic shocks all pose significant threats to realising national potential in hydrogen and other emerging technologies.

Europe, despite its technological expertise and ambitious climate agenda, often struggles with regulatory complexity that slows decision-making and fragments markets. The United States, historically a global leader in innovation, faces its own form of fragility through political gridlock, where long-term strategy can be derailed by short-term electoral swings. India, with its vast talent pool and demographic advantage, contends with infrastructure gaps and logistical bottlenecks that complicate industrial-scale deployment.

These issues point toward a broader truth: the hydrogen transition is simply too vast for any single nation to "lead" the world in every dimension. The scale of infrastructure required, the diversity of resources, and the interconnectedness of supply chains mean leadership will be distributed — specialised, regional, and conditional. Nations may lead in electrolysers, or pipelines, or export markets, or technology, but none will dominate the entire system. This is not a weakness of the transition; it is a defining characteristic.

Yet distribution does not mean irrelevance. If the United States occupies a diminished role, the system will not stall — but it will evolve differently. Standards may be shaped more heavily by European regulatory philosophy. Cost curves may be driven more aggressively by Chinese scale. Trade corridors may consolidate around Gulf production and Asian demand. The hydrogen economy would still advance, but its rules, supply chains, and industrial geography would reflect the priorities of those who commit consistently.

The emerging geopolitical landscape further complicates this picture. As John Zubrzycki notes in 'The Shortest History of India', many developing nations increasingly view China as the model for rapid economic ascent — admiring its ability to mobilise capital, direct industry, and scale infrastructure with extraordinary speed. At the same time, Zubrzycki highlights that the stand-off between New Delhi and Beijing may be "the most dangerous great-power relationship in the world." This rivalry underscores how different governance models convert genius into industrial capability in different ways — and how geopolitical friction can shape, accelerate, or destabilise national hydrogen strategies.

These examples underscore a critical lesson: abundant talent and technical genius are necessary but not sufficient. Success in hydrogen requires resilience — the ability to absorb shocks without losing momentum; adaptability — the capacity to pivot in response to evolving circumstances; and strategic foresight — the vision to align innovation with policy, investment, and infrastructure over the long term.

Across the globe, the energy transition is already revealing the contours of national genius. Some nations turn talent into industrial scale with remarkable efficiency; others struggle to connect policy, infrastructure, and innovation. The Genius Principle reminds us that success is not determined solely by resources or ambition but by the ability to create, empower, and deploy the people capable of navigating technological and geopolitical complexity.

In the hydrogen era, understanding where and how these geniuses operate is the key to understanding not just who leads—but how leadership itself is defined.

13. Energy Today vs. Energy Tomorrow – Changing Infrastructure

"The future is already here — it's just not evenly distributed."
— William Gibson

"Every system is perfectly designed to get the results it gets."
— W. Edwards Deming

One of the biggest challenges in the hydrogen economy today is not technological—it's conceptual. Too many organisations are still treating hydrogen as if it were a purely chemical business. Historically this makes sense: hydrogen production came out of the chemical industry, and at its core sits a chemical reaction. But when this mindset becomes the primary lens through which companies approach hydrogen, it leads them down the wrong path.

An electrolyser is not simply another piece of chemical-processing kit. It behaves far more like an engine or a generator—taking in one form of energy (electricity) and outputting another (hydrogen). It is a machine, a modular building block of a future energy system. And, as discussed earlier in this book, hydrogen is not merely a fuel but a connector, an energy carrier that binds together sectors, vectors, and time scales.

This is why industry needs to stop looking backwards—to the legacy of chemical plants—and learn instead to look forward: towards hydrogen as an energy system component. The mindset shift is subtle yet profound. Hydrogen should be treated as an integral part of a multi-vector energy architecture, not as a niche chemical product.

To get there, we must begin developing system-level solutions designed specifically for hydrogen. That means assembling capabilities across wind, solar, battery storage, microgrid design, and—eventually—new nuclear technologies such as

SMRs. The future of hydrogen isn't a single discipline; it is the intersection of many.

And as we shift toward these integrated systems, something else becomes clear: hydrogen will blur the traditional boundaries between upstream and downstream energy. For the first time, the same molecule links electricity systems, mobility, industrial heat, long-duration storage, and even backup power. The downstream world will be reshaped just as profoundly as the upstream one, because hydrogen forces sectors that were once separate to interact as a single system.

And when you look at the market today, you see early signals of this shift. Roughly 50% of all hydrogen projects globally are under 50 MW. Energy-wise, this is tiny. Some will inevitably scale up over time, but many will not—and that's the point. A growing number of organisations are not looking at hydrogen through an environmental lens at all; they are looking through an energy security lens. Hospitals, data centres, logistics depots, military facilities—they see hydrogen as a path to produce local, controllable energy at the place they need it most. It is a means not to "go green" but to keep the lights on when everything else goes dark.

These organisations are not only thinking about how hydrogen is produced, but how it is stored, converted, and used downstream. The energy security value of hydrogen lies equally in these downstream functions—compression, storage, fuelling, backup power, microgrids—not merely in the electrolysers themselves. This marks a clear departure from the traditional chemical-plant mindset.

This is where the architecture of the energy system becomes central to understanding what comes next.

Today's grid is built around a hub-and-spoke model: one very large generation plant pushing energy outwards across long distances, distributing it into ever smaller packets as it gets closer to demand. It is, in many ways, like a heart pumping blood.

But the system of tomorrow will look very different. What we are moving toward is what I call a hub-and-island system. We will still have the large, centralised facilities—the big hearts—but we will also have many smaller hearts producing energy locally, capable of operating independently when required. And in today's political climate, you can see immediately why this idea appeals to policymakers, business leaders, and the public. It offers resilience. It offers fallback. It offers control.

Crucially, these "energy islands" are not only generation islands—they are downstream service islands. They generate, distribute, store, convert, and even retail

energy at the point of use. Ports, industrial parks, logistics hubs, airports—many of these locations will become fully fledged energy ecosystems in their own right, performing roles that once belonged only to utilities and midstream operators.

But this shift isn't only about backup power. It is also about transforming business models.

Hydrogen enables companies to become producers of their own energy ecosystems. Imagine a logistics centre with 8 MW of rooftop solar and a mid-sized wind turbine. That onsite renewable energy can produce not only electricity for operations but also hydrogen for vehicle fleets. When the economics align—and they will—this becomes a practical, rational business choice. The line between "energy consumer" and "energy producer" begins to blur.

And this blurring continues downstream. Hydrogen allows companies to procure energy not as a commodity, but as a service. Instead of buying fuel or electricity, they can purchase a fully integrated energy solution—electrolyser capacity, storage, fuelling, and backup generation—as a managed service with fixed uptime guarantees. This is the beginning of Hydrogen-as-a-Service, or more broadly, Energy-as-a-Service (EaaS).

And this shift brings another advantage: cost structure.

Today's power stations are enormous, capital-intensive, high-complexity assets. They require turbines the size of buildings, vast fuel supply lines, extensive land, and sophisticated distributed control systems. But the "islands" of the future will not resemble this. Electrolyser-based systems are smaller, modular, and fed by local renewables. Their control systems can be handled by compact programmable logic controllers rather than sprawling control rooms.

I'm not suggesting these future systems are simple—but they will be cheaper, modular, repeatable, and more accessible than the mega-infrastructure of the past.

This is where modularity becomes transformative. Returning to our 8 MW logistics-centre example: every part of that system—solar generation, wind, electrolyser skids, water treatment, compression, battery storage, even hydrogen storage—can be delivered as standardised, plug-and-play elements. With the Modular Type Packaging (MTP) concept, these systems can connect seamlessly. Equipment manufacturers can deliver their units as black boxes, and EPCs can assemble them far more easily than the bespoke plants of today.

This modularity extends downstream too. Hydrogen refuelling stations, compression units, storage pods, ammonia crackers, and MW-scale fuel-cell generators are increasingly available as containerised blocks. Downstream infrastructure, which traditionally required custom-built solutions, is moving toward a "Lego-like" assembly model—accelerating deployment while reducing engineering cost.

One overlooked benefit of hydrogen in this downstream context is its ability to carry time. Batteries store energy for hours; hydrogen can store it for days, weeks, or seasons. This gives downstream users—ports, industrial clusters, data centres—the ability to stabilise their own micro-systems against intermittency or price volatility. Hydrogen becomes not just an energy vector, but a temporal buffer, shifting surplus energy from when it is generated to when it is truly needed.

As these downstream roles expand, regulation must adapt as well. You cannot apply petrochemical safety frameworks to a hydrogen-powered campus, hospital, or neighbourhood-scale microgrid. New permitting regimes, operational standards, and digital oversight structures will be needed for this distributed hydrogen future to function safely and efficiently.

In practice, this misalignment is already visible. In several recent small-scale hydrogen projects — modular installations well below 20 MW — regulators have required the same distributed control system (DCS) architectures used in large petrochemical complexes. On paper, this may appear prudent. In reality, it imposes the cost structure and complexity of a refinery onto what is essentially a modular, PLC-driven energy unit.

The result is not increased safety — it is distorted economics.

A four-electrolyser site serving a local energy island does not behave like a sprawling chemical plant with hundreds of process interdependencies. Treating it as such reflects a regulatory framework designed for yesterday's infrastructure, not tomorrow's.

This is precisely the challenge of Energy Tomorrow: energy systems are moving closer to end users, becoming smaller, modular, and distributed. Yet governance structures often remain calibrated for centralised, heavy-industrial assets. Until this gap is addressed, innovation will repeatedly collide with inherited assumptions.

None of this eliminates the need for large-scale hydrogen projects—far from it. About half of all hydrogen developments remain above 50 MW, and many gigawatt-scale installations will continue to drive industrial decarbonisation and

global trade. But the future is not one or the other—it is both. A dual system: centralised hubs and decentralised islands, coexisting and reinforcing each other.

And right now, industry is not fully preparing for this mix.

This chapter is about exactly that transformation: how energy infrastructure—physical, digital, and conceptual—will evolve as we move from the world of Energy Today to the decentralised, multi-vector, resilient world of Energy Tomorrow.

The System We Inherited: How Today's Energy Architecture Works

The energy system we rely on today is the product of more than a century of engineering choices shaped around fossil fuels. It emerged gradually, layer by layer, as coal, oil, and gas became the dominant forces behind industrialisation, electrification, and global trade. Over time, these fuels shaped not only our technologies, but the very architecture of the system itself. As fossil fuels are energy-dense, easily stored, and relatively straightforward to transport, the most efficient way to use them was to concentrate production in a small number of large facilities and then distribute the resulting electricity, heat, or refined products outwards to consumers.

One aspect of traditional energy infrastructure that is often overlooked is the extent to which it was shaped by the variability of the raw materials themselves. We often speak simply of "oil" or "natural gas", as though each were a uniform product. In reality, they are anything but. North Sea crude differs significantly from West Texas Intermediate; both differ again from the light "sweet" crudes of the Middle East or the heavy oils of Venezuela and Canada. Each possesses different physical and chemical characteristics that influence extraction, transport, storage, refining, and ultimately the design of the facilities built to process them.

As a result, much of the bespoke nature of the oil and gas industry was not a matter of engineering preference but of engineering necessity. Every field, every refinery, and every processing facility had to be designed around the specific characteristics of the feedstock available. Engineers were adapting their infrastructure to the geology beneath their feet.

Green hydrogen presents a fundamentally different proposition. Rather than extracting a naturally variable resource, electrolysis manufactures the same basic molecule every time. Regardless of whether the electricity originates from offshore wind, solar PV, hydroelectricity or nuclear power, the hydrogen leaving the electrolyser is the same elemental building block. For the first time, the industry

begins with a manufactured feedstock whose composition is defined by engineering rather than geology.

That seemingly simple distinction has profound implications. If the feedstock itself is consistent, much of the surrounding infrastructure can become more consistent as well. Electrolysers, compression systems, storage, refuelling equipment and many downstream processes become increasingly amenable to modular design, repeatable engineering and standardised interfaces. Even where hydrogen is subsequently converted into ammonia, methanol or synthetic fuels, those conversion processes begin from the same well-defined foundation rather than from hydrocarbons whose composition varies from field to field.

This does not eliminate engineering complexity—hydrogen systems remain sophisticated pieces of infrastructure—but it shifts where that complexity resides. Instead of continually adapting equipment to the variability of nature, engineers can increasingly focus on optimising standard building blocks within larger energy systems. The opportunity is therefore not simply to build cleaner infrastructure, but to industrialise its delivery through modular, repeatable building blocks in ways that were rarely possible in an industry continually adapting to the variability of nature.

Understanding this difference also helps explain why the inherited energy system evolved into the highly centralised architecture that still dominates today.

The logic that defines this system is unmistakably centralised. Gigawatt-scale power plants, massive refineries, LNG terminals, and national transmission networks all follow the same organising principle: a few large production assets generate power or process fuels, and these outputs are pushed along carefully engineered pathways—pipelines, transmission lines, rail corridors, tanker routes—to millions of passive end users. Energy flows in one direction. The assets at the centre drive everything around them.

For decades, this architecture served us exceptionally well. By concentrating production, countries achieved remarkable economies of scale. A small number of professionally operated plants could produce electricity or refined fuels at low cost, under high levels of control, and with predictable performance. Centralisation reduced operational complexity: system operators could dispatch large volumes of power from a handful of controllable generators, confident that demand would follow stable daily and seasonal patterns. The entire structure was optimised for reliability, affordability, and the regulated, top-down coordination that defined 20th-century energy planning.

But embedded in this success were certain assumptions—assumptions about the fuels, the technologies, and the behaviour of demand. Fossil fuels were inherently stable and controllable, which meant the system could be rigid without consequence. Operators could turn generation up or down as needed. Fuels could be stockpiled for weeks or months. Pipelines and grids could operate with predictable flows because the inputs never fluctuated. The physical properties of the fuels allowed the architecture to remain linear, hierarchical, and centralised.

Yet this very structure also made the system rigid. It depends on large, centralised assets operating continuously and predictably. It depends on long-distance networks functioning flawlessly. It also depends on a passive consumer base that takes whatever energy is delivered, whenever it is delivered. The system works brilliantly under stable conditions, but its rigidity becomes a weakness when confronted with volatility, decentralised generation, or the need for rapid structural adaptation.

The system we inherited was an elegant solution to the challenges of the fossil-fuel era. But it was designed for a world that no longer exists. It solved the problems of the past, not the demands of the future. The assumptions that once made this architecture the pinnacle of efficiency are the same assumptions that now constrain its ability to evolve.

The Assumptions That Built It — and Why They No Longer Hold

The architecture of the modern energy system did not emerge from conscious design so much as from a set of implicit assumptions that shaped every technical and economic decision for more than a century. These assumptions were so deeply embedded in our thinking that they barely needed to be stated. They were the background logic that made the hub-and-spoke model seem inevitable, sensible, and efficient. But they were also the invisible scaffolding that now limits the system's ability to evolve.

The first assumption was that energy would always be dispatchable. In the fossil-fuel paradigm, energy came from fuels that could be burned whenever needed. If demand increased, you burned more coal or opened the gas valve a little wider. If demand fell, you dialled it down. Flexibility was built into the molecules themselves. Dispatchability was not a feature of the system; it was a property of the fuels. This meant the entire architecture—from grid operations to market design—was built on the premise that supply would effortlessly follow demand.

The second assumption was that energy carriers would always be easy to store and cheap to transport. Coal could sit in a pile for months with no degradation. Oil could be stored indefinitely in tanks or moved across oceans at a low cost per unit of energy. Gas could be compressed or liquefied and sent through pipelines that operated with little variability. Everything about fossil fuels favoured centralised production and long-distance distribution. Transporting energy was inexpensive relative to generating it, so building ever-larger plants made sense: the bigger the plant, the lower the unit cost, and the infrastructure to move the product was already in place.

The third assumption—rarely acknowledged but absolutely foundational—was that carbon had no cost. Emissions were externalities in the true economic sense of the word: consequences that did not appear on the balance sheet. Since carbon pricing did not exist and environmental constraints were minimal, energy planning could focus exclusively on cost, reliability, and scale. The environmental footprint of the system simply did not shape its architecture.

For decades, these assumptions held true. They made the system stable, predictable, and economically aligned with the priorities of industrial growth. They also justified the design choices that prioritised centralisation, baseload generation, and large capital-intensive assets. The hub-and-spoke architecture was, in every sense, a rational response to the world as it was.

But today, each of those foundational assumptions is breaking down.

Dispatchability is no longer guaranteed. The growth of wind and solar—now the fastest-growing sources of new energy capacity—means that supply is increasingly variable. Instead of generation following demand, demand now must adapt to supply. The system must respond to weather patterns, seasonal variability, and local intermittency. The architecture, which once relied on operators simply turning dials, must now rely on flexibility, storage, and digital coordination.

Storage and transport can no longer be taken for granted. Electricity is difficult and expensive to store over long durations. Hydrogen is more complex to move than methane. Renewable electricity is often cheapest far from demand centres but difficult to transport without major grid reinforcement. The assumption that you can always move energy from where it is produced to where it is consumed no longer holds with the same ease or cost. A related misconception is that renewable resources automatically translate into usable energy.

In reality, generation and connectivity are increasingly becoming separate challenges. Some regions possess abundant renewable potential but lack the trans-

mission capacity needed to move that energy to demand centres. During a visit to Canada, I was struck by discussions around Quebec. On paper, the province possesses vast quantities of low-carbon hydroelectric power. Yet for industrial developers seeking new projects, access to that power is often constrained not by generation itself but by the availability of grid connections and transmission infrastructure. The lesson is an important one. Having energy available somewhere in the system is not the same as being able to access it where and when it is needed. As renewable deployment accelerates, grid capacity may increasingly become a more valuable resource than generation capacity itself. In a multi-vector system, moving energy becomes both a technical challenge and an economic decision.

Increasingly, the challenge is no longer simply one of building additional generation, but of extracting greater value from the infrastructure already in place. Modern grid operators are deploying digital substations, dynamic line rating, advanced distribution management systems, digital twins, AI-assisted forecasting and increasingly sophisticated optimisation tools to unlock capacity that previously remained inaccessible. In many cases, significant improvements can be achieved without immediately constructing new transmission lines. The future electricity grid will therefore not simply be larger than today's—it will also be significantly smarter, with digital intelligence becoming as important as physical infrastructure itself.

And carbon now has a real cost—financial, political, and societal. Emissions pricing, environmental regulation, and corporate commitments mean that carbon cannot be treated as a free by-product of growth. The system must optimise not only for affordability and reliability, but also for decarbonisation. This changes what “efficiency” means. A system designed to minimise cost per megawatt-hour is not the same as a system designed to minimise emissions, maximise resilience, and integrate diverse energy carriers.

These shifts do not just challenge the old assumptions—they invalidate them. The energy system is being asked to operate in a fundamentally different environment, and the architecture built for the fossil era cannot simply stretch to accommodate the new realities. The assumptions that once made the system robust now make it brittle.

In their place, new assumptions are emerging:

- that variability is intrinsic, not a deviation;
- that storage is fundamental, not optional;
- that energy will often be produced close to where it is needed;

- that carbon must be minimised or captured;
- and that resilience requires diversity, not just scale.

The task now is to redesign the energy system for these new assumptions—because the old architecture, no matter how successful it was in its time, cannot deliver what the future requires.

The Pressures Reshaping the System

While it is tempting to attribute the transformation of the world's energy architecture solely to climate commitments or net-zero targets, the truth is far more complex. The traditional system is not being disrupted by a single force, but by the convergence of multiple pressures—technological, economic, political, and physical. Each of these pressures might be manageable on its own, but together they stretch the old design past its limits. The hub-and-spoke model, once celebrated for its simplicity and predictability, finds itself increasingly mismatched to the world it must now serve.

The first major pressure is the rise of renewable electricity, which fundamentally alters the behaviour of the energy system. Fossil plants provided dispatchable output, following demand with remarkable precision. Wind and solar, by contrast, follow the weather. Their variability is not a flaw—it is an inherent property of the resource. As renewables shift from being supplementary to becoming central, their intermittency begins to dictate system behaviour. The entire architecture must adapt to supply that ebbs and flows rather than obediently meeting demand.

The second pressure is electrification, which changes not only how much energy we use, but where and when we use it. Electric vehicles, heat pumps, electrolysers, and digital infrastructure all reshape demand curves, producing sharp peaks, deeper troughs, and new geographical concentrations of load. Critically, electrification also pushes energy consumption closer to the consumer and closer to the end-use equipment. Industrial users increasingly want the ability to produce and manage energy on-site—a profound shift from the passive consumption model of the past.

Linked to this is a third, rapidly intensifying pressure: the growing emphasis on energy security and resilience. In a world marked by geopolitical tension, supply chain risk, and infrastructure vulnerabilities, many organisations—from hospitals and data centres to logistics hubs and military facilities—are no longer willing to depend solely on a distant central grid. They are seeking ways to secure their own operations through on-site generation, local storage, and autonomous energy systems. Hydrogen plays a pivotal role in this trend, enabling long-duration stor-

age and multi-vector flexibility that traditional backup systems cannot match. The result is a move toward producing energy where it is needed, not only where it is cheapest.

A fourth pressure comes from hydrogen and synthetic molecules, which introduce entirely new energy carriers with different physics and economics. Hydrogen's high transport cost and conversion losses encourage siting production close to the off-taker. In many cases, it becomes more rational to place electrolysers beside industrial facilities, fuelling depots, ports, and logistics centres rather than in remote regions, unless the hydrogen is converted into a more transportable energy carrier, such as ammonia or methanol, where the economics of long-distance shipping become far more favourable. This marks a stark departure from the old model, where fuels were produced centrally and shipped across long distances. Hydrogen nudges the system toward distributed production by default.

A fifth pressure is digitalisation, which converts the energy system from a mechanical machine into a digital organism. Energy flows must now be forecasted, optimised, and controlled in real time across thousands of distributed assets. This introduces powerful new capabilities—but also new demands. The old top-down control model cannot cope with a system where generation and consumption happen at the edges as well as the centre.

Finally, the sixth pressure—geopolitical volatility—has forced governments and industries to rethink the fragility of global energy supply chains. Hydrogen hubs, microgrids, and distributed storage assets are no longer viewed only as decarbonisation tools; they are now strategic resilience assets. The energy island is not merely a technological vision—it is becoming a policy objective.

What unifies these pressures is that they each challenge a different foundational assumption of the 20th-century system. Renewables challenge dispatchability. Electrification challenges predictable demand. Distributed hydrogen challenges the logic of centralised production. Digitalisation challenges hierarchical control. Geopolitical risk challenges the assumption of stable global supply chains, while energy security challenges the very idea that consumers will remain passive.

Individually, these pressures strain the old architecture; together, they dissolve it.

The system cannot simply be patched. It must be re-architected to reflect a world where energy is produced closer to where it is consumed, where resilience is as important as efficiency, and where multi-vector flexibility replaces single-vector

simplicity. These pressures point inexorably toward a new design paradigm—one built not around a single hub, but around a network of interconnected energy islands.

This shift toward system-level thinking is not merely theoretical. In practice, it is already reflected in the way modern energy systems are analysed. Integrated modelling frameworks such as the ETHOS model suite are designed to explore how electricity supply, hydrogen production, storage, infrastructure constraints, and patterns of demand interact over time. Rather than treating hydrogen as an isolated technology with a single cost curve, ETHOS places it within a wider energy system where outcomes emerge from interactions between assets. What such models consistently reveal is that performance is rarely determined by one dominant variable. Changes in utilisation, grid access, demand timing, or system integration often have a greater impact on results than incremental improvements in component efficiency. The value of this approach is not prediction to the last decimal place, but clarity about structure: it shows that energy systems deliver the outcomes they are designed for, and that redesigning relationships between assets often matters more than optimising the assets themselves.

From Hub-and-Spoke to Hub-and-Island

The hub-and-spoke model dominated the 20th century because it perfectly matched the physics and economics of fossil fuels. Large, centralised plants produced vast quantities of energy at low marginal cost, and the system's job was simply to deliver that energy—through pipelines, wires, and shipping routes—to wherever it was needed. Everything radiated outward from a handful of major hubs. Consumers were passive endpoints. Production was separate from use and energy security was equated with having sufficiently robust supply chains feeding into these central hubs.

However, the forces reshaping today's energy system have eroded the foundations of that logic. Renewables decentralise generation; hydrogen decentralises storage; electrification decentralises demand. The system is no longer defined by a single centre of gravity. Instead, multiple centres of gravity are emerging—each shaped by local resources, local consumption, local constraints, and local strategic priorities. This requires a fundamentally different architecture.

The alternative is not to abandon hubs, but to multiply them and connect them differently. This is the essence of the hub-and-island model: a system in which large, centralised assets continue to play major roles, but are complemented by a growing constellation of smaller, distributed, semi-autonomous "islands" of

energy production and storage. These islands can draw from the grid when it is available, support the grid when it is stressed, and operate independently when the situation demands it.

The shift is not ideological—it is a response to physics, economics, and geopolitics.

Renewables distribute generation because the resource is distributed. You cannot move the sun or wind to a convenient location; you must go to where they naturally exist. Hydrogen, owing to its transport costs and conversion losses, pushes production closer to the off-taker unless converted to ammonia, methanol, or another energy-dense carrier. Electrification pushes consumption into new corners of the system—charging depots, data centres, heat pumps, electrolyser clusters. And energy security pushes organisations to take more control of their own energy destinies.

An energy island can take many forms. It might be a logistics centre with solar, batteries, and an electrolyser producing hydrogen for its fleet. It might be a port that handles ammonia imports, cracks hydrogen on-site, supplies bunkering fuel, and feeds power to vessels at berth. It might be a hospital, a military base, a semiconductor fab, or a data centre—each integrating multiple vectors to ensure that operations can continue under any circumstances. It might be an industrial park where electricity, hydrogen, heat, and CO_2 circulate in closed loops to maximise efficiency and minimise exposure to external disruptions.

What these islands share is functional autonomy. They are not off-grid communities in the romantic sense; they are industrial-strength micro-systems capable of running in grid-connected mode or in island mode depending on need. They reduce dependence on distant infrastructure. They reduce vulnerability to upstream disruptions and increase the resilience of the wider system by providing distributed flexibility and buffering capacity.

In the hub-and-island architecture, the directionality of energy flow changes. Instead of energy being pushed outward from a centre, energy moves in multiple directions between nodes. Some islands will primarily consume energy, others will produce it, and many will do both at different moments. Hydrogen will move between islands as a store of time; electrons will move quickly to balance instantaneous needs; heat may circulate locally within industrial clusters; CO_2 may flow backward toward sequestration hubs. The system behaves less like a tree and more like a mesh network.

This new architecture also changes the role of central hubs. Instead of being the single point of control, hubs become stabilisers and coordinators. Large offshore wind farms, SMRs, or gigawatt-scale hydrogen facilities will still anchor the system, but they will interface with dozens or hundreds of distributed islands rather than with millions of passive consumers. The relationship becomes interactive, not hierarchical.

The shift to hub-and-island is inevitable because it is the only architecture that can reconcile the forces of decentralisation with the need for system-level coordination. It preserves the strengths of large-scale assets while incorporating the resilience, flexibility, and localisation required by the emerging energy landscape. It transforms the consumer from a passive recipient into an active participant. It also aligns the energy system with a world defined increasingly by volatility, decentralisation, and the desire for autonomy.

In short, hub-and-island is not a theoretical model. It is the blueprint for an energy system where hydrogen, renewables, storage, and digitalisation work together—not as isolated technologies, but as interlocking components of a dynamic, multi-vector, resilient infrastructure.

Energy Islands: What They Are and Why They Matter

Before we think of energy islands as futuristic concepts, it is worth recognising that some industries have been forced to operate this way for decades. Early in my oil and gas career, I was in a meeting with a Chevron manager discussing operations in a country that will remain unnamed. At one point he leaned back in his chair and said, almost casually, "we actually have more power-generating capacity in that country than the national power provider." They had built their own parallel energy system—generation, backup, distribution—simply because the national grid could not provide the reliability that their operations required.

What was once an extreme necessity for oil and gas operators is now a preview of the future for energy-intensive industries everywhere. The logic behind that Chevron system—self-sufficiency, resilience, and autonomy—is the same logic that underpins the emerging architecture of energy islands.

Energy islands are not geographical curiosities, nor are they miniature versions of the old power system. They represent a fundamentally different way of organising energy: modular, multi-vector ecosystems capable of generating, storing, converting, and consuming energy locally while still participating in the wider network.

At the heart of an energy island is functional autonomy. Unlike traditional consumers that rely entirely on distant hubs, an island can meet some or all of its own energy needs on-site. It may still draw power from the grid under normal conditions, but it is not dependent on it. When the grid is stressed, unstable, or unavailable, the island can continue operating by drawing on its own generation and storage. This ability to "island" when necessary is more than a technical capability—it is a strategic asset in a world defined by resource volatility and geopolitical uncertainty.

What makes energy islands compelling is not a single technology but the integration of multiple technologies into a cohesive system. A typical energy island may include:

- Renewable generation (solar PV, wind turbines, waste heat recovery, geothermal)
- Short duration storage (batteries for minutes to hours)
- Long-duration storage (hydrogen tanks, ammonia tanks, underground caverns)
- Hydrogen production and conversion (electrolysers, fuel cells, ammonia crackers)
- Flexible loads (EV fleets, data centres, industrial processes)
- Digital control systems (AI forecasting, automation, microgrid controllers)

Each island becomes a microcosm of the future energy system: electrons, molecules, heat, and data moving in synchrony to deliver reliability, resilience, and economic value.

Crucially, energy islands are emerging because they solve real-world problems that the central grid cannot easily address. For industries with continuous operations—semiconductor fabs, chemical plants, food processing facilities, hospitals, airports—the cost of energy disruption can dwarf the cost of energy itself. These organisations cannot afford to tolerate blackout risk, voltage instability, or upstream supply disruptions. Hydrogen-enabled islands provide long-duration backup and operational continuity in a way that diesel generators or batteries alone cannot.

Energy islands also respond to a growing desire for energy sovereignty at the local and organisational level. After years of geopolitical volatility, supply chain fragility, and price shocks, governments and companies are positioning themselves to take more control of their energy fate. A logistics hub with its own solar fields, battery storage, and hydrogen system is far less vulnerable to external shocks. A

port with ammonia imports capability and hydrogen bunkering infrastructure becomes a strategic node in global energy trade. A military base powered by an integrated microgrid gains resilience against both physical and digital threats. In each case, the island provides a level of operational independence the old system rarely allowed.

Equally important is the role of islands in enabling off-taker proximity, one of the defining characteristics of the hydrogen economy. As hydrogen is expensive to move in its pure form, energy islands place production and consumption close together. Electrolysers sit beside industrial furnaces, refuelling stations beside depots, ammonia crackers beside bunkering facilities. This dramatically reduces infrastructure needs and unlocks use cases that are uneconomical in a centralised model.

Energy islands are also a practical response to the modularity of modern energy technologies. Electrolysers come in prefabricated skids; refuelling stations arrive in containers; battery banks are plug-and-play; fuel cells can be added as drop-in units. These components integrate far more easily at the local level than in the large, bespoke infrastructure projects of the past. Islands are therefore not exceptions—they are the natural way the energy transition scales.

And they are not all the same. Energy islands differ depending on context:

- Industrial islands optimise heat, power, hydrogen, and carbon loops.
- Port islands handle multimodal energy vectors for shipping, bunkering, and import/export.
- Campus islands (universities, hospitals, city districts) prioritise resilience and cost control.
- Logistics and fleet islands integrate renewable production with refuelling networks.
- Remote or off-grid islands use hydrogen to store renewable energy that would otherwise be stranded.

Each type of island emerges from a specific combination of needs, constraints, and opportunities. Yet all share a common design principle: local optimisation feeding into a global system, rather than global optimisation imposed uniformly on every user.

The importance of energy islands cannot be overstated. They are:

- resilience mechanisms for critical industries,
- flexibility providers for the wider grid,
- decarbonisation platforms for hard-to-abate sectors,

- innovation hubs for hydrogen and synthetic fuels,
- and strategic assets in national energy planning.

They are not a fringe idea—they are the building blocks of the emerging energy system. If the fossil-fuel era was defined by massive central plants and long-distance pipelines, the hydrogen era will be defined by interconnected islands, each self-sufficient yet networked, each capable of supporting the system and being supported by it.

Energy islands matter because they reflect the world as it is becoming, specifically distributed, digital, multi-vector, and strategically uncertain. They turn energy from something delivered to something shaped, managed, and secured at the point of need. In doing so, they unlock a resilience and adaptability that the old architecture simply could not provide.

The New Downstream: Storage, Conversion, and Energy-as-a-Service

The downstream side of the energy system has traditionally been the simplest to describe and the least complex to manage. Fuels arrived from far away, grids delivered electricity from distant plants, and consumers used what they were given. Downstream was a one-way street. It involved logistics, metering, and billing—not system design.

Hydrogen changes that completely.

In a hydrogen-enabled world, the downstream becomes the most dynamic part of the entire energy ecosystem. It is where generation, storage, and conversion meet actual use cases. It is where energy is reshaped into the form that specific users need, in the moment they need it. It is also where the biggest operational and commercial transformations will take place.

At the centre of this shift is storage—particularly long-duration storage. Batteries can smooth fluctuations over seconds or hours, but hydrogen introduces temporal reach the fossil system never needed. A plant, a port, or a data centre can now stabilise its own energy environment not just for an afternoon, but potentially for an entire season. That capability turns energy from a commodity into a tool of operational control.

But hydrogen's impact goes further. It introduces conversion as a normal downstream activity rather than an upstream specialty reserved for refineries and chemical plants. Electricity becomes hydrogen; hydrogen becomes electricity;

ammonia becomes hydrogen; hydrogen becomes heat or mobility fuel. Downstream no longer means "fuel at the end of the pipe." It becomes a place where energy vectors can change shape depending on the use case.

This is where a crucial part of the energy transition often gets overlooked: the enormous and rapidly growing market for temporary and mobile power. For years, diesel generators have dominated this space—construction sites, festivals, mines, remote operations, disaster recovery, military bases, temporary data installations. This market remains huge, and despite the sharp rise of batteries and renewables, it is nowhere near being displaced. Not everything can be electrified by a battery pack dragged onto a site. Many applications require multi-day resilience, high power density, or fuels capable of long storage and rapid deployment.

This is precisely where hydrogen, fuel cells, methanol systems, and hybrid battery-hydrogen units begin to play. The temporary and mobile power sector is a downstream frontier waiting to be transformed. It is a perfect example of "form follows use case": sometimes batteries will be ideal, sometimes hydrogen, sometimes methanol and sometimes a hybrid of all three. The point is that downstream users gain choice, whereas the fossil system forced homogeneity. Instead of a one-size-fits-all diesel genset, future mobile power becomes modular, clean, silent, and configurable.

This downstream shift is being accelerated by modularity. Electrolysers arrive on skids. Hydrogen storage and compression units can be dropped onto concrete pads. Fuel cells can be containerised and delivered as turnkey assets. Battery banks click into place digitally through standardised interfaces like MTP. These technologies aren't waiting for multi-year centralised planning cycles—they can be deployed quickly, expanded easily, and integrated into existing operations.

As the downstream becomes more complex, a new service model gains traction: Energy-as-a-Service. Most companies do not want to own or maintain their own electrolysers, hydrogen buffers, ammonia crackers, battery units, or predictive microgrid controllers. They want guaranteed uptime, predictable energy cost, stable backup, and resilience against volatility or blackout. The complexity is outsourced, while the performance is internalised. Energy stops being something purchased purely as a commodity and becomes something purchased as an outcome.

This shift is profound. It means that downstream users—industrial clusters, ports, data centres, logistics hubs, remote sites—no longer wait for the central grid to modernise. They can modernise themselves. They can generate, store, and con-

vert energy on site; they can hedge against outages and price spikes; and they can build resilience that is tailored to their specific operating reality.

The downstream, once passive, becomes the centre of agency in the energy system. It is where new models of resilience emerge. It is where hydrogen becomes not just a fuel but a strategic asset. It is also where the architecture of Energy Tomorrow becomes visible first: decentralised, modular, flexible, and deeply attuned to the use case rather than the assumptions of yesterday's system.

Infrastructure Reimagined

The shift from a fossil-based system to a renewable, hydrogen-enabled one does more than change how we make energy—it changes the shape of the infrastructure that supports it. The pipes, ports, storage sites, and control systems that we inherited were built for a world of stable fuels, predictable flows, and centralised production. But the emerging system behaves differently. It is multi-vector rather than single vector, distributed rather than purely centralised, and digitally coordinated rather than mechanically scheduled.

As the logic of the system changes, the physical backbone must change with it.

In the old model, pipelines carried one molecule in one direction, ports handled a narrow range of fuels, storage was peripheral, and digital systems played a supporting role. In the new model, each of those components becomes more versatile. Pipelines must adapt to new molecules and new flow patterns. Ports must evolve into gateways for hydrogen carriers, synthetic fuels, and CO_2. Storage becomes a structural pillar rather than a contingency, and digital systems shift from optional to essential as the coordinating intelligence of a far more dynamic network.

This transformation is not theoretical—it is already taking shape. Across the world, infrastructure developers, port authorities, pipeline operators, industrial clusters, and energy companies are beginning to reconfigure assets to handle hydrogen, ammonia, methanol, CO_2, synthetic methane, and multi-vector power flows. Some projects repurpose existing infrastructure; others involve entirely new builds. All reflect the same underlying reality: the energy system is becoming more flexible, more integrated, and more modular.

The sections that follow examine how this plays out in practice.

Not conceptually, but concretely.

- What does a hydrogen-ready pipeline network look like?

- How are ports redesigning themselves to handle hydrogen carriers and new bunkering fuels?
- Where are hydrogen storage caverns, ammonia terminals, and synthetic fuel hubs already being built?
- How are digital platforms turning complex multi-vector systems into something operable in real time?

This is where Energy Tomorrow becomes tangible—not as policy or theory, but as steel, concrete, software, and real projects moving into operation.

From Architecture to Reality — What Is Being Built Today

The ideas in the previous sections only matter if they materialise in the real world. Across the globe, this is already happening. Pipelines are being repurposed or newly routed for hydrogen and CO_2. Ports are redesigning themselves to handle new energy carriers. Caverns and terminals are being converted into long-duration storage sites. Digital platforms are emerging to coordinate multi-vector systems. Entirely new forms of infrastructure — refuelling depots, hydrogen logistics hubs, mobile power systems — are appearing in places the old energy system never reached.

What makes this moment unique is that these developments are not theoretical or speculative. They are concrete projects, funded and under construction, and together they offer a preview of the energy landscape that will define the next several decades.

For each major part of the future system — pipelines, ports, storage, digital coordination, and distributed infrastructure — we look at the most significant global developments and then take a closer look at one or two projects that illustrate the direction of travel. The aim is simple: to show where the transition is already taking shape, and how these early choices will influence the energy system that we ultimately end up with.

Pipelines: From Single-Purpose to Multi-Vector

Pipelines have long been the quiet infrastructure of the modern world—steel arteries transporting hydrocarbons across continents with remarkable reliability. But they were engineered for a simpler paradigm. Traditional networks carried a single commodity, usually methane, in a single direction, under stable and predictable conditions. Their purpose was narrow, and their behaviour well understood.

In the emerging energy system, that simplicity disappears. Hydrogen, CO_2, synthetic methane, and district heat are beginning to play interconnected roles within industrial clusters and regional networks. Pipelines cease to be one-purpose conduits and instead become part of a multi-vector circulatory system. Existing gas corridors are being examined for hydrogen repurposing; new routes are forming to connect production sites with storage caverns and demand centres; CO_2 capture hubs increasingly require dedicated transport routes to utilisation or sequestration sites. In some regions, heat networks linking industrial sites and urban districts are also being explored.

This transformation is more than a change of materials or pressure regimes—it represents a redefinition of what a pipeline is. Instead of serving a single fuel moving in one direction, pipeline infrastructure is evolving into a flexible backbone capable of supporting multiple molecules and flows across the energy ecosystem.

This shift is already visible in projects taking shape across Europe, North America, the Middle East, and East Asia. Operators are beginning to design pipeline networks that look very different from the natural gas grids of the 20th century: systems capable of transporting hydrogen, CO_2, synthetic methane, and even heat, often with bidirectional flows linking industrial clusters that once functioned in isolation. Nowhere is this transformation more advanced than in Europe.

Deep Dive 1: The European Hydrogen Backbone — From Vision to Reality

When the first outline of the European Hydrogen Backbone (EHB) appeared in 2020, many observers assumed it was a conceptual sketch — a high-level vision of what a hydrogen-ready Europe might look like. Yet within just a few years, it hardened into a coordinated infrastructure plan spanning more than thirty gas transmission operators across nearly every major EU economy. The Backbone now charts a network of over 40,000 kilometres of hydrogen pipelines targeted for 2040, with the majority expected to be repurposed from existing natural gas infrastructure. That detail alone is transformative. Repurposing a gas line for hydrogen is typically a fraction of the cost of a greenfield build, and far faster to deliver. Europe has discovered that building a new energy system does not mean discarding the old one — it means reshaping its skeleton.

What makes the EHB so consequential is not its scale but its logic. It connects the resource basins of the North Sea — home to vast offshore wind plans — with industrial clusters in Germany, the Netherlands, Belgium, France, and Northern

Italy. It ties storage caverns in the Netherlands and Germany to emerging import terminals along the coast. It threads southern corridors into the network to allow hydrogen from North Africa to reach central European demand. Energy flows that for decades ran from east to west begin, for the first time, to run north-south, west-east, and cross-border in entirely new configurations.

The Backbone expresses a simple truth: Europe's heavy industry will need hydrogen at scale, and the only way to supply it reliably is through a continental pipeline architecture that mirrors, but does not simply replicate, the natural gas grids that preceded it.

During discussions at the Hydrogen Dialogue in Munich, one analogy was used that I felt captured the current stage of hydrogen infrastructure development particularly well. The emerging European Hydrogen Core Network was compared to the motorway network of a country. It is an analogy that is both simple and remarkably accurate.

The motorway network forms the primary transport arteries of a transport system, connecting major cities, industrial centres and ports. However, motorways alone do not deliver goods to factories, warehouses or individual businesses. That requires a second layer of infrastructure: regional roads, industrial access routes and local streets that complete the journey.

The same principle applies to hydrogen. The Hydrogen Core Network currently under development represents the primary transmission system, enabling large volumes of hydrogen to move efficiently between production centres, storage facilities, import terminals and industrial clusters. However, like any transport network, it cannot fulfil its purpose on its own. Much of the regional distribution network still remains to be built before hydrogen can reach many of the businesses that will ultimately depend upon it.

This should not be viewed as a criticism of the Hydrogen Core Network. Every infrastructure network develops in phases, with the primary arteries constructed before the secondary distribution systems. Rather, it is a reminder that today's projects represent the beginning of Europe's hydrogen transport system, not its completion. The backbone is beginning to take shape, but many of the roads that will ultimately connect hydrogen to the factories, businesses and communities that use it are still to come.

As important as this reconfiguration is, another project offers a glimpse into how Europe's hydrogen geography is being rewritten even more dramatically — not by repurposing old corridors, but by drawing entirely new ones across the sea.

Deep Dive 2: The Scotland–Germany Interconnector — A New North Sea Energy Axis

Among the emerging hydrogen corridors in Europe, the proposed Scotland–Germany interconnector is perhaps the most strategically symbolic. It represents more than a technical project; it signals a realignment of Europe's energy relationships and a rebalancing of where industrial power may lie in a decarbonised world.

Scotland is endowed with extraordinary renewable resources. Its offshore wind potential exceeds domestic electricity demand many times over, offering a scale of generation that simply cannot be absorbed by its own grid. Germany, by contrast, is a land of industry: steelworks, chemical parks, automotive plants, and heavy manufacturing centres that collectively form one of the largest industrial hydrogen markets in the world. Yet Germany's own renewable land availability is limited, and its decarbonisation deadlines unusually tight.

The interconnector between the two countries is, in essence, an energy bridge across the North Sea. It envisions hydrogen produced in Scottish coastal regions — initially from offshore wind, later potentially from floating wind or other marine renewables — transported directly into Germany's industrial heartlands. In doing so, it replaces the fossil-era logic of continental pipelines with a new kind of relationship: a clean-energy corridor between two advanced economies linked less by geography than by complementary strengths.

This project shifts Europe's energy map in subtle but powerful ways. For half a century, Germany relied on pipeline gas arriving overland from Russia and Norway. The interconnector replaces that dependency with a maritime corridor from a democratic neighbour whose energy resources are expanding, not depleting. It interprets Scotland's wind-rich coastline not as a peripheral region but as a strategic export platform — echoing, in a very modern form, the transformation triggered by North Sea oil in the late 20th century.

The project also fits neatly into the broader backbone strategy. Hydrogen arriving in northern Germany can integrate seamlessly into the continental network, flowing south into Bavaria's steel sector, west into the Netherlands' ports, or east into the growing industrial clusters of Central Europe. It allows Europe to diversify its hydrogen supply not only in terms of geography, but also in terms of energy vectors, blending domestic production, maritime imports, and cross-border pipelines into a cohesive system.

What makes the Scotland–Germany corridor so compelling is that it demonstrates how hydrogen infrastructure is no longer merely a technical enterprise. It is geopolitical architecture. It changes trade patterns. It affects industrial competitiveness. It anchors long-term relationships between exporting and importing regions. In short, it is energy strategy embodied in steel.

Other Emerging Pipeline Corridors Around the World

Europe may be the most advanced, but the pipeline transition is not confined to its borders. Germany's North–South Hydrogen Spine is reshaping internal flows by connecting wind-rich northern ports with the energy-intensive industries of the south. The Netherlands' HyWay27 represents one of the earliest and most systematic efforts to repurpose existing gas infrastructure. In the United States, hydrogen hubs are beginning to cluster into regional corridors that hint at an eventual backbone. In the Middle East, projects in NEOM, Oman, and the UAE are developing internal hydrogen pipeline systems designed primarily for export infrastructure rather than domestic consumption—a reversal of the fossil-era logic in which pipelines served local industries first. Meanwhile Japan and South Korea, constrained by geography, are testing early-stage hydrogen distribution networks that could one day support their industrial valleys.

Across these regions the pattern is clear: pipelines are no longer simply transporting a single fuel but reshaping how energy moves. They connect production to storage, storage to industry, and countries to one another in configurations that were unimaginable under the fossil system. In the 20th century pipelines determined where energy flowed and who held influence; in the 21st century hydrogen pipelines will play the same role with a different logic. Instead of reinforcing old dependencies around a single fuel, they will underpin flexible, multi-vector systems and new forms of energy partnership. The European Hydrogen Backbone and the Scotland–Germany interconnector already illustrate this shift: a continental network and a cross-sea corridor that together hint at the future geography of energy connectivity.

Ports: The New Energy Gateways

Ports have always been fulcrums of global energy flows, serving as gateways for coal, oil, and LNG. They are where fuels change hands, where global trade connects to local industry, and where energy systems reveal their true shape. But the fuels of tomorrow behave differently from those of yesterday. Hydrogen will rarely cross oceans as hydrogen itself, but as ammonia, methanol, liquid organic hydrogen carriers (LOHCs), synthetic hydrocarbons, or occasionally liquid hydro-

gen. At the same time, CO_2 will begin travelling outward from industrial clusters toward sequestration sites, and electricity may increasingly arrive through high-voltage subsea cables.

This shift transforms the identity of the port. No longer a simple transit hub, it becomes an energy conversion centre—an industrial complex where molecules change form and value is added. Future ports must receive and store multiple hydrogen carriers, crack or reform them when required, liquefy or load CO_2, provide zero-carbon bunkering fuels, and integrate operations with local renewable generation, storage systems, and backbone pipeline networks. Tanks must be redesigned, safety systems expanded, and workforce capabilities extended to include chemical processing, digital monitoring, and multi-vector asset management.

In this sense, ports become energy islands at global scale: places where supply chains begin, where carriers change state, and where the economic logic of long-distance hydrogen transport is determined. How ports reinvent themselves over the next two decades will shape not only maritime fuels but the structure of global energy trade for the rest of the century.

Deep Dive 1: Port of Rotterdam — Europe's First Full-Spectrum Hydrogen Hub

Few places capture the scale of this transformation more vividly than the Port of Rotterdam. Long a cornerstone of Europe's oil and chemical economy, Rotterdam is now positioning itself to become the continent's primary gateway for hydrogen and hydrogen-derived fuels. The transition is visible everywhere: in the redevelopment of storage terminals, in the construction of ammonia import infrastructure, in the building of cracking facilities, and in the integration of port operations with the emerging Dutch hydrogen backbone.

Rotterdam is not simply preparing to import hydrogen; it is preparing to convert it, store it, reform it, crack it, distribute it, and use it. The port's strategy recognises that hydrogen will not arrive in one form. Some will come as ammonia from the Middle East and Australia. Some will come as methanol or synthetic hydrocarbons. Some may arrive via pipelines or offshore hydrogen production connected to the port's industrial cluster. Rotterdam's infrastructure therefore reflects a systemic mindset: the ability to receive multiple carriers, shift between them depending on market needs, and feed hydrogen-rich molecules into adjacent chemical parks, refineries, and industrial users.

What is striking is not only Rotterdam's commitment to handling hydrogen carriers, but the way the port is embedding itself into Europe's broader energy network. The Dutch national hydrogen pipeline system begins in and around Rotterdam. Cavern storage in the north of the Netherlands will be tied into the port's flows. German industrial demand is expected to absorb large volumes routed through the port. Rotterdam, in effect, is becoming the western anchor of Europe's hydrogen economy, playing the same role for hydrogen that it played for oil — but with infinitely more complexity.

Rotterdam's transformation shows what it means for a port to become a "conversion hub": a place where energy arrives in one form and leaves in another, where infrastructure is designed for flexibility rather than a single fuel, and where the port acts not as a passive waypoint but as an active part of the industrial system. It is, in every sense, the blueprint for what an energy gateway looks like in a multi-vector world.

Deep Dive 2: Singapore — The Maritime Fulcrum of the Future Fuel System

If Rotterdam is the archetype of the hydrogen-importing industrial supercluster, Singapore represents something quite different: the marine pivot point through which global shipping will decide its future fuels. Singapore is the busiest bunkering port in the world, the strategic hinge of maritime logistics, and a city-state whose energy strategy must always be aligned with international trade. It is uniquely positioned — and deeply compelled — to shape how ships refuel in the decades ahead.

Singapore's ambition is not merely to offer hydrogen or ammonia bunkering; it is to build the world's first multi-fuel bunkering ecosystem, capable of safely handling ammonia, methanol, synthetic fuels, and potentially hydrogen derivatives. Its strategy begins not with physical infrastructure but with safety frameworks. Singapore recognises that no global shipping transition will succeed without internationally trusted rules for storing, transferring, and bunkering hazardous new fuels. It has therefore become a global centre for standards development, simulation studies, training protocols, and risk assessment. The port is effectively writing the rulebook for how the next generation of marine fuels will be handled.

Singapore's approach stands in contrast to Rotterdam's. Rotterdam begins with infrastructure; Singapore begins with governance. Yet both strategies reflect their respective roles. Singapore's energy system is not built on heavy industry but on logistics, finance, and maritime services. Its influence derives from the

sheer volume of ships it serves and from its reputation for operational excellence. If Singapore certifies a fuel, the world listens. If Singapore refuses a fuel, the world hesitates. Its decisions effectively guide global shipping behaviour.

In recent years, Singapore has conducted pilot operations for ammonia bunkering, explored methanol supply chains, and initiated partnerships with Australia and the Middle East for hydrogen derivatives. The port is preparing not only to supply these fuels, but to ensure that vessels, terminals, tug operators, and emergency services can work safely with them. It is maritime decarbonisation through systems design — quiet, rigorous, and incredibly influential.

Singapore shows that the future of hydrogen and its derivatives is not just about infrastructure. It is about trust, rules, and the choreography of global transport. Without Singapore's adoption and standard-setting, no marine hydrogen carrier — ammonia, methanol, synthetic fuels — will ever scale globally.

A Tale of Two Ports — and Two Futures

Rotterdam and Singapore embody two complementary sides of the emerging hydrogen economy. Rotterdam shows how a port can become an industrial energy hub — a place where hydrogen carriers are transformed, distributed, and consumed within a vast industrial cluster. Singapore shows how a port can become a maritime rule-setter — a place where global shipping experiments with, validates, and eventually standardises new fuels.

Both ports are building infrastructure, but more importantly, both are crafting roles. Rotterdam is anchoring Europe's internal energy flows; Singapore is shaping the world's external ones. Rotterdam demonstrates scale; Singapore demonstrates influence. Taken together, they foreshadow a future where the global energy system is not defined by oil terminals and LNG berths, but by ports capable of handling an entire family of hydrogen-based molecules.

Ports, in this new world, are no longer transit points. They are the gateways through which the energy transition will move — literally and strategically.

Global Momentum — A New Geography of Hydrogen Ports

While Rotterdam and Singapore illustrate the depth of change at two iconic ports, the transformation is unfolding across the world.

In Belgium, the Port of Antwerp–Bruges is emerging as a central node in Europe's hydrogen and CO_2 networks. It is building large-scale ammonia handling systems,

expanding liquid bulk storage, and developing infrastructure to export captured CO_2 to the North Sea for sequestration. Closely integrated with neighbouring chemical clusters, it complements Rotterdam's role with a stronger emphasis on circular carbon systems.

In Asia, ports such as Yokohama and Kawasaki are pioneering early liquid-hydrogen import chains, demonstrating what cryogenic hydrogen logistics could look like at scale. South Korea's port of Ulsan is moving along a similar path, preparing infrastructure for ammonia co-firing fuels and supporting the country's growing hydrogen mobility and industrial sectors.

Across the Pacific, Australian terminals in places such as Gladstone and the Pilbara are being redesigned as export gateways for hydrogen-derived fuels produced from vast solar and wind resources—sending molecules across oceans rather than across national grids.

Meanwhile, ports in Oman, the UAE, and Saudi Arabia are developing large ammonia export capacity aimed at supplying Europe and Asia. For fossil-fuel exporters, this marks a strategic shift: maritime infrastructure once built for hydrocarbons is being reoriented toward hydrogen-based trade.

Together these developments point to a new geography of energy, where ports are no longer simple throughput points but strategic nodes shaping supply chains, investment flows, and geopolitical alignment.

Storage: The Backbone of Flexibility

In the fossil energy system, storage played a largely peripheral role. Coal sat in piles, crude oil in tanks, and gas flowed through pipelines that could absorb small variations in supply and demand. The fuels themselves acted as storage, allowing energy systems to operate with relatively little attention to buffering across time.

When renewables dominate the energy mix, this logic reverses. Storage becomes a structural requirement rather than a convenience. Instead of simply holding fuel in reserve, storage must actively absorb volatility, balance renewable production, and ensure that energy is available precisely when and where it is needed.

Different technologies address different timescales. Batteries smooth second-to-second and hour-to-hour fluctuations. Hydrogen provides medium- to long-duration flexibility capable of stabilising systems across days or weeks. Synthetic fuels extend storage into seasonal balancing and allow energy to move across continents. At the largest scale, geological formations—salt caverns, depleted

fields, and porous reservoirs—offer vast storage volumes capable of supporting entire industrial clusters.

This shift elevates storage from contingency to foundational design principle. It aligns intermittent generation with firm demand, stabilises industrial operations, and provides resilience across multiple timescales. In the fossil world we stored fuels; in the emerging system we store flexibility.

Storage therefore appears at multiple scales. Geological formations provide deep seasonal buffers for regional systems, while compact on-site storage supports the day-to-day operation of trains, logistics hubs, and mobile power systems. Both are essential. Teesside and the Humber illustrate the first; emerging hydrogen rail and logistics depots across Germany illustrate the second.

Deep Dive 1: Teesside & the Humber — Caverns as Industrial Anchors

Few regions illustrate the industrial logic of hydrogen storage more clearly than Teesside and the Humber. These are dense industrial ecosystems — refineries, chemical parks, steelworks, and energy facilities sitting within a geography shaped by a century of heavy industry. Critically, they also lie above extensive salt formations that have been used for decades to store natural gas. This geological inheritance gives them a unique advantage as the UK transitions toward hydrogen.

The emerging plan is straightforward but transformative. Existing gas caverns are being assessed for hydrogen suitability, while new caverns are being planned specifically for hydrogen. Once operational, these caverns will act as the deep buffers that allow gigawatt-scale electrolysis to operate flexibly. When offshore wind is abundant, hydrogen can be produced cheaply and stored. When industrial demand rises, or when the grid is strained, hydrogen can be withdrawn as needed.

This decoupling of production and consumption is the heart of hydrogen system resilience. Industrial plants cannot adjust their processes every time the wind shifts. Electrolysers cannot run at high capacity without a place to put their output, while the grid cannot depend on electrolysis without confidence that hydrogen can be flexibly absorbed. Cavern storage solves all three problems simultaneously. It turns Teesside and the Humber into true energy islands — clusters capable of managing their own internal balance across seasons and market conditions.

The importance of these caverns reaches beyond local industry. They underpin the UK's broader hydrogen strategy by providing the headroom required for both domestic production and future imports. In a world where hydrogen flows will be more dynamic and varied than natural gas ever was, such geological buffers become not just useful but indispensable.

Deep Dive 2: Distributed Storage — Rail Depots, Logistics Hubs, and the Infrastructure We Overlook

While cavern storage anchors regional systems, the real day-to-day visibility of hydrogen often emerges in smaller, more localised forms of storage that keep mobility and logistics functioning smoothly. Nowhere is this more evident than in the new hydrogen rail infrastructure taking shape in southern Germany.

In Baden-Württemberg, Siemens Mobility's Mireo Plus H trains are being deployed as part of Deutsche Bahn's H2goesRail programme. These trains are supported by a purpose-built depot designed not around demonstration timelines, but around the precision of commercial rail operations. The depot incorporates high-pressure storage tanks, compression systems, and refuelling equipment engineered to guarantee consistent, reliable service.

Here, storage performs a very different role from the cavern systems of Teesside and the Humber. It does not buffer seasonal imbalances or stabilise multi-gigawatt industrial clusters. Instead, it ensures that trains running tight timetables can refuel predictably, regardless of upstream delivery patterns. Hydrogen may arrive in irregular batches, but the service frequency must remain fixed. On-site storage bridges that gap. It smooths operational rhythms, providing hydrogen as needed, in the right quantities, at the right pressures, without forcing operators to adjust timetables or depend on just-in-time logistics.

A similar pattern is emerging in logistics centres where hydrogen powers forklift fleets, yard tractors, heavy trucks, and stationary fuel-cell generators. Facilities that adopt hydrogen as part of their operations quickly discover that storage — whether in compressed cylinders, modular tank systems, or trailer-based mobile units — becomes the quiet enabler of reliability. It allows local hydrogen production to fluctuate, deliveries to be batched, and consumption to follow the natural workflow of the site.

This distributed storage logic is also reshaping the temporary and mobile power markets. Hydrogen fuel-cell units used for construction sites, events, film production, remote operations, or emergency response rely on modular storage systems that can be transported as easily as diesel once was. These systems may be

small compared to caverns, but they are transformative. They allow hydrogen to function as a practical substitute for diesel in environments where noise, emissions, or transportability matter as much as cost.

What unites these examples is the way storage enables independence. Rail depots become independent of upstream variations. Logistics hubs become independent of external delivery patterns. Mobile power systems become independent of grid connections and fuel trucks. Distributed hydrogen storage turns hydrogen into a usable, dependable energy carrier for everyday operations.

A System That Stores Flexibility, Not Just Energy

Taken together, Teesside, the Humber, the Baden-Württemberg hydrogen rail network, and the countless emerging logistics hubs reveal a key truth: hydrogen storage is not a monolithic concept. It is a layered capability. Caverns provide strategic resilience. Depot-scale storage provides operational certainty. Mobile storage provides flexibility at the edge of the system.

The old energy system worked because fuels were easy to store and move.

The new system will work because hydrogen storage — at every scale — provides the flexibility the system itself lacks.

Hydrogen is not the only energy carrier that relies on storage, but it is the one that transforms storage into a structural cornerstone rather than a peripheral afterthought. In that sense, storage will do for the hydrogen economy what refineries once did for oil: tie the system together.

A Broader Global Picture — Caverns, Tanks, and the Geology of Advantage

The United Kingdom's cavern assets are not unique. Other countries with suitable geology are beginning to make similar moves, and one project in particular demonstrates how this model is scaling internationally.

In the United States, the ACES Delta project in Utah is becoming one of the world's first utility-scale hydrogen cavern storage systems. Designed to pair gigawatt-scale electrolysis with vast salt formations, ACES provides seasonal balancing for the western US grid — exactly the function hydrogen storage is expected to provide more broadly in highly renewable systems. Unlike many proposed megaprojects, ACES is moving forward decisively, with construction well underway and commercial operation expected imminently. The project stands as proof

that large-scale hydrogen storage is not only technically viable but economically aligned with renewable integration.

Across the Nordic region, countries with strong offshore wind ambitions and favourable geology — Denmark, Norway and parts of Sweden — are exploring cavern storage as a natural extension of their renewable build-out. These caverns would not merely store hydrogen for industry; they would serve as seasonal reservoirs for national energy systems with high renewable penetration.

In Germany, operators are assessing which of the country's existing natural gas caverns can be safely repurposed for hydrogen, and where new caverns might fill strategic gaps. The motivation is clear: Germany's industrial load is enormous, and long-duration balancing cannot rely solely on batteries or imports.

Meanwhile, in countries without cavern geology, innovation takes a different form. Japan and South Korea are advancing tank-based storage for liquid hydrogen and hydrogen bound in liquid organic carriers. These systems do not provide seasonal storage, but they offer strategic and operational flexibility where underground options are limited. Similarly, Europe's mobility hubs rely on high-pressure tank systems that cycle rapidly, creating a storage layer optimised for daily operations rather than long-term buffering.

Across these examples, a common theme emerges, specifically that hydrogen storage adapts to the geology, industry profile and operational needs of each region. Some nations gain natural advantage from salt formations; others compensate with chemical, cryogenic or high-pressure surface technologies. The result is a global mosaic of storage solutions, each reflecting the contours of local energy systems.

Distributed and Mobile Infrastructure: Refuelling, Logistics, and the Edges of the New System

If pipelines and ports form the durable backbone of the hydrogen economy, distributed and mobile infrastructure defines its reach — the points where hydrogen is actually used, handled, and judged in day-to-day operations. This is where abstract strategies meet practical realities: at refuelling stations, fleet depots, construction sites, logistics hubs, railway yards, and remote locations where the power grid cannot fully serve. These edge environments determine whether hydrogen becomes a working solution or remains a theoretical one.

Not all infrastructure of the future will be large, fixed, or centrally located. Some of the most significant assets will emerge precisely at these edges of the system—

in logistics depots, along rail lines, on construction sites, in mining regions, at film sets, and on military installations. These distributed and mobile energy nodes represent a new frontier: modular, flexible systems that deliver energy exactly where it is needed, even when conventional infrastructure is absent.

Hydrogen refuelling stations illustrate this shift clearly. They are not simple pumps but compact energy ecosystems that integrate multi-pressure storage, compression, thermal management, safety systems, and digital control. Some stations generate hydrogen on-site; others receive it through pipelines or delivery. Each acts as a miniature energy island capable of supporting heavy-duty transport reliably and efficiently.

Rail and logistics hubs extend this idea. Hydrogen-powered freight corridors require dependable on-site refuelling and storage aligned with tight operational windows and high throughput. These hubs become hybrid energy and logistics centres—places where goods and molecules move together. They store hydrogen, compress it, dispense it rapidly, and sometimes produce it locally to avoid long-distance trucking.

The temporary power market reveals another expression of this distributed paradigm. Diesel generators dominate construction sites, events, remote operations, and grid-constrained regions, but hydrogen fuel cells paired with batteries are beginning to provide cleaner, quieter, digitally managed alternatives. These mobile units operate as self-contained microgrids, deployable anywhere and coordinated through remote monitoring systems.

This logic extends naturally to off-grid and remote operations—mines, island communities, polar research stations, forward military bases, offshore installations. Hydrogen, paired with renewables, batteries, or small modular reactors, enables these sites to build energy islands entirely independent of fragile diesel supply chains. In such environments, modularity and self-sufficiency are not conveniences but operational necessities.

Even the by-products of hydrogen systems—oxygen, heat, and water—take on new value when infrastructure becomes distributed. Oxygen supports industry and water treatment; heat serves buildings or industrial processes; water feeds local loops. In decentralised systems, by-products become resources rather than waste.

Together, these distributed and mobile energy nodes round out the architecture of Energy Tomorrow. The future energy system will not be defined solely by pipelines, ports, caverns, and centralised plants, but also by millions of modular,

flexible nodes extending energy access into places the old system could never reach. They represent the adaptive, practical edge of the transition—the frontier where the future becomes tangible.

These systems are also where hydrogen will ultimately be judged: not by policymakers or planners, but by drivers, technicians, and fleet operators who rely on them every day. If hydrogen infrastructure proves reliable and predictable in these environments, the transition will succeed.

Every Energy Transition Looks Wasteful Before It Looks Inevitable - The Station Waiting for Its Queue

On a grey morning outside Poznań, trucks roll past in a steady rhythm towards warehouses and industrial estates that have been expanding here for years. Beyond a low fence, a new piece of the future stands ready: a hydrogen refuelling station, all clean lines and stainless steel, its canopy still looking almost too bright against the sky. To most people passing by, it looks like a slightly unusual extension of a regular fuel station. To Roy van Son, stepping out of his car with a folder under his arm, it looks like something more complicated: a bet placed on a market that is still deciding what it wants to be.

Roy works for Resato Hydrogen Technology, a Dutch company that designs and builds refuelling stations across Europe. He spends much of his time travelling between early adopters, cautious municipalities, curious logistics managers and investors trying to understand what hydrogen might mean for their business. To sell a hydrogen station, he realised early on, is not simply to sell a piece of equipment. It is to sell a piece of the future — and then help people figure out how to live with it.

Poland is a fitting setting for this moment. It is already one of Europe's largest hydrogen producers, historically for refineries and chemical plants, and it has declared ambitions to pivot toward low-carbon hydrogen as part of its wider energy transition. Consortia and oil majors are planning networks of hydrogen refuelling points stretching from Warsaw and Poznań to Gdańsk, Wrocław and Katowice. Public stations are beginning to appear; one serving buses and cars in Poznań, another in Katowice, and others in Warsaw and Gdańsk under the NESO brand.

But these are still early days. The stations are like punctuation marks in a paragraph that has not yet been written.

Roy walks the perimeter of the site with the operator, a local company that has just taken delivery of the Resato station. The hardware is installed. The safety systems are tested. The first trial fills have been completed. On paper, everything is ready.

What is missing is volume.

“How many vehicles do you have now?” Roy asks.

“Four buses,” the operator replies. “Two trucks on order. A few cars that come by now and then.”

Roy nods. It is an answer that he has heard across the continent. The station has a theoretical capacity that could support a sizeable fleet. The actual daily throughput is only a fraction of that.

This is the chicken-and-egg problem of hydrogen infrastructure made concrete. Vehicles will not come without stations; stations cannot be profitable without vehicles. Bankers ask for utilisation projections. Fleet operators ask for price guarantees. Policy makers talk about 2030 and 2040 targets. Roy lives in the narrower world of the present: this quarter, this tender, this customer who must justify the decision to a board that still thinks mostly in diesel and spreadsheets.

Yet he is not discouraged. Every station begins like this — overbuilt for today, calibrated for tomorrow. Resato’s systems are modular for that reason. A site can start small, supplying a few hundred kilograms a day, and then expand toward several tonnes as fleets grow. Technically, the model works well. Financially, it can feel unnerving. Someone must pay for capacity that sits idle until the market arrives.

Inside the small office next to the station, Roy sits down with the operator and two representatives from the city transport authority. They talk through the numbers: capital costs, operating expenses, potential grants. The project has attracted support from EU programmes and national funds designed to de-risk early infrastructure. Poland’s hydrogen strategy envisages dozens of such stations by 2030. Even so, the first movers feel exposed. They watch the price of hydrogen at the dispenser, compare it to diesel, and remind themselves that the gap reflects timing as much as technology.

Roy’s role in conversations like this sits somewhere between engineer and translator. He explains how higher utilisation will eventually drive costs down, how the same compression and storage system can serve buses, trucks and cars, and how fixed costs spread across larger fleets. He points to projects in Germany and the

Netherlands where utilisation is slowly rising, and to emerging refuelling systems for hydrogen rail that hint at future freight corridors. But he never pretends the numbers already work on their own. Instead, he frames them for what they are: the cost of being early.

Outside, a hydrogen bus pulls into the station. It is one of the new vehicles serving a municipal route, its livery bright against the low Polish sky. The driver steps out, connects the nozzle, and within minutes the tank is full. The bus pulls away again, quiet compared with its diesel predecessors, leaving almost no trace of what just happened.

For Roy, this moment is the argument in its purest form. Hydrogen refuelling already works this simply. The hardware is no longer the limiting factor. The challenge lies in everything around it: financing models, policy consistency, fleet procurement cycles and fuel price trajectories.

He thinks about how different the site may look in five or ten years if the plans on paper become reality — if networks announced by companies like Orlen expand across the country, if bus fleets multiply from dozens to hundreds, and if hydrogen trucks begin to appear along the motorway corridors linking major cities. On that day, this station will look like an obvious investment. Few will remember that for its first years it stood slightly ahead of demand.

He has seen this pattern before. Electric vehicle chargers that once stood empty now face queues at motorway rest stops. Wind turbines once criticised as uneconomic now form the backbone of northern Europe's power system. In every transition, someone must build first, endure the period of scepticism and underutilisation, and then watch as the world gradually catches up.

As Roy prepares to leave, he walks one final round of the station, checking status screens and listening to the steady rhythm of the compressor. There is pride in this kind of work — infrastructure that functions quietly when everything is working well. There is also humility: the understanding that even the best technology ultimately depends on decisions made far beyond the site itself — in ministries, corporate boardrooms and freight depots.

He starts the car and pulls onto the access road. Trucks thunder past toward the motorway. One day, one of them may pull into this station instead of a diesel pump. When that happens, no one will think about the spreadsheets, the cautious meetings or the early stations that waited patiently for demand.

They will simply see a truck refuelling and a station with a queue.

The lesson is not that early hydrogen infrastructure is misguided, nor that the economics are broken. It is that timing is the hardest variable to model. Infrastructure must exist before fleets commit, before corridors are fully operational, and before demand becomes visible at scale. Someone must always build first — and absorb the cost of waiting.

This is what the transition phase looks like in practice: not a smooth curve from pilot to scale, but a long middle period where assets are technically ready, commercially fragile and strategically essential. The hydrogen system does not fail in this phase; it learns, and the companies operating in this gap are not misaligned with the future. They are holding space for it, long before the volumes arrive.

Deep Dive: Digital Systems for Mobility — The Next Frontier of Hydrogen Use

The mobility sector shows why digitalisation is equally essential at the scale of day-to-day operations. Hydrogen refuelling stations, fleet depots, and logistics hubs still operate today in a remarkably simple way: stations store hydrogen, vehicles arrive, refuelling happens, and operators manage the rest manually. It is a workable model for pilot projects, but it cannot support a system expected to serve buses, trucks, rail, and specialised fleets with high uptime requirements.

Transport infrastructure has no tolerance for uncertainty. Buses must depart on schedule. Freight must move predictably. Trains cannot wait for hydrogen deliveries. As hydrogen scales into mainstream mobility, the sector will need digital systems that can predict demand, coordinate refuelling, manage storage levels, and optimise station operations. The physical assets alone will not be enough.

This next generation of digital platforms is beginning to take shape. Operators and technology developers are quietly building software environments that will coordinate every part of the hydrogen mobility chain. These tools will integrate vehicle telematics, route planning, station capacity, storage state, and (where relevant) on-site electrolysis. A fleet manager will no longer simply decide when a bus refuels; the system will propose — or even automate — the optimal pattern based on upcoming service schedules, pressure levels at nearby stations, and expected hydrogen availability.

At logistics centres, similar intelligence will emerge. Hydrogen-powered forklifts, yard tractors, or heavy trucks operate in patterns that repeat daily. Digital systems will learn these patterns and balance hydrogen production, storage, and deliveries accordingly. Even mobile hydrogen generators used for construction sites or event power will eventually be tied into predictive digital tools that mon-

itor consumption, pressure, and runtime, triggering resupply before an interruption occurs.

None of this exists at scale today, and that is precisely why it must be addressed now. The hydrogen mobility sector will not scale through hardware alone. It needs digital systems that hide complexity from operators, making hydrogen as reliable and frictionless as diesel once was. Without this digital layer, hydrogen mobility remains a pilot. With it, hydrogen becomes a genuinely practical transport solution.

Germany's Hydrogen Refuelling Network — Early Infrastructure, Realigned to Real Demand

No country illustrates the evolution of hydrogen mobility better than Germany. In the early 2010s, Germany took a pioneering step with the launch of H2 Mobility, a consortium of major industrial players dedicated to building a nationwide hydrogen refuelling network. The vision was long-term: hydrogen passenger vehicles would eventually take their place alongside battery-electric cars, and building the infrastructure early would accelerate adoption.

The stations were built to a high technical standard, with advanced compression, automated safety systems, and high-pressure dispensing. However, the timing was ahead of the market. Automakers ultimately focused their early efforts on battery-electric platforms, leaving the hydrogen passenger segment smaller and slower to develop than expected. The result was not a strategic error — the long-term logic for hydrogen passenger mobility may yet re-emerge — but a sequencing mismatch. Germany built the network before the vehicles arrived.

What followed is an instructive example of strategic adaptation. As the early hydrogen market matured, it became clear that the first segment to scale was not private cars but heavy-duty transport. Trucks, buses, regional delivery fleets and specialised utility vehicles need rapid refuelling, long range, predictable uptime and continuous duty cycles — demands that hydrogen is uniquely capable of meeting.

Germany responded by re-orienting its network toward this reality. Stations originally designed for passenger vehicles are being upgraded, expanded, or repurposed to support heavy-duty refuelling. New stations are being built with significantly higher throughput, larger storage volumes, increased compression capacity, wider access for articulated trucks, and integration points with logistics centres and highway corridors. What began as an early foundation for passenger

cars is now evolving into a strategically aligned, high-value asset for freight decarbonisation.

This shift aligns directly with the European Commission's AFIR regulations, which envision a continuous hydrogen corridor running from southern Italy through central Europe and Germany, up through Denmark and Sweden and into northern Norway. Along this spine, heavy-duty refuelling must be available at regular intervals. Germany sits at the logistical centre of this corridor; without a robust network there, the entire northern European hydrogen mobility system cannot function. In effect, the infrastructure built early is becoming the backbone of Europe's future freight network.

Germany's experience captures a larger truth about the hydrogen transition: infrastructure must be flexible enough to adapt as the market reveals itself. Building early was not the problem. The key is the ability to evolve.

Distributed Power and Temporary Infrastructure — Hydrogen at the System's Edge

Refuelling stations represent only one aspect of distributed hydrogen systems. Equally important, though far less visible, is the role of hydrogen in providing temporary and mobile power. The modern world runs on diesel generators far more than most people realise. They power construction sites, rail maintenance compounds, events, film production, telecommunications nodes, emergency operations and countless off-grid activities. Diesel dominates this landscape because it is portable, reliable, and universally understood.

Hydrogen is beginning to enter this domain, not only as a cleaner alternative, but in many cases as a technically superior one. Fuel-cell generators deliver stable, noise-free, vibration-free power without the particulate emissions associated with diesel. They can operate indoors or in environmentally sensitive areas where diesel fumes would be unacceptable. They offer modularity: a compact unit can power small equipment, while multiple units can be combined to support multi-megawatt loads for major construction or industrial applications.

Companies such as Geopura have been instrumental in demonstrating the technical viability of these systems. Their generators integrate fuel cells, sophisticated power electronics, cooling and thermal management, safety systems and modular storage interfaces in a way that allows rapid deployment. The flexibility to use cylinder racks, bundled storage, or trailer-mounted hydrogen supplies shows how mobile hydrogen infrastructure must adapt to diverse operational environments. Embedded monitoring and control systems track load, consumption,

runtime, internal temperatures and pressure levels, allowing operators to treat these units like standard industrial equipment rather than experimental prototypes.

This mobility of energy is one of hydrogen's most intuitive advantages. While the electricity grid attempts to follow where demand arises, hydrogen can be brought directly to the point of use. Temporary power applications therefore become a proving ground for hydrogen's practicality: if it can replace diesel generators — one of the most ubiquitous and resilient technologies in the world — it signals a maturity that extends beyond niche demonstration projects into real operational value.

Digital Systems: The Operating Layer of Tomorrow's Energy System

The fossil energy system could operate with relatively simple control logic. Fuels were dispatchable, demand was predictable, and assets were centralised. Digital systems were useful but not indispensable. A distributed, multi-vector energy system changes that entirely.

If pipelines, ports, and storage tanks form the skeleton of the future energy system, digital systems become its nervous system — the coordinating intelligence that ensures the whole structure behaves coherently. As energy becomes more distributed, modular, and multi-vector, the challenge is no longer simply producing hydrogen or moving it from one place to another. The challenge is orchestrating thousands of assets, each with its own variability, constraints, and operating rhythm: electrolysers, fuel cells, batteries, compressors, pipelines, storage caverns, heat networks, and mobility systems.

Digital infrastructure provides the intelligence required to coordinate this complexity. It forecasts renewable output, optimises electrolyser operation, balances heat and power loads, manages storage across multiple vectors, and responds to market and grid signals in real time. Digital twins allow planners to model entire clusters or ports before construction begins, while advanced optimisation and AI tools help operators manage systems that would otherwise be impossible to control manually. Cybersecurity becomes just as critical as physical safety.

In practice, this digital layer is already emerging across energy infrastructure operators, even if it rarely attracts attention. Transmission system operators, port authorities, and industrial cluster managers increasingly rely on high-fidelity system models to understand how new assets interact with legacy infrastructure. These models simulate flows, pressures, storage behaviour, and demand patterns

that have never existed before, allowing operators to test future operating conditions long before they occur.

What matters is not the software itself, but the shift in mindset it enables. Infrastructure is no longer treated as a static asset to be dispatched, but as a dynamic system whose behaviour must be anticipated and coordinated. Without digital coherence, a multi-vector energy system would be unmanageable. With it, the system becomes adaptive, resilient, and capable of operating across levels of complexity that would have been impossible in the fossil era.

Digitalisation therefore becomes the invisible architecture of the hydrogen economy. It allows thousands of decentralised assets to behave as a coherent system, transforming variability into stability and modular infrastructure into a functioning network.

From Bespoke Engineering to System Architecture

For more than a century, energy projects were defined by bespoke engineering. Every refinery, power station, and pipeline system was a singular piece of infrastructure, designed from the ground up, tailored to its geography and fuel supply, and optimised for a lifetime of predictable operation. The intellectual centre of gravity lived in mechanical design, civil works, and customised control systems. Projects succeeded because the engineering was meticulous; they failed when bespoke complexity overwhelmed budgets and schedules.

Hydrogen disrupts that pattern—not because the technology is new, but because the system it belongs to is fundamentally different. Where the fossil world rewarded bespoke optimisation, the emerging system rewards architecture: the ability to integrate multiple energy vectors, orchestrate distributed assets, and adapt the system as new technologies mature. The pivot is subtle but profound. We are moving from a world where engineers built plants to a world where they design systems instead.

In the traditional fossil paradigm, a power plant was largely self-contained. Fuel arrived, electricity left, and the control systems in between were complex but conceptually linear. Hydrogen plants, energy islands and multi-vector hubs do not work this way. Their performance depends on how they interact with the wider environment—the grid, the gas network, local renewables, industrial loads, and the storage assets around them. A poorly integrated electrolyser can destabilise a microgrid; a misaligned storage strategy can erase project economics; an overly simplistic control philosophy can turn a multi-vector asset into

something stranded and underutilised. The complexity is no longer in the equipment itself; it is in the relationships between components.

This is why system architecture becomes the defining engineering discipline of Energy Tomorrow. It governs how technologies combine, how they balance each other, and how they adapt as conditions change. It requires an understanding of when to prioritise electrons over molecules, how to balance short-duration and long-duration storage, how to manage flows of electricity, heat, hydrogen, ammonia, and CO_2 in real time, and how to align hydrogen production with renewables, demand profiles, market signals and pricing structures. It also means designing assets that can grow, evolve, or repurpose over a twenty- or thirty-year lifecycle. The dominant challenge is no longer technical performance, but systemic coordination.

This transition is not just intellectual—it is grounded in hard experience. In large oil and gas projects, EPC contractors traditionally offered two commercial models: fixed-price turnkey or cost-plus. In my career, I have never seen a fixed-price turnkey project where all of the engineering was genuinely complete, aligned, and properly coordinated when construction began. Inevitably, gaps appeared. Interfaces didn't behave as designed. Assumptions made early in the process no longer held once equipment arrived on site. The contractor had priced what they believed the project was, not what it actually needed to be. The result was always the same: remedial work, redesign, delay, and friction.

Likewise, I have never seen a cost-rate project that didn't run over budget. When engineering effort is open-ended, bespoke complexity has free rein. Scope grows, integration challenges multiply, and costs escalate. It is not incompetence—it is structural inevitability. Bespoke systems generate bespoke risk.

This is exactly the trap that the hydrogen industry must avoid. If every hydrogen plant is custom engineered from scratch, it will simply replicate the overruns, delays, and rework of the fossil era. But when hydrogen systems are modular, interoperable, and governed by coherent architectural principles, uncertainty collapses. Interfaces become predictable. Engineering becomes repeatable. EPCs shift from reinventing the wheel to assembling well-defined modules. Project risk transforms accordingly.

A deeper cultural change follows. The fossil era fostered engineering disciplines built around certainty: design a plant once, build it once, and operate it for decades. The new system is defined by uncertainty. Technologies evolve at different speeds; regulations shift; markets fluctuate; energy vectors diversify; and geopolitical events reshape supply chains. Infrastructure must be adaptable to a world

we cannot fully predict. Bespoke optimisation creates fragility; architectural thinking creates resilience.

A helpful analogy comes from computing. Early computers were bespoke machines—hand-wired, custom-built, enormous, and designed for a single purpose. Over time, computing shifted to layered architectures: modular components, standardised protocols, interoperable systems, and the ability to upgrade endlessly without rebuilding the machine. Energy is undergoing the same transformation. We are moving from custom installations to configurable systems that evolve.

The companies that thrive will not be those that design the most elegant standalone electrolyser plant, but those that design the most adaptable energy system. They will think in terms of use cases, load profiles, digital orchestration, and lifecycle evolution. They will treat hydrogen not as an isolated project but as a strategic capability within a broader operational ecosystem. An energy system is no longer built—it is assembled, configured, and continuously optimised.

This transformation widens the competitive landscape. In the bespoke model, only large incumbents with deep engineering capacity could deliver complex energy infrastructure. In the architectural model, smaller innovators, digital integrators, automation specialists, and microgrid developers can all contribute meaningfully. Innovation accelerates because the system becomes a platform rather than a fortress.

The deeper message is simple: the energy transition is not a contest between technologies but a contest between architectures. Those who continue thinking in silos will build stranded assets. Those who think in systems will build the backbone of Energy Tomorrow.

The Cost of Complexity and the Price of Transition

As the energy system shifts from centralised, fossil-based assets to a modular, multi-vector architecture, it enters a period where efficiency temporarily declines before it rises again. Every major infrastructural transformation passes through this phase: the new system must be built while the old one continues to operate, and neither performs at its best. The result is duplication, friction, and investment that does not immediately translate into lower costs or greater reliability.

For more than a century, the fossil system delivered low-cost energy because it was mature, standardised, and stable. The emerging hydrogen-enabled system is

not yet any of those things. Equipment is improving rapidly but is not yet fully mass-manufactured; integration practices vary between projects; and most new assets must coexist with infrastructure never designed for them. Costs rise before they fall, interfaces take time to harmonise, and standards inevitably lag deployment. This is not failure—it is simply the reality of building a new system from within an old one.

The transition also does not eliminate oil and gas; it reshapes their role. Even in a deeply decarbonised system, hydrocarbons will remain relevant for industrial heat, petrochemicals, aviation fuels, shipping blends, peaking power, and strategic reserves. They will no longer define the system, but they will remain part of it. The transition is therefore not a binary switch but a rebalancing of energy vectors, where molecules and electrons coexist in new proportions.

Much of the transitional cost comes from operating two architectures simultaneously. Pipelines, ports, grids, storage sites, and digital platforms must continue serving legacy fuels while preparing for new ones. Hydrogen corridors emerge alongside natural gas networks. Ports add ammonia and methanol infrastructure while still handling oil and LNG. Grids integrate distributed generation while relying on large central plants. Dual operation is expensive, but unavoidable.

At the same time, cost is migrating from hardware to integration. In the fossil system, value lay primarily in large physical assets. In the emerging system, the challenge is making diverse components work together—electrons with molecules, short-duration storage with long-duration buffers, digital control systems with legacy infrastructure. The competitive advantage increasingly lies not in owning the largest equipment, but in orchestrating the system effectively.

Modularity helps reduce friction, but it does not eliminate the learning curve. Early projects show that modular designs contain mistakes rather than amplifying them, allowing systems to evolve without locking in early imperfections. Even so, learning—standardising interfaces, refining safety frameworks, and developing new automation practices—carries its own cost.

Another transitional challenge lies in capability. A multi-vector system requires different skills: digital orchestration rather than linear control, cross-vector safety rather than siloed disciplines, and system thinking rather than component optimisation. The workforce of the fossil era was not trained for this reality, and organisations must invest heavily in building the capabilities the new architecture demands.

Despite these frictions, the direction is unavoidable. The world cannot meet decarbonisation goals or resilience needs with the architecture of the past. The cost of transition is real, but the cost of remaining inside a brittle, centralised, carbon-dependent system would be far greater.

What we must accept is unevenness. Some sectors and organisations will become system-capable sooner than others. Some infrastructures will adopt modularity quickly; others will evolve gradually as legacy assets retire. The system will not transform everywhere at once—it will move in layers.

And that is enough. Each layer built today strengthens the architecture of tomorrow. The true price of transition is not hydrogen or electrolysers themselves, but the temporary cost of operating two systems while constructing the one we ultimately need.

End-Point Infrastructure: Where Hydrogen Meets the Real World

Across the rest of the world, similar patterns are emerging as hydrogen mobility and distributed power begin to take shape in diverse and sometimes surprising environments. In California, one of the earliest hydrogen refuelling networks outside Europe demonstrated both the promise and the pitfalls of first-generation infrastructure: technically advanced stations, uneven utilisation, and a transition now moving firmly toward supporting heavy-duty vehicles in the state's logistics corridors. Japan, by contrast, has quietly built a dense network of hydrogen bus and rail depots designed around metropolitan mobility rather than long-distance freight, showing how depot-centric refuelling can anchor early demand.

In the UK, plans in Teesside and parts of Scotland are preparing integrated depots where hydrogen trains and buses can refuel alongside stationary fuel-cell generators, mirroring the multi-use logic seen in Germany. Even outside civilian transport, military organisations, such as the US Army and NATO partners, are testing hydrogen generators to replace diesel units in field operations where silence, reduced heat signature and improved air quality offer operational advantages. Mining operations in Australia and Canada are beginning to trial hybrid hydrogen–battery microgrids to power remote sites where diesel delivery is costly and emissions-intensive. Across these varied examples, the direction of travel is unmistakable: hydrogen is progressing from isolated pilot infrastructure to a distributed ecosystem of depots, refuelling hubs and mobile generators that bring energy directly to the point of use.

The Transition Phase — Operating Between Today and Tomorrow

Everything described so far in this chapter — pipelines, ports, storage, digital systems, distributed assets — represents the infrastructure of the future. But the truth is far more complicated: for the next decade, we will live in a hybrid system, half belonging to the fossil world, half emerging from the renewable one. This period is not simply an engineering challenge. It is an operational and commercial puzzle that must be solved in real time, with imperfect tools, incomplete infrastructure, and economics that shift month by month.

The transition is not a switch. It is a negotiation — between electrons and molecules, between legacy systems and new ones, between what is optimal today and what is essential tomorrow.

A System Learning to Become Itself

What emerges from all these developments is an energy system learning, almost awkwardly at first, to operate across two realities at once. It is attempting to behave like the system of tomorrow while still bound to the constraints, rigidities and legacy structures of today. Elegant theoretical models — with their smooth curves, perfect optimisation and idealised load profiles — meet the friction of messy practice. Slowly, sometimes reluctantly, the two begin to reconcile.

In this evolving landscape, curtailment and negative pricing stop being statistical anomalies and instead become signals — markers of where flexibility is scarce and therefore valuable. Grid fees, once a marginal line item, suddenly dictate whether connecting to the transmission network is a pragmatic decision or a financial impossibility. Power purchase agreements, meanwhile, act as stabilising anchors in a volatile environment, allowing developers to secure green electricity long before the grid itself is ready to deliver it reliably. Electrolysers, for their part, begin to behave not as rigid industrial machines, but as instruments responding to the weather, modulating output in harmony with shifting winds and sunlight. Across all of this, costs decline not merely because assets grow larger or factories more efficient, but because the system becomes smarter, as operational intelligence begins to substitute for brute-force engineering.

Yet intelligence brings new responsibilities. As energy systems become more digital, interconnected and autonomous, cybersecurity moves from being an operational concern to a strategic necessity. The critical infrastructure of the future will no longer consist solely of power stations, pipelines and substations, but also of millions of connected controllers, sensors, distributed energy islands

and cloud-based optimisation platforms. This expanded digital landscape increases the potential attack surface for malicious actors, making cyber resilience as fundamental to energy security as physical protection. Designing tomorrow's energy system therefore means engineering not only for efficiency and flexibility, but also for trust, resilience and security from the outset.

This is the transition phase. It is uneven. It is untidy. It frequently contradicts itself. Yet it is also entirely normal. Every major energy transformation in history — from coal to oil, from steam to electricity — has passed through a liminal period where the emerging logic must coexist with the remnants of the old. We are living through that period now: a system that is not what it was, not yet what it will be, but unmistakably on its way.

Closing Reflections — The Shape of What Comes Next

The story of hydrogen infrastructure is not a story of instant reinvention, but of gradual realignment — a system turning, piece by piece, toward a new operating logic. Pipelines once built for methane are being taught to carry new molecules. Ports designed around oil are learning the choreography of ammonia, methanol and synthetic fuels. Storage is shifting from a passive reserve to an active balancing mechanism. Digital systems, once peripheral, now function as the quiet intelligence binding everything together. At the edges of the network — in depots, refuelling stations, construction sites and remote locations — hydrogen is proving itself in small, practical ways that matter more than any blueprint.

This is how energy transitions really unfold; not through a single breakthrough or declaration, but through the slow accumulation of capabilities that make the old system less necessary and the new one more inevitable. The infrastructure being built today is imperfect, uneven and often experimental — but it is also irreversible. Each pipe repurposed, each port redesigned, each storage cavern converted and each digital platform deployed nudges the system closer to its next equilibrium.

We are watching the scaffolding of Energy Tomorrow take shape. Not as a replacement of what came before, but as its evolution — more flexible, more interconnected, and ultimately more resilient. The transformation is underway, and although it will take decades to mature, its direction is already unmistakable. The future energy system does not wait to be discovered; it is being built, tested and refined right now, in steel and software, at every scale from continental backbones to the smallest mobile power unit.

14. What Business Can Do

"Culture eats strategy for breakfast." — Peter Drucker

A Quiet Machine in a Very Loud Place

It is a warm Friday evening in late June on the Isle of Wight, just off the southern coast of the UK. The Isle of Wight Festival is underway, and near the town of Newport, at Seaclose Park, tens of thousands of people are gathering as daylight begins to fade. By the time the crowd swells toward its peak, close to 60,000 revellers are packed in front of the main stage. Anticipation hangs in the air — the familiar mix of music, movement, dust, and expectation that defines a major summer festival.

As The Prodigy walk on stage, the mood shifts instantly. The opening notes of Invaders Must Die hit, and the field erupts. Bass rolls across the site, lights snap into life, and the weekend properly begins. From the perspective of the crowd, everything feels exactly as it always has. The sound is immense. The energy is raw. The festival machine is operating at full throttle.

What no one in the audience is thinking about — and crucially does not need to think about — is where the power is coming from.

Behind the scenes, out of sight of the crowd, a quiet system is doing its job. Parts of the festival's power demand are being supplied not by diesel generators, but by mobile fuel-cell power units provided by GeoPura. There is no spectacle attached to them. No banners explaining the technology. No attempt to draw attention to the fact that hydrogen is involved at all. The equipment simply delivers electricity where and when it is needed, reliably and without drama.

This is not a technology demonstration. It is not a pilot framed as an experiment. It is rented infrastructure performing a function — powering operations that must not fail in an environment where reliability is non-negotiable. Artists, crews, vendors, and organisers depend on consistent power. The crowd expects the

lights to come on and the speakers to work. There is no tolerance for excuses, and no appetite for complexity.

And that is precisely the point.

GeoPura is not interesting here because it uses hydrogen. It is interesting because hydrogen disappears entirely from the customer conversation. The festival is not buying fuel cells. It is not managing energy systems. It is buying power — as a service — and transferring the technical risk elsewhere. The outcome is what matters. The technology is simply the means.

In a field full of noise, colour, and motion, the most important innovation is the one no one notices at all.

When the Conversation Is the Constraint

Elsewhere in this book, I return to a well-known observation attributed to Charles Darwin: it is not the strongest of the species that survive, nor the most intelligent, but the ones most responsive to change. The same principle applies uncomfortably well to the hydrogen economy.

Many of the challenges facing hydrogen today are not rooted in technical feasibility. The core technologies exist. Performance has been demonstrated. In many regions, manufacturing capacity, infrastructure plans, and supply chains are already in place. Nonetheless, progress often feels slower and more fragile than the underlying fundamentals would suggest. The reason is not that hydrogen does not work. It is that too often it is framed in ways that fail to adapt to how people, businesses, and institutions really make decisions.

As discussed earlier in the context of social acceptance, adoption rarely hinges on technical merit alone. People do not reject new systems because they are inefficient; they hesitate because they feel unfamiliar, risky, or misaligned with everyday experience. The same dynamic applies in business. When hydrogen is introduced primarily as a technology — something novel that must be understood, evaluated, and believed in — it immediately creates distance. Customers are asked to step outside their core activities and engage with unfamiliar complexity. Project developers are asked to shoulder new categories of risk. Bankers are asked to finance technologies rather than revenues. Each additional explanation, however well-intentioned, increases perceived uncertainty.

Hydrogen struggles most when it is presented as exceptional.

By contrast, hydrogen succeeds when it is embedded inside services that people already recognise. Power, logistics, uptime, resilience, and reliability are not new ideas. They fit comfortably within existing procurement processes, financial models, and operational thinking. When hydrogen sits quietly behind those outcomes, it stops being a conceptual leap and becomes an implementation detail. The focus shifts away from how the energy is produced and toward what the service delivers.

This is where many businesses fail to adapt. Faced with a new market, the instinct is often to educate — to explain why hydrogen matters, how it works, and why it will be important in the future. However, education is not the same as adoption. Most customers do not want to become energy experts. Most financiers do not want to underwrite learning curves. They want solutions that fit into the world as it already exists, not invitations to redesign it from first principles.

Business success in the hydrogen economy increasingly depends on changing the conversation rather than winning the argument. That means speaking in the language of the customer rather than the language of the technology. It means framing value in terms of reduced risk, operational certainty, and delivered outcomes — not system architectures or chemical pathways. It also means recognising that normalisation is more powerful than persuasion.

The companies making the fastest progress are not necessarily those with the most advanced technology. They are the ones that have adapted how that technology is presented, contracted, and absorbed by the market. They understand that hydrogen does not need to be celebrated to succeed. It needs to feel familiar. When hydrogen becomes something that simply works in the background — something that delivers the expected outcome without demanding attention — resistance falls away naturally.

In that sense, the challenge facing hydrogen is not to become more impressive. It is to become less remarkable.

Demand, Confidence, and Normality

Discussions about hydrogen frequently return to the idea of a “chicken and egg” problem: demand will not materialise without supply, and supply will not scale without demand. In practice, this framing is increasingly unhelpful. In many parts of the system, supply-side readiness is no longer the primary constraint. Manufacturing capacity exists. Infrastructure plans are advancing. Technology risk has been reduced. What is missing is not capability, but commitment.

The real bottleneck is demand — and more specifically, confidence.

Organisations hesitate to commit not because hydrogen is unavailable, but because it still feels unfamiliar. Novelty amplifies perceived risk. New contractual structures, new operational dependencies, and new failure modes all sit outside established experience. Even when the technical case is strong, uncertainty around delivery, reliability, and responsibility can stall decisions. Demand, in this context, is less a question of price than of trust.

Confidence emerges when systems behave in ways people already recognise. Familiarity lowers the cognitive and institutional barriers to action. When hydrogen is presented as something that requires new expertise, bespoke contracts, or deep technical engagement, it remains exceptional. When it is delivered as part of a service that mirrors existing arrangements — power provision, fuel supply, logistics, or uptime guarantees — it becomes legible to customers, financiers, and insurers alike.

This is why service-based models matter so much. They do not remove technical complexity; they absorb it. Risk is retained by those best equipped to manage it, rather than being passed downstream to customers who have neither the incentive nor the capability to do so. From the customer's perspective, the transaction looks familiar. Performance is measured in outcomes, not inputs. Responsibility is clear. The hydrogen itself becomes incidental.

Crucially, this shift accelerates confidence long before cost parity is achieved. Projects move forward not because hydrogen is cheap, but because it is understandable. Banks finance revenues rather than technologies. Customers sign contracts for services rather than systems. Demand appears not through persuasion, but through normalisation.

In that sense, hydrogen does not suffer from a classic chicken-and-egg problem. It suffers from an excess of explanation. Demand does not need to be convinced into existence; it needs to be enabled. When confidence replaces novelty, commitment follows — and when commitment appears, the rest of the system begins to move.

Three Companies, One Pattern

When I started writing this book, this was one of the sections I most looked forward to. Not because it would be easy, but because the companies I am about to discuss give me genuine energy. In an industry that can often feel stuck between ambition and delivery, seeing how certain players are approaching the market

gives me real confidence that hydrogen's promise might actually materialise. They are not waiting for perfect conditions; they are adapting to the market as it exists today.

Let's start with the company that we met in the opening: GeoPura.

I first visited GeoPura in Newcastle in November 2022. When I heard the vision, I'll be honest — my initial reaction was that they had bitten off more than they could chew. They were a relatively small team, entering a market that is operationally unforgiving, commercially conservative, and intolerant of failure. Temporary power is not an arena where "good enough" is acceptable. However, it quickly became clear that what looked like overreach was, in fact, focus. They weren't trying to win the hydrogen debate. They were trying to solve a customer problem.

GeoPura deliberately targeted the temporary power market — one of the largest remaining strongholds of diesel generation, and one of the most visible. It is also a market where, until very recently, few hydrogen players were even attempting to compete seriously. From the outset, their differentiation was not framed around the technology itself. It was rooted in how customers interact with it.

Rather than manufacturing containers and selling them into a market that would then need to operate, maintain, and assume performance risk, GeoPura chose a different path: 'Energy as a Service'. Customers do not buy hydrogen equipment. They come with a requirement — power for a film set, a live broadcast, a construction site, or a music festival — and GeoPura delivers the outcome. The fact that hydrogen sits behind that service is secondary to the customer experience.

That framing changes the entire conversation. It removes the need for customers to become energy specialists. It shifts technical and operational risk to the party best positioned to manage it. It simplifies procurement, clarifies accountability, and reduces the psychological barrier to adoption. Supported by favourable policy conditions in the UK — which we have already explored earlier in this book — this approach has allowed GeoPura to grow at a pace that aligns with both market demand and operational reality.

Step by step over the last few years, the company has expanded its fleet, strengthened its delivery capability, and worked toward securing its own green hydrogen supply to support deployments. For me, GeoPura remains one of the clearest examples of what progress looks like when a business stops trying to persuade the market about hydrogen — and instead focuses on delivering something that feels familiar, dependable, and normal.

If GeoPura made hydrogen invisible in power, HTEC did the same in logistics.

HTEC is a hydrogen fuelling and mobility company based in Vancouver, Canada. Unlike some newer entrants, it began in a relatively conventional way: developing fuelling stations, securing local hydrogen supply, and building a stable base of early customers. For a time, the model worked. Operations were running, assets were utilised, and the business had traction.

The challenge emerged when HTEC began to think seriously about expansion.

As the network grew, the familiar problem resurfaced. New fuelling infrastructure required confidence in future demand, yet potential customers were hesitant to commit without proven availability and reliability. The result was a stalemate that we have already encountered several times in this book: no one was willing to finance infrastructure without guaranteed offtake, and no one was willing to commit to offtake without assured infrastructure. In the context of heavy-duty transport, this deadlock was amplified by high upfront capital costs and lingering uncertainty around new vehicle technologies.

HTEC responded not by waiting for the market to resolve itself, but by changing its role within it.

Rather than expecting customers to purchase expensive hydrogen trucks and absorb the associated risk, HTEC took that risk onto its own balance sheet. The company invested directly in the vehicles, retaining ownership and responsibility, and offered them to customers as part of an integrated service. In doing so, HTEC shifted the customer conversation away from fuel choice and capital expenditure and toward something far more familiar: operating a logistics business.

This ‘Logistics-as-a-Service’ approach fundamentally altered the risk profile. Customers no longer needed to believe in hydrogen as a technology or justify large upfront investments. They needed only to assess whether the logistics service met their operational and commercial needs. Hydrogen became embedded in the solution rather than presented as the decision itself.

By adapting the commercial model rather than the technology, HTEC unlocked stalled expansion plans and re-established momentum. The lesson is a consistent one. Progress did not come from convincing customers to adopt hydrogen, but from removing the reasons they had to hesitate.

If GeoPura made hydrogen invisible in power, and HTEC did the same in logistics, Hydropulse is doing it in industrial supply.

Hydropulse emerged as a carve-out from ITM Power with a clear understanding of where many hydrogen projects stall. For industrial customers, the barrier is rarely a lack of interest in hydrogen itself. It is the complexity and risk associated with developing, owning, and operating production assets. Project delivery risk, operational responsibility, and long-term performance uncertainty are often enough to delay decisions indefinitely.

Rather than selling equipment or expecting customers to become hydrogen producers, Hydropulse offers a different proposition. The company builds, owns, and operates the hydrogen production plant, while the customer simply purchases the hydrogen they require. The commercial relationship is structured around supply, not assets. From the customer's perspective, hydrogen becomes a feedstock, not a project.

This model shifts risk to the party best equipped to manage it. Technology integration, operations, and performance all sit with Hydropulse. Customers avoid upfront capital expenditure and the organisational burden of running unfamiliar infrastructure. As with the other examples in this chapter, the hydrogen itself fades into the background. What remains is a familiar transaction: a long-term supply agreement with defined responsibilities and outcomes.

Hydropulse does not change the physics of hydrogen production. It changes who carries the risk. In doing so, it lowers the barrier to commitment and allows projects to move forward that might otherwise remain on paper.

Together, GeoPura, HTEC, and Hydropulse illustrate the same underlying pattern. Progress comes not from asking customers to adapt to hydrogen, but from adapting hydrogen to fit the way customers already operate.

While the contexts differ — temporary power, logistics, and industrial supply — the underlying approach is the same. None of these companies succeed by persuading customers to adopt hydrogen. They succeed by redesigning the commercial interface so that hydrogen is no longer the decision. Risk is absorbed upstream, complexity is hidden, and value is framed in outcomes that feel familiar. In doing so, these models transform technical uncertainty into contractual clarity — the same shift that, as discussed earlier, allows capital to flow. Progress follows not from technological novelty, but from making hydrogen fit naturally into existing ways of operating.

Why “As-a-Service” Changes Everything

The significance of the models described in the previous section is not limited to hydrogen. They reflect a broader shift in how complex systems are commercialised when uncertainty is high and traditional ownership structures no longer work. At its core, the move toward “as-a-service” models is a reallocation of risk — away from customers and toward those best equipped to manage it.

In conventional project structures, technical, operational, and market risks are often pushed downstream to end users. Customers are expected to invest capital, build new capabilities, and carry unfamiliar liabilities in exchange for future benefits that remain uncertain. Unsurprisingly, many choose not to. By contrast, service-based models invert this logic. Technology risk, integration complexity, and performance responsibility sit with the provider, while customers pay for delivered outcomes. This does not eliminate risk; it concentrates it where experience, scale, and incentives align.

This shift is also what makes these models legible to finance. Banks and institutional investors do not fund transitions; they fund cash flows. When revenue is tied to long-term service contracts rather than speculative demand, uncertainty becomes measurable. As discussed earlier in this book, the difference between a promising technology and a bankable project is rarely technical. It is contractual. Service models turn ambition into something that can be priced, insured, and financed.

For customers, the appeal is simpler. They do not buy hydrogen, infrastructure, or systems. They buy power, mobility, or supply certainty. The energy transition becomes something that happens around them, not something they must actively manage. That distinction matters. It lowers internal resistance, shortens decision cycles, and allows organisations to move forward without rewriting their identity or business model.

None of this works in isolation. Service-based approaches depend on partnerships and ecosystems rather than stand-alone projects. No single company optimises the entire value chain. Equipment suppliers, operators, financiers, and end users must align around shared incentives and clearly defined responsibilities. Trust becomes a commercial asset, built through delivery rather than declarations. Risk-sharing replaces risk-shifting, and coordination begins to matter as much as scale.

As hydrogen markets mature, this need for collaboration extends beyond individual projects and into the international ecosystem itself. It is often assumed that

creating a global hydrogen economy will require complete harmonisation of regulations across countries. In reality, that may be neither practical nor necessary. A more achievable objective is to develop sufficient common standards, trusted certification, interoperability and mutual recognition to allow different regional hydrogen markets to work together effectively. Businesses have an important role to play in this process, not simply by complying with standards, but by actively helping to develop them, participating in industry initiatives and building the relationships that underpin trust. The goal is not to make every market identical, but to make them sufficiently compatible that hydrogen can be traded and deployed with confidence across borders.

These structures also accelerate learning. When companies operate assets on behalf of customers, feedback is immediate. Performance data, failure modes, and operational realities surface quickly and visibly. Learning occurs through use, not theory. That learning compounds over time, creating advantages that are organisational rather than technological. Early movers gain not because they guessed correctly, but because they adapted faster.

If there is a practical playbook for business in the hydrogen transition, it is remarkably short:

1. Don't sell hydrogen.
Sell what customers already want and let hydrogen support that outcome quietly.

2. Don't educate customers unless they ask.
Complexity increases friction; familiarity reduces it.

3. Don't wait for universal understanding before acting.
Normalisation follows deployment, not the other way around.

Everything else — policy, finance, standards, and infrastructure — works better once these principles are in place.

From Insight to Action

The lesson from these examples is not that hydrogen is ready, or that the transition will be straightforward. It is that progress depends on how quickly business adapts its own behaviour. Models that hide complexity, absorb risk, and deliver familiar outcomes are already moving faster than those waiting for perfect alignment across policy, finance, and infrastructure.

The question facing most companies is no longer whether hydrogen will play a role in the energy transition. That question is largely settled. The more immediate challenge is whether businesses are willing to change how they engage with it — shifting from selling technologies to delivering services, from transferring risk to managing it, and from waiting for certainty to learning through action.

Those that adapt early will not do so because they predicted the future correctly, but because they positioned themselves to learn faster than others. In a transition defined by uncertainty, that ability may prove to be the most valuable advantage of all.

15. Digitalisation and AI — Accelerating the Hydrogen Transition

"It is not the strongest of the species that survive, nor the most intelligent, but the one most responsive to change." – **Charles Darwin**

Understanding AI

It would be difficult to write about the future of hydrogen without addressing the role AI will play in shaping it.

However, beyond understanding AI as a technological concept, it is more important to recognise how it integrates with hydrogen as a system. AI can act as connective tissue across the value chain — from production through storage, distribution, and utilisation. It does not replace expertise; it amplifies it. By analysing large datasets from electrolysers, storage tanks, pipelines, and end-use applications, AI can identify patterns invisible to human operators, predict maintenance needs, optimise energy flows, and suggest design improvements. In that sense, AI is not a standalone marvel, but a co-pilot guiding a complex energy transition.

To provide some context, there are two books I would recommend for understanding the broader digital transformation we are living through: *The Fourth Industrial Revolution* by Klaus Schwab and *Race Against the Machine* by Erik Brynjolfsson and Andrew McAfee. Both are now over a decade old, but they remain highly relevant in explaining the forces shaping today's technological landscape.

In *Race Against the Machine*, Brynjolfsson and McAfee illustrate the nature of technological acceleration using the well-known chessboard metaphor. A king offers a reward for the invention of chess, and the inventor asks for one grain of rice on the first square, doubling on each subsequent square. What begins as a modest request quickly becomes incomprehensible, ultimately exceeding all the

rice on Earth. The lesson is simple: exponential growth is deceptive. It appears slow at first, but once it reaches scale, it accelerates beyond intuition.

For much of modern history, progress in computing and automation felt incremental — the early squares of the chessboard. But as capabilities compound, each doubling produces disproportionately larger gains.

Crucially, Brynjolfsson and McAfee argue that humanity is now entering the early stages of the second half of this chessboard. The exponential curve has begun to steepen but has not yet reached its full intensity. This is an important distinction. It suggests that the technological breakthroughs we have already witnessed — AI mastering complex strategy games, algorithms designing new materials, machines learning to communicate in natural language — are not the peak, but the beginning.

Each successive doubling will bring advances that are faster, more powerful, and more disruptive than those before. The pace of change that we are experiencing today, remarkable as it may seem, could soon feel slow in comparison to what lies ahead.

This shift can feel unsettling. The idea that technological progress is accelerating beyond our control often raises concerns about displacement or loss of agency. However, as Klaus Schwab argues, these same forces also offer the potential to enhance human capability. The key is not whether technology advances, but how it is applied.

Schwab illustrates this through a second use of chess — not as a metaphor for growth, but for collaboration. While computers can outperform even the best human players, the strongest results are achieved when humans and machines work together. Human intuition, judgement, and adaptability combine with computational speed and precision to produce outcomes neither could achieve alone. This is a useful way to understand AI in practice: not as a replacement, but as an extension of human capability.

For Klaus Schwab, this example encapsulates the opportunity of the Fourth Industrial Revolution. The future he describes is not one in which humans are replaced by intelligent machines, but one in which we work alongside them — each enhancing the other's strengths.

When applied thoughtfully, technology can free people from routine tasks, expand intellectual and creative horizons, and allow individuals to focus on what truly requires human insight and empathy. In this sense, the accelerating pace of

change, far from being a threat, becomes a catalyst for a deeper and more meaningful form of human progress.

This shift has also changed how I think about the skills required for the future. I remember encouraging my son to focus on learning to code, believing it would be essential for the world ahead.

Today, I would give different advice. Coding remains useful, but it is no longer the defining skill. What matters more is a deep understanding of the underlying systems within a field — the physics, the chemistry, the operational realities — combined with the ability to work effectively alongside AI.

Jensen Huang, the CEO of NVIDIA, has expressed this clearly — the future will favour those who combine domain expertise with the ability to work effectively alongside AI. AI can increasingly handle coding and routine tasks, but it is those who understand the physics, chemistry, or operational realities of a system — and can guide AI accordingly — who will create the most value.

Taken together, these perspectives point to a consistent conclusion: technological progress does not replace human capability; it amplifies it when applied intelligently.

At the same time, it is important to be clear about what AI actually is. Not everything labelled "AI" reflects genuine intelligence. Increasingly, systems are presented as intelligent interfaces — capable of responding, guiding, or even "conversing" — yet when examined more closely, they are often little more than structured responses or pattern-matching systems presented through a convincing front end.

At first glance, these tools can feel impressive, even persuasive. But under scrutiny, they can produce outputs that are inconsistent, irrelevant, or simply incorrect. The issue is not that they are useless — many are helpful within defined boundaries — but that they are often mistaken for something more capable than they truly are.

This distinction matters. True AI involves adaptability, learning, and the ability to operate across complex and variable conditions. Without recognising this, there is a risk of overestimating capability and misapplying these tools in ways that introduce inefficiency — or worse, misplaced confidence.

A useful way to understand this in practice is through the process of writing this book. AI did play a role — but not in the way many might assume.

The structure, arguments, and perspectives are mine. What AI offered was leverage: a way to work faster, check myself more rigorously, and express ideas more clearly than I might have done alone. It acted as a sounding board — helping refine ideas, test phrasing, and identify gaps — but it did not determine the direction of the narrative.

Of course, AI is not perfect. Anyone who has tried creating infographics, diagrams, or more complex outputs with it will recognise the reality: it requires guidance, iteration, and patience. It can enhance your work, but it cannot replace the judgement, curiosity, and critical thinking that come from the human mind. Used wisely, it amplifies your capabilities; misused, it can mislead.

For me, the experience of writing this book with AI as a partner has been both a practical tool and a learning process — one that reflects the broader potential, and limitations, of these technologies.

Ultimately, this book reflects the broader argument of this chapter: when technology is applied as an extension of human judgement rather than a substitute for it, the result is not dilution, but amplification.

Understanding artificial intelligence is therefore only part of the story. In practice, AI is not a standalone technology applied in isolation, but one element within a much broader digital engineering toolkit that also includes simulation, digital twins, physics-based modelling, optimisation and advanced data analytics. While AI often captures the headlines, these complementary technologies are already transforming how complex industrial systems are designed, validated and operated. Together, they allow engineers to move beyond intuition and experience alone, replacing uncertainty with evidence long before physical assets are built.

Reducing Business Case Risk with Simulation

One of the biggest challenges facing hydrogen projects today is the risk associated with bringing them online. As discussed in the previous chapters, financing is closely tied to confidence — and confidence is, in turn, shaped by how well risk can be understood and managed. If you follow developments in hydrogen as closely as I do, you'll notice that every quarter brings the announcement of new prospective projects. As discussed in earlier chapters, this growth is essential—transitioning to a hydrogen-based energy system requires replacing vast quantities of existing energy supply, measured in exajoules.

However, with this wave of new project announcements comes a more concerning, and arguably more important, statistic: how many of these projects actually reach Final Investment Decision (FID)?

The reality is concerning, in recent years—up to 2025—fewer than 10% of announced hydrogen projects have reached FID. Despite the urgency of global decarbonisation, momentum to push these projects forward often falls short, as economic reality intervenes.

Hydrogen technology remains relatively nascent, and significant uncertainty still surrounds its commercial and operational use cases. Investors are understandably cautious; few are willing to commit multi-million-euro investments to projects that might not deliver on their expected outcomes.

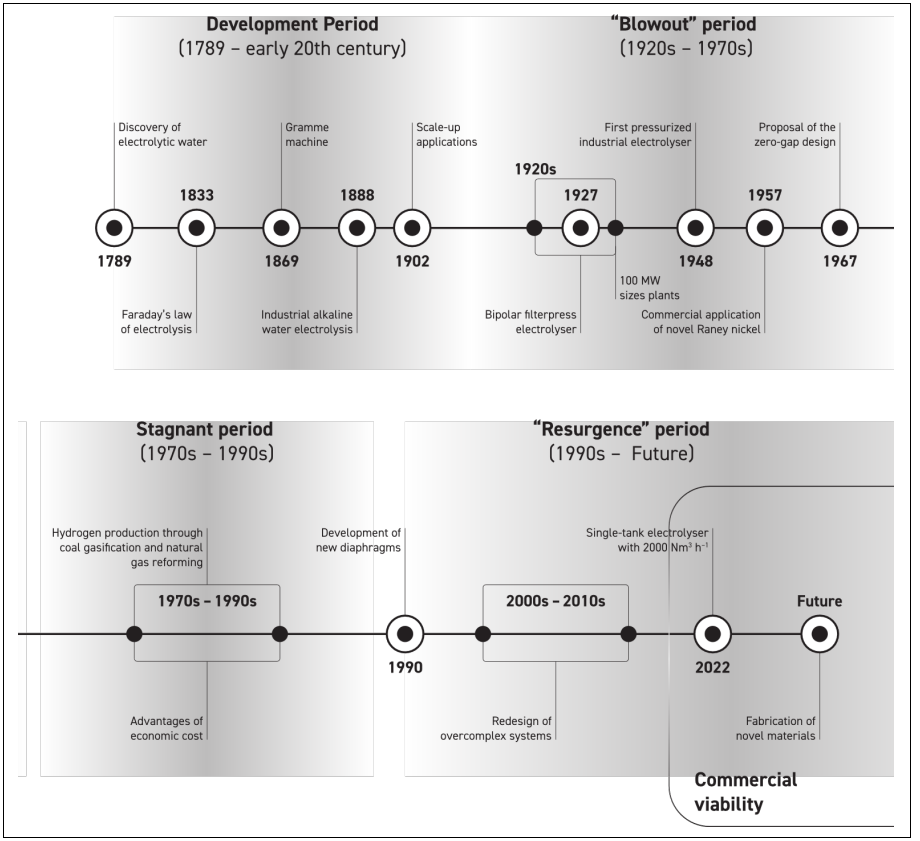
This is where digitalisation and simulation become valuable. By creating detailed digital models of hydrogen systems—before a single euro is spent on physical construction—developers can explore thousands of scenarios, test assumptions, and quantify risk in a controlled environment. Simulation enables project teams to visualise how design choices, technology selections, or market conditions will impact performance and economics over the project's lifetime. In doing so, it helps to replace uncertainty with data-driven insight, improving confidence in the business case and, ultimately, the likelihood of reaching FID.

Moreover, simulation allows decision-makers to explore "what-if" scenarios, stretching beyond technical risk into commercial and strategic dimensions. By modelling variations in electricity prices, hydrogen demand, or policy incentives, teams can quantify sensitivities and make more informed investment decisions. In this sense, digital tools do not simply mimic reality—they extend foresight, enabling planners to navigate uncertainty and translate complex, interdependent variables into actionable insight.

Digital Tools in Hydrogen Manufacturing

The value of digitalisation becomes clearer when we move from theory to application. Across the hydrogen value chain, simulation and AI can reduce risk, improve design, and accelerate learning. Let us start in the factory.

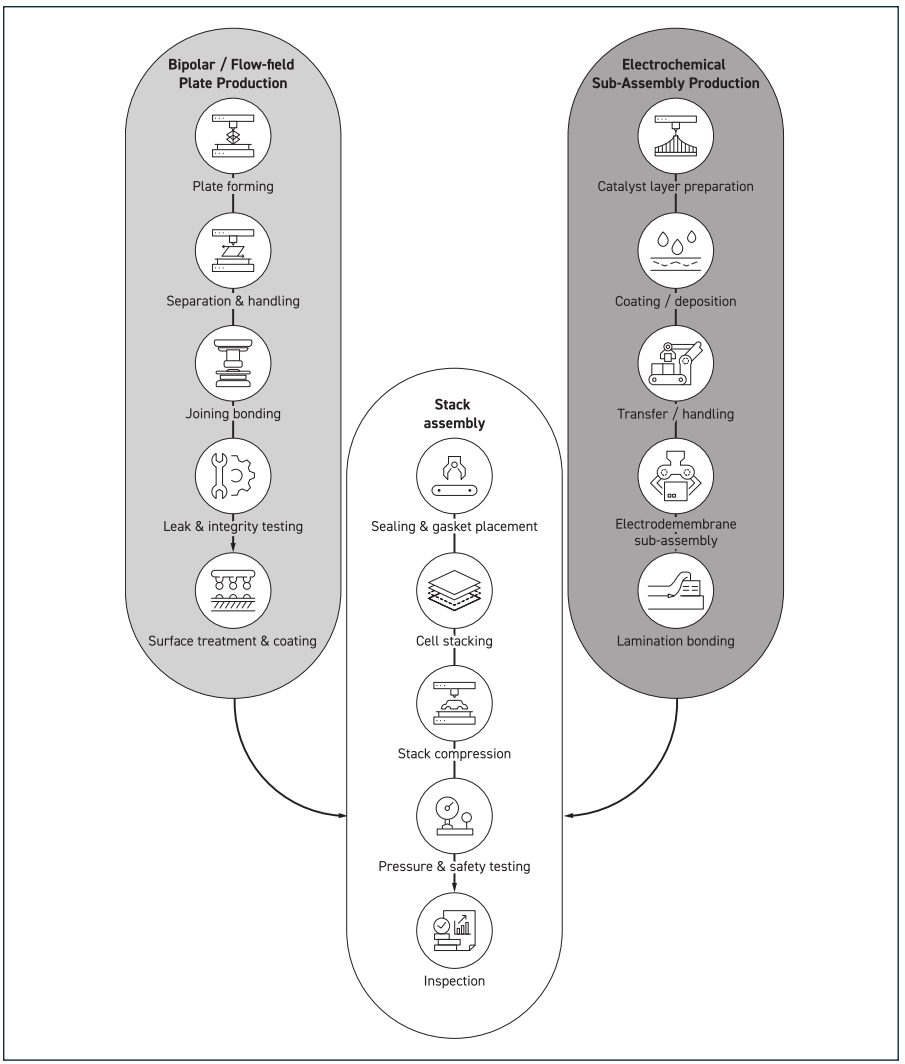
Earlier, I introduced you to the various electrolyser types—PEM, Alkaline, SOEC. Now, to understand why digitalisation is so important, it is necessary to grasp where we are in the lifecycle development of electrolysis.



Adapted from Xu et al. (2023), Rational Design of Hydrogen Evolution Reaction Electrocatalysts for Commercial Alkaline Water Electrolysis; additional information added and illustration redesigned for this book.

Understanding the maturity of electrolyser technology is only the first step. Just as the earliest automobiles required repeated testing, redesign, and adaptation, hydrogen infrastructure benefits from a similar iterative approach—but at a pace accelerated by digital tools. Simulation and AI allow engineers to experiment with new configurations virtually, revealing potential bottlenecks, predicting energy efficiency, and testing responses to fluctuating renewable inputs. In effect, the plant becomes a living laboratory where hypotheses can be tested safely and economically before committing capital, turning risk into manageable insight.

Even though the process of electrolysis has been known for hundreds of years, it has only been commercially viable in the last decade. To put that into perspective, think about the car you are driving today. The electrolyser of today is comparable to the very first vehicles that Carl Benz put on the road in Karlsruhe in the 1880s. Just as the modern car has gone through countless design stages and iterations to reach the level of sophistication and reliability we now take for granted, the same is true for the electrolyser and fuel cell industry. The product of today will not be the product being built in two, five, or even ten years from now.



As advancements in technology continue, both the product and the manufacturing process will evolve. This means that today's OEMs must be prepared to change and adapt their factories regularly to keep their products competitive. Of course, any pause in production comes with significant costs and risks—a single unplanned stoppage can have knock-on effects across the entire plant.

This is precisely where a digital twin becomes invaluable. By creating a digital twin of a plant as it exists today, an OEM can visualise and simulate changes before they are executed in the physical world. Replacing a single machine may seem straightforward, but it can have cascading effects on throughput, sequencing, or integration with other equipment. Running the change through a digital twin allows these issues to be identified and addressed in advance. It might reveal that additional machines need to be adjusted, recalibrated, or replaced. The result is that all necessary changes can be executed in a single, coordinated shutdown, rather than performing one change, restarting production, and discovering the next problem weeks later.

The value of a digital twin extends beyond single plants. By integrating fleet-wide data, operators can benchmark performance across multiple sites, learn from one facility to improve another, and dynamically allocate resources where they are most needed. This collective intelligence transforms isolated plants into a coherent, adaptive network capable of responding to supply fluctuations, maintenance events, or policy changes with unprecedented agility.

Digital twins, therefore, allow factories to remain agile, adaptive, and resilient. They are not simply a tool for efficiency—they are an essential mechanism for keeping pace with a rapidly evolving industry, ensuring that the hydrogen manufacturing process can continue to advance safely, cost-effectively, and competitively.

Everything we've discussed so far has focused on the manufacturer. But the real value of digitalisation extends far beyond the factory walls. Ultimately, these technologies only become meaningful when they help operators make better decisions in the field. Once electrolysers leave the production line and become part of larger energy systems, the challenge shifts from designing individual products to orchestrating entire infrastructures. It is here that digital twins, simulation and AI begin to demonstrate their greatest value—not by optimising a single machine, but by coordinating complete hydrogen ecosystems.

A Practical Use Case: Nürnberg Airport

Let's explore a specific example.

As I write this, I'm sitting in the Albrecht Dürer Airport in Nürnberg, Germany. Imagine that the airport decides to convert its entire ground fleet to operate on hydrogen. To make that transition, it would need an energy source, an electrolyser to produce hydrogen, adequate storage capacity, a refuelling system, and finally, a fleet of vehicles fitted with fuel cells. Clearly, this represents a major investment.

There are already digital tools capable of performing the electro-mechanical and thermodynamic calculations required to assess the feasibility of such a system. In this case, however, it's best to work backwards—from demand to supply.

The process might begin by analysing the vehicle fleet's operational data over a representative year. Understanding vehicle usage patterns allows us to estimate hydrogen consumption under real operating conditions. Once we know the required energy demand, we can select the appropriate fuel cell technology and calculate the volume of hydrogen needed to sustain operations.

From there, attention shifts to hydrogen production and supply. Since installing wind turbines at an airport is not feasible, solar power would likely be the preferred renewable source. Historical weather and irradiance data for Nürnberg could be used to size both the solar array and the electrolyser capacity required to meet demand.

Of course, several other factors must also be considered: Nürnberg experiences cold winters and variable sunlight; hydrogen has storage and leakage challenges; and seasonal variations could affect supply. A complete system model would therefore need to include energy storage, backup supply agreements with the grid, and contingency scenarios for reliability.

By simulating all these interactions digitally, project teams can quantify risk before committing capital. A few hundred thousand euros spent on simulation and analysis can prevent the loss of millions in a poorly performing real-world deployment. In essence, simulation helps replace uncertainty with evidence—strengthening the business case and increasing confidence on the path to FID.

Multi-Domain Dynamic Simulation for Hydrogen Systems

This is where the airport example becomes truly useful: not in modelling single pieces of equipment, but in understanding how the entire hydrogen system behaves as an integrated whole.

A hydrogen system is not a single process. It is electrical, mechanical, thermal, and fluid all at once — and each part influences the others. A change in solar input affects the electrolyser. That affects hydrogen production, which in turn affects storage pressure, refuelling availability, and ultimately vehicle operation.

Multi-domain simulation allows all of this to be captured in a single model.

In this case, the model might include:

- the vehicle fleet, with its varying demand profiles and usage patterns,
- the electrolyser, responding dynamically to fluctuating power input,
- storage and delivery systems — compressors, pipelines, valves, refuelling stations — with their pressure and flow constraints,
- and the supporting systems such as cooling, pumps, and energy management.

Bringing these elements together makes it possible to simulate how the system behaves over time, not just at a fixed point. You can see whether hydrogen production actually matches demand throughout the day, how the system responds to variable solar generation, and where bottlenecks begin to emerge — often in places that would not be obvious when looking at components in isolation.

This is where the real value appears. Instead of reacting to problems once the system is built, operators can anticipate them. Maintenance can be planned before degradation becomes failure. Control strategies can be adjusted before inefficiencies become embedded.

Over time, this turns the system from something that is operated into something that is continuously learned from. Each cycle adds information, improving performance and reliability rather than simply maintaining it.

It also allows different operating strategies to be tested before committing to them. What happens if demand spikes during peak hours? What if renewable input drops suddenly? How should the electrolyser respond — ramp, pause, or continue at reduced load? These are not theoretical questions; they directly affect efficiency, cost, and reliability.

By answering them in a virtual environment, the system can be designed with far greater confidence. Components can be sized correctly, control strategies can be validated, and risks can be identified early — before they become expensive problems in the real world.

Digitalisation in Product Design

As I've mentioned a few times already, today's commercial electrolysers are still relatively new to the market. This means that even though we have solid expectations about how they should perform, the reality is that long-term operational data is still limited. Until these systems have been running in the field for several years, no one can state with complete confidence exactly how they will behave under different conditions—be it performance degradation, efficiency drift, or maintenance needs.

This uncertainty poses a challenge for everyone involved. It keeps the cappuccino drinkers in R&D and product development awake at night as they try to refine designs and predict performance curves. It's also equally concerning for the Red Bull drinkers in operations, who are responsible for ensuring that these products deliver consistently once deployed in the field. Both sides are grappling with the same core issue: how do we design, test, and validate systems when the full operational picture is still emerging?

By embedding AI and simulation into product development, manufacturers can compress decades of learning into months. They can virtually trial new materials, cell designs, or stack geometries and immediately see the projected impact on efficiency, lifetime, and operational flexibility. In this way, digitalisation becomes a creative partner—one that can rapidly explore permutations far beyond the capacity of traditional prototyping, enabling faster innovation cycles and more resilient designs.

Advanced Process Simulation for Electrolyser Power Stacks

Up to this point, we have been looking at systems as a whole. Yet the same logic applies one level deeper — inside the electrolyser itself.

The challenge here is straightforward: today's electrolysers have not yet been operating at scale for decades. That means long-term behaviour — efficiency, degradation, lifetime — is still not fully understood. Yet these are exactly the parameters that determine whether a project succeeds or fails.

This is where detailed process simulation becomes critical.

Rather than relying solely on physical prototypes, engineers can build a physics-based model of the electrolyser stack itself. This model captures what is actually happening inside the system — the electrochemical reactions, heat generation, fluid flows, and how these interact under different operating conditions.

The advantage is not just accuracy, but speed of learning. Instead of building, testing, and redesigning hardware repeatedly, different configurations can be explored virtually. Changes to materials, geometry, or operating conditions can be tested quickly, allowing engineers to see how they affect efficiency, durability, and hydrogen output.

This becomes especially important when electrolysers are coupled to renewable energy. In reality, power input is rarely stable. Solar and wind introduce variability, meaning systems must cope with ramping, start-stop cycles, and partial loading. Simulation makes it possible to test these conditions in advance — not just whether the system works, but how it degrades over time.

Once these models are calibrated with real data, they evolve into something more powerful: a digital representation that can be used to predict long-term behaviour. This allows teams to estimate maintenance needs, understand degradation patterns, and make more informed decisions about design and operation.

The real value, however, lies in understanding trade-offs. There is no single “best” design — only balances between efficiency, cost, and lifetime. Simulation allows these trade-offs to be explored explicitly. What happens if efficiency is pushed higher? What does that do to degradation? How sensitive is the system to electricity price or operating temperature?

By answering these questions early, much of the uncertainty can be removed before physical systems are built. That reduces the need for costly trial-and-error, shortens development cycles, and increases confidence that the final design will perform as expected.

In practical terms, this means three things. It accelerates product development, ensures that components are properly integrated at system level, and provides a clearer understanding of how the electrolyser will behave over its lifetime.

Ultimately, it shifts development from a process of iteration in the field to learning in advance — where mistakes are far cheaper, and insights can be applied before they become risks.

Example Application: Designing a Renewable-Powered Electrolyser System

Consider an engineering team developing a new electrolyser system intended to operate alongside a solar photovoltaic installation. Before committing to physical prototypes, they build a detailed digital model of the power stack and its supporting components. Using local solar irradiance data, the team simulates how variable electricity supply affects hydrogen production, cell temperatures, and system efficiency throughout a typical year.

The simulation reveals that frequent start-stop cycles during cloudy periods lead to elevated thermal stresses in the stack, which could accelerate degradation. By adjusting the control logic and modifying the heat management system within the model, the team identifies an operating strategy that maintains efficiency while reducing mechanical and thermal strain.

They also use the model to evaluate storage requirements, determining how much buffer capacity is needed to maintain continuous hydrogen output during periods of low sunlight. Finally, an economic analysis—based on simulated performance, energy costs, and component lifetimes—helps quantify levelized hydrogen cost and investment risk.

Through this virtual testing process, the development team gains a clear understanding of how design decisions, environmental factors, and control strategies influence both technical and economic performance. By resolving many uncertainties digitally, they can make more informed design choices and present a stronger, data-backed business case for the project's next phase.

Digital Performance Management for Hydrogen Projects

The previous examples have focused mainly on individual systems and plants, however the same logic applies at a larger scale. As hydrogen projects grow from single facilities into megawatt- and gigawatt-scale assets — and, increasingly, into networks of sites across different regions — the challenge changes. It is no longer just about making one plant work well. It becomes about understanding and managing performance across an entire portfolio of assets.

A single facility is already complex. Multiple facilities, operating under different conditions and in different locations, quickly become difficult to compare, optimise, or even fully understand. Without a clear view, inefficiencies accumulate, performance drifts, and decisions become reactive rather than deliberate.

This is where digital performance management becomes important — not as an additional layer of complexity, but as a way of making that complexity visible.

At its simplest, it means bringing together data from across the system: production, storage, compression, energy input, and equipment behaviour. When that information is seen in one place, patterns begin to emerge. You can see how efficiently a plant is operating, where energy is being lost, and how equipment is really performing over time.

At a single-site level, this helps operators run the plant more effectively — adjusting schedules, identifying inefficiencies, and addressing maintenance before it becomes a failure. But the real shift happens when multiple sites are involved.

Now, performance can be compared. One plant can inform another. Practices that work well in one location can be applied elsewhere, while underperformance becomes visible rather than hidden. Over time, this creates a more consistent and predictable operation across the entire fleet.

Equally important is the ability to track stack degradation in electrolysis units. Over time, stacks lose efficiency, affecting output, energy use, and maintenance planning. When that degradation is monitored continuously, operators can intervene earlier, schedule maintenance more intelligently, and make better decisions about replacement and load allocation across multiple plants.

It also changes how decisions are made. Instead of reacting to issues after they occur, operators can anticipate them. Bottlenecks can be identified early, degradation can be tracked, and different operating strategies can be tested against real data. The result is not just better performance, but more informed decision-making.

There is also a practical dimension to this. As projects scale, expectations around reporting, efficiency, and compliance increase. Being able to demonstrate how a plant is performing — not just in theory, but in practice — becomes increasingly important, whether for internal management, investors, or regulators.

Ultimately, digital performance management is less about adding new capability and more about creating clarity. It allows operators to see what is happening across their assets, learn from it, and act on it.

At the level of whole energy ecosystems, this logic extends even further. AI can help coordinate interactions between hydrogen plants, renewables, storage, and demand-side consumers, turning scattered assets into a more responsive net-

work. In that world, hydrogen production stops being a series of isolated plants and becomes part of a wider, adaptive system.

As hydrogen infrastructure grows from individual projects into interconnected networks, that clarity becomes essential — not just for efficiency, but for confidence, resilience, and scale.

Extending Simulation to Ammonia and Methanol Production

Once hydrogen production is understood and optimised, the next step in many projects is its conversion into chemicals, such as ammonia or methanol. These act as energy carriers, storage media, or industrial feedstocks, and their production is tightly coupled to the hydrogen supply.

Simulation allows these downstream processes to be explored in the same way as hydrogen production itself. For ammonia, this means modelling the Haber–Bosch process, where hydrogen reacts with nitrogen under high temperature and pressure. For methanol, it involves reactions with carbon dioxide or carbon monoxide under catalytic conditions. In both cases, performance depends on reaction kinetics, heat management, and system constraints — all of which can be captured in a model.

The real advantage comes when these models are connected. By linking hydrogen production with chemical synthesis, the entire value chain can be studied as a single system rather than as separate steps.

This allows key questions to be addressed early:

- how hydrogen supply matches chemical demand under variable renewable input,
- how intermittent operation affects reactor behaviour and product yield,
- where operating conditions should be set to balance efficiency, cost, and stability,
- and how sensitive the system is to changes in input conditions or component performance.

Answering these questions before construction reduces both technical and commercial uncertainty. It allows reactor design, operating strategy, and system integration to be tested without relying on expensive trial-and-error in the field.

It also highlights one of the central challenges of hydrogen-based systems: variability. When hydrogen production is tied to renewable energy, supply is rarely constant. Simulation makes it possible to understand how that variability propag-

ates through the system — and what needs to be added, whether storage, buffering, or control strategies, to maintain stable operation.

In this sense, simulation bridges the gap between hydrogen production and its utilisation. It turns what would otherwise be a series of loosely connected processes into a coherent system that can be designed, tested, and scaled with far greater confidence.

Designing Fluid Systems: Pipelines and Tanking

Hydrogen is, as discussed earlier, a particularly elusive molecule — the Harry Houdini of the periodic table. It has a tendency to escape, leak, and behave in ways that make it both fascinating and challenging to handle. Whether in pipelines, storage vessels, compressors, or refuelling systems, understanding how hydrogen behaves dynamically is critical. This is where simulation becomes invaluable.

Fluid flow simulation allows engineers to model how gases and liquids move through complex piping networks. It captures pressure drops, flow rates, transient behaviour, thermal effects, and compressibility — all of which are particularly relevant when working with hydrogen.

In a hydrogen plant, this can be applied to:

- pipelines and distribution systems, to identify pressure losses and bottlenecks,
- compressors and valves, to understand behaviour during start-up, shutdown, or rapid demand changes,
- storage and refuelling systems, to ensure safe and reliable filling and delivery,
- and safety analysis, to understand how hydrogen may accumulate or disperse under different conditions.

The key advantage is visibility before anything is built. Instead of discovering issues during operation, engineers can test how the entire network behaves under realistic conditions. For example, they can assess how the system responds to fluctuating renewable input or sudden changes in demand from a vehicle fleet.

By creating a digital representation of the network, pipe sizes, valve settings, compressor schedules and storage capacity can all be adjusted in advance. “What-if” scenarios — a blocked line, partial storage, or unexpected demand spikes — can be explored safely in a virtual environment.

Returning to the Nürnberg Airport example, this becomes particularly clear. Simulation can model how hydrogen flows from production to storage and through to refuelling stations as vehicles move throughout the day. It can identify pressure drops when multiple refuelling points operate simultaneously, assess how compressors respond to fluctuating supply, and ensure that storage remains sufficient during peak demand.

It also allows edge cases to be tested — situations that are difficult, or risky, to explore in reality. By doing so, engineers can design systems that are not only efficient, but resilient.

Ultimately, this reduces the likelihood of costly redesigns, improves operational reliability, and provides confidence that the hydrogen infrastructure will perform as expected under real-world conditions.

Digital tools are often dismissed as a "nice to have." In fact, they provide a clear advantage when applied alongside human judgement — not replacing it but strengthening it.

Across this chapter, we have seen how simulation can be used to optimise production, explore chemical conversion, and understand system-level behaviour. But their real value lies in something more fundamental: reducing uncertainty.

That said, digitalisation must be applied selectively. Introducing complex tools too early can add cost and slow progress. These capabilities are powerful, but they need to be used at the right stage, in the right way, if they are to remain an enabler rather than a burden.

When applied well, they allow projects to be tested before they are built — turning assumptions into evidence. Design decisions can be challenged, operational strategies explored, and risks understood long before significant capital is committed.

This changes the nature of hydrogen projects. Instead of relying on expectation, they can be developed on the basis of insight.

Digitalisation, in this sense, is not just about efficiency. It is a way of moving faster with greater confidence — bridging the gap between what can be imagined and what can be delivered.

The Foundations Beneath the Future

Beneath all this technological progress lies an essential truth: none of it works without shared foundations. Digitalisation only accelerates hydrogen if the data is secure, the models are trusted, and the systems can communicate.

Standards will determine whether a digital twin built by one OEM can interface with a model built by another. Interoperability will decide whether multi-vendor plants operate as coherent systems or awkward, disconnected collections of equipment. Cybersecurity will define whether hydrogen assets become resilient infrastructure or vulnerable targets. Additionally, data governance will shape who controls the insights, how they are shared, and how they are validated.

As regulators, insurers, financiers, and operators increasingly depend on digital twins and AI-assisted analysis, a new form of certification will emerge — one that validates not only the physical asset, but the digital model used to design, operate, and predict its behaviour. This is not an administrative detail; it is the social licence that allows digitalisation to scale responsibly.

Hydrogen's digital future will be built not only on algorithms, but on trust.

16. Action Today

"In theory, there is no difference between theory and practice. In practice, there is." — Yogi Berra

Why "Today" Is the Decisive Window

The energy transition is often framed as a challenge of the future. Targets are set for 2040 or 2050, roadmaps stretch across decades, and difficult decisions are routinely deferred to later stages of the journey. That framing is misleading. The outcome of the transition will not be determined in 2050, but by what is decided — and built — in the years immediately ahead.

At this stage, the central risks are no longer technical. They are institutional. Delay fragments momentum. Fragmentation multiplies cost and complexity. Misaligned incentives reward optimisation within silos rather than performance at the system level. None of these failures appear dramatic in isolation, yet over time they compound, locking in pathways that become increasingly difficult to unwind.

One of the recurring themes throughout this book has been the tendency to compare ourselves with others. We look to China's industrial planning, India's manufacturing ambitions, Europe's evolving policy framework, or the policies of other regions, often asking whether they have chosen the better path. These comparisons are valuable because they broaden our perspective and allow us to learn from different approaches. They should not, however, become a distraction.

No nation can determine the long-term strategy of another. Every nation can, however, determine its own. The real objective is therefore not to replicate another country's approach, but to develop a clear, consistent strategy that reflects domestic strengths, provides confidence for industry, and remains sufficiently stable to endure beyond individual political cycles. Long-term competitiveness is built through sustained execution rather than continual comparison. Energy transitions are measured in decades, not election cycles, and strategies

must be resilient enough to outlast both changing governments and changing headlines.

The 2020s therefore represent a decisive window. Infrastructure choices made now will shape energy flows for decades. Regulatory frameworks established in this period will influence capital allocation long after political cycles have turned. Skills developed — or neglected — today will determine what can realistically be delivered later, regardless of how ambitious future targets may appear.

This decade is not about discovering new answers, but about committing to known directions. Progress depends less on innovation than on coordination: between policy and markets, between power systems and industrial demand, and between near-term actions and long-term intent. Where alignment is achieved, scale follows. Where it is not, complexity accumulates and momentum dissipates.

The task of this chapter is therefore not to restate the case for hydrogen or the wider energy transition. That case has already been made. Its purpose is more immediate and more demanding: to examine what must be done now, by those in a position to act, if the transition is to move from aspiration to execution.

Invest — Move from Experiments to Platforms

Investment has not been absent from the hydrogen transition. Capital has flowed steadily into pilots, demonstrations, and early commercial projects across multiple regions. What has been missing is not interest, but intent at scale. Much of today's investment activity remains framed as experimentation — cautious, fragmented, and optimised to limit downside rather than to enable learning at the system level.

That posture was understandable in the early stages of the transition. It is no longer sufficient. At this point, incrementalism has become its own form of risk. Projects designed primarily to prove narrow technical feasibility do little to reduce the uncertainties that now matter most: integration, replication, supply-chain maturity, and bankability at scale. Those uncertainties cannot be resolved in spreadsheets or pilot programmes. They can only be resolved by building platforms rather than one-off assets.

The distinction is critical. Platforms create learning effects that individual projects cannot. They force standardisation, expose real operational constraints, and allow costs to fall through repetition rather than bespoke optimisation. This is where earlier discussions of modularity become relevant not as a design philosophy, but as an investment logic. Modular systems — whether in electrolysis,

balance-of-plant, or supporting infrastructure — reduce construction risk, shorten learning cycles, and allow capacity to scale through replication rather than reinvention. Their value lies less in technical elegance than in speed, predictability, and repeatability.

For hydrogen, this shift implies a clear change in investment priorities. Capital must increasingly flow toward infrastructure and capability, not just production assets. Manufacturing capacity, storage, ports, pipelines, grid connections, and system integration matter more to long-term outcomes than the marginal efficiency gains of any single plant. Infrastructure built ahead of full utilisation may appear inefficient in the short term, but history shows that it is often the only way to unlock demand, reduce risk, and enable scale.

Digitalisation plays a similar, if less visible, role. Digital systems increasingly act as enabling infrastructure, allowing assets to be coordinated, optimised, and maintained at scale. Their contribution is not headline-grabbing, but it is material: reducing operational uncertainty, improving asset performance, and making complex systems investable. As hydrogen systems grow in size and interdependence, digital capability becomes a prerequisite for bankability rather than a discretionary add-on.

Equally important is the need to invest deliberately on the demand side. Hydrogen does not scale because supply exists; it scales because credible offtake exists. Industrial anchors — in steel, chemicals, refining, shipping, and dispatchable power — provide the volume and predictability required for financing. Without them, projects remain exposed to market volatility, policy shifts, and short-term pricing dynamics that deter long-term capital.

This is where public intervention has a legitimate and focused role. Mechanisms such as hydrogen banks, contracts for difference, and offtake guarantees are best understood not as subsidies, but as tools to reshape early-stage risk. Their effectiveness lies less in the amount of public capital deployed than in their ability to mobilise multiples of private investment by stabilising revenues and reducing uncertainty. Used well, they accelerate scale; used poorly, they fragment it.

This shift also requires a different attitude to early economics. The first generation of platforms will not be cost-optimal, and they do not need to be. Their value lies in what they enable: mature supply chains, skilled workforces, and markets that function with less intervention over time. Waiting for fully competitive economics before committing capital may feel prudent, but it delays the very conditions required to achieve them.

The hydrogen transition will not be built through a collection of optimised projects assembled after the fact. It will be built through deliberate investment in platforms that allow repetition, coordination, and learning. The question facing investors today is not whether hydrogen will play a role in the future energy system, but whether they are prepared to invest in the structures that allow that role to emerge at scale.

Collaborate — Break the Silos

By the late 2020s, the hydrogen landscape is crowded with advanced plans. Electrolysers are technically ready. Industrial off takers express interest. Infrastructure routes are mapped. Policy frameworks exist in draft or in law. However, across many regions, projects struggle to move from intention to execution.

The problem is rarely a single missing component. More often, it is timing. Power generation waits for offtake certainty. Industrial demand waits for infrastructure. Infrastructure waits for committed volume. Each actor's position is rational when viewed in isolation. Taken together, they cause delay.

This is not a failure of ambition or competence. It is a failure of coordination. Hydrogen sits between sectors that were never designed to move in step, and it exposes the cost of treating interdependence as someone else's problem. At this stage of the transition, collaboration is no longer a matter of goodwill or partnership. It is a system requirement.

A common failure mode is sequential rather than parallel development. Power assets are planned without clarity on hydrogen offtake. Hydrogen production is developed without certainty on transport or storage. Infrastructure is delayed until demand is proven, while demand remains constrained by the absence of infrastructure. Each decision appears defensible in isolation. Together, they produce inertia.

Another failure mode lies in misaligned time horizons. Infrastructure requires long-term commitment; industrial investment depends on medium-term visibility; financial markets often operate on shorter cycles. Without mechanisms to align these timelines, projects stall in the gaps between them. The result is a proliferation of advanced plans that struggle to reach final investment decision, despite apparent readiness.

Collaboration failures are often framed as cultural or organisational problems. In practice, they are structural. Effective collaboration does not require consensus

or central control. It requires interfaces that allow decisions made in one part of the system to be credible to others. Standardised contracts, interoperable technical standards, shared planning assumptions and transparent risk allocation matter far more than informal partnerships or memoranda of understanding.

The role of system integrators is therefore critical. Engineering firms, network operators, and platform developers increasingly act as translators between domains, aligning technical design with commercial reality and regulatory constraint. Where this capability is weak, projects become bespoke and fragile. Where it is strong, complexity becomes manageable and repeatable.

Public actors also have a role, but not by coordinating every decision. Their most valuable contribution is to reduce uncertainty at the interfaces: aligning permitting timelines, clarifying grid access, harmonising standards, and signalling long-term intent consistently. Where public policy adds layers of complexity or frequent redesign, it amplifies fragmentation rather than resolving it.

The risk of failing to collaborate is not immediate collapse, but gradual underperformance. Costs remain high because learning is isolated. Infrastructure is underutilised because timing is misaligned. Capital demands higher returns because risks remain opaque. Over time, these inefficiencies erode confidence and slow deployment, even as underlying technologies improve.

Breaking silos does not require perfect coordination. It requires sufficient alignment to allow decisions to be made in parallel rather than in sequence. Hydrogen will only function as the connector it is meant to be if interdependence is treated as a design constraint, not an inconvenience. Where collaboration is optional, the system fragments. Where it is embedded from the outset, scale becomes possible.

Regulate — Enable Speed

Regulation is often treated as the primary driver of the energy transition. In practice, its role is narrower and more decisive. Regulation cannot build infrastructure or operate systems. What it can do is determine whether rational actors move — or wait.

At this stage of the transition, the central regulatory challenge is not ambition, but durability. Frequent redesigns, shifting eligibility criteria, and short political time horizons create hesitation even among willing participants. For capital-intensive infrastructure, uncertainty is more damaging than imperfect policy. Stability enables commitment; volatility rewards delay.

The most effective regulatory frameworks focus on direction rather than outcomes. They define boundaries, signals, and guardrails, while allowing markets and industry to determine optimal configurations. Where regulation attempts to prescribe technologies or architectures, it tends to slow deployment and lock in assumptions that quickly become outdated.

Revenue stability is a recurring fault line. Hydrogen projects are exposed not because costs are inherently prohibitive, but because future cash flows remain uncertain. Regulatory approaches that stabilise revenues — regardless of their specific form — reduce risk more effectively than headline support levels. Their value lies in predictability, not generosity.

Standards and definitions play a similarly critical role. Fragmented certification schemes, inconsistent safety requirements, and incompatible technical specifications introduce friction that compounds across borders and supply chains. Interoperability matters more than uniformity. Capital flows more easily when assets, fuels, and guarantees do not need to be reinterpreted at every interface.

Permitting and planning remain persistent constraints. In many regions, approval timelines are misaligned with stated policy urgency. Accelerating these processes does not require lowering environmental or social safeguards. It requires recognising that delay itself carries cost — economic, climatic, and strategic.

A final challenge lies in how regulation treats infrastructure built ahead of demand. Energy systems have rarely scaled by waiting for full utilisation before investing. Frameworks that penalise early infrastructure in the name of short-term efficiency risk reinforcing the very fragmentation they seek to avoid.

Regulation cannot remove risk, nor should it attempt to. Its purpose is to ensure that risk sits where it can be managed, rather than where it paralyses decision-making. In the hydrogen transition, regulation succeeds not when it controls outcomes, but when it enables speed. The question is no longer how much to regulate, but whether regulation is allowing decisions to be made in parallel rather than in sequence.

Communicate — Precision Over Promotion

Communication has played an enabling role in the early stages of the hydrogen transition, helping to surface interest, align stakeholders, and secure initial support. As deployment moves closer to scale, however, the function of communication changes. The priority is no longer visibility, but precision.

At this stage, the primary risk is not silence, but overstatement. Hydrogen has often been framed in absolute terms — as inevitable, universal, or imminent — despite timelines and constraints that are necessarily uneven. When expectations run ahead of delivery, credibility is weakened and decision-making becomes more cautious rather than more confident.

Effective communication therefore requires discipline. This means being explicit about trade-offs, sequencing, and limits. It means acknowledging where hydrogen is well suited, where alternatives dominate, and where uncertainty remains. Just as importantly, it means resisting the temptation to promote solutions outside the contexts in which they are most effective. Deploying the right technology at the wrong scale, in the wrong sector, or at the wrong point in the transition can slow progress rather than accelerate it — not because the technology is flawed, but because expectations and system readiness are misaligned. Precision does not slow progress; it reduces friction by aligning ambition with reality.

This matters particularly for coordination across the system. Inconsistent language around definitions, standards, and performance — whether between policymakers, developers, or financiers — creates avoidable complexity. Misunderstandings propagate quickly into contracts, permitting processes, and investment decisions. Shared terminology and clear framing are not cosmetic; they are operational necessities.

Communication also shapes risk perception. Capital-intensive infrastructure depends on confidence that commitments are durable and signals are consistent. Overpromising, shifting narratives, or reframing objectives midstream increases perceived risk and raises the cost of capital. Conversely, measured, technically grounded communication supports longer-term alignment, even when progress is incremental.

At this stage of the transition, communication should do less persuading and more aligning. Its role is not to inspire belief, but to support execution by reducing ambiguity and maintaining credibility across the system. The work of generating urgency belongs elsewhere. The work here is simpler, and harder: to speak accurately, consistently, and with restraint.

Action Depends on Where You Sit

By this point, it should be clear that no single actor controls the energy transition. Progress does not hinge on one decisive intervention, but on a series of aligned decisions made by actors with very different incentives, constraints, and time horizons. Each sees only part of the system. Each acts rationally within their own

frame. The challenge is not to eliminate these differences, but to take responsibility for them.

The more useful question, then, is not what should happen in theory, but what it is rational to do today, given where you sit.

If I were a Senior Policymaker

The central responsibility of policymakers at this stage is not to perfect policy, but to make direction durable and capability cumulative. Most of the necessary instruments already exist. What remains fragile is confidence — not only that rules will persist, but that policy is sufficiently coherent to align supply, infrastructure, and demand rather than treating them as separate problems.

This requires policy that looks across the market, not just at individual projects. Building production capacity alone is not enough if credible offtake is absent. At present, capital is often available to build assets, yet projects stall because demand remains uncontracted or insufficiently supported. The role of policy is therefore not simply to accelerate build-out, but to create conditions in which supply and demand can move together — through consistent frameworks that support offtake, revenue stability, and infrastructure in parallel.

Equally important is the need to reduce unnecessary complexity. Much of today's regulatory architecture was designed for energy systems that evolved over decades, not for transitions that must be delivered within them. Policymakers should continually ask whether existing regulations are helping the future system emerge or inadvertently protecting assumptions from the past. Permitting processes, grid access procedures, certification schemes, and planning frameworks must maintain appropriate safeguards, but they must also recognise that speed has become a strategic requirement. Complexity rarely stops projects outright; more often it delays them until opportunities disappear.

That means looking beyond legislation alone. Universities, technical institutes, and vocational systems are not peripheral to the transition; they are part of its infrastructure. Energy systems scale only as fast as the people who design, permit, build, operate, and regulate them. Supporting education and training in engineering, chemistry, physics, data science, and systems integration is therefore not an educational add-on, but a strategic necessity.

This is a challenge that I suspect many developed economies are underestimating. A few years ago, my son was selecting his subjects at school. Looking across the choices being made by his year group of roughly 120 students, I was struck by

how heavily the cohort gravitated towards economics and business-related pathways. Approximately 90% followed that route, while only a small minority chose the natural sciences and engineering-focused subjects.

There was one encouraging observation. Among those choosing science-based pathways, a significant proportion were female. For industries that have historically struggled with diversity, that is undoubtedly positive. Yet what stayed with me was not the gender split, but the overall numbers. If we are serious about building the infrastructure described throughout this book, we will need far more scientists, engineers, technicians, electricians, software specialists, and skilled tradespeople than we are currently producing.

The energy transition is ultimately a physical transformation of society. It requires people to design electrolysers, reinforce electrical grids, construct pipelines, develop new materials, optimise industrial processes, write software, operate complex facilities, and maintain them for decades afterwards. We often speak about gigawatts, billions of euros of investment, and national strategies, but every one of those ambitions ultimately depends on people with the right skills.

A shortage of capital can be addressed by markets. A shortage of technology can often be solved through innovation. A shortage of skilled people is far harder to overcome because it cannot be manufactured overnight. The engineers needed in 2035 are sitting in classrooms today. If we fail to develop that pipeline of talent, then many of the targets and roadmaps discussed throughout this book will remain aspirations rather than achievements.

The transition therefore requires not only investment in infrastructure, but investment in human capital. In the long run, the most important energy infrastructure a country possesses may not be its pipelines, ports, or power stations, but its ability to educate the people who will build the next generation of them.

The history of the nuclear industry offers an important lesson in this regard. During the second half of the twentieth century, nuclear power benefited from sustained programmes of research, engineering and construction that created deep reservoirs of expertise across universities, national laboratories, equipment manufacturers and operating companies. When investment slowed, however, much of that capability gradually dispersed. Experienced engineers retired, specialist suppliers disappeared, research programmes contracted, and practical knowledge—much of it accumulated through decades of experience rather than written documentation—became increasingly difficult to replace. Today's renewed interest in nuclear therefore faces not only the challenge of building new

reactors, but of rebuilding an industrial ecosystem that has been allowed to diminish over time.

The same lesson applies across the wider energy transition. Industrial capability cannot simply be switched off and then switched back on when priorities change. Electrolysers, fuel cells, grid technologies, digital engineering and many other strategic industries will only mature if investment in skills, supply chains and practical experience remains continuous. Technologies evolve through sustained application as much as through scientific discovery. If continuity is lost, future generations are forced not simply to continue the journey, but to relearn parts of it.

Policymakers should prioritise stability over optimisation. Resisting the temptation to constantly redesign frameworks in response to political cycles or early imperfections is critical. For capital-intensive systems, consistency matters more than elegance. An imperfect policy that endures — and that connects projects to markets — will mobilise more investment than a theoretically superior one that is repeatedly revised. Industry can adapt to almost any set of rules; what it struggles to adapt to is a set of rules that change faster than investment cycles.

Finally, policymakers must accept uneven progress. Regions will move at different speeds. Some projects will fail. This is not evidence of policy failure, but of learning. The greater risk lies in waiting for uniformity before acting.

If I Were the Head of a Large EPC or System Integrator

For EPCs and system integrators, the defining challenge of the coming decade is not winning individual projects but building repeatable delivery capability. Hydrogen will not scale through bespoke engineering alone. It will scale through modularisation, standardisation, and disciplined execution across portfolios of projects.

This is where modular design philosophy and digitalisation intersect. Modular systems reduce construction risk, shorten schedules, and enable learning through repetition. Digital tools — from design automation and digital twins to construction sequencing and operational optimisation — turn that modularity into a force multiplier. Together, they allow complexity to be managed rather than reinvented each time.

Crucially, this also changes the business model of integration itself. As delivery becomes more standardised and digitally enabled, value shifts away from one-off engineering margin toward lifecycle performance, availability, and repeatable

execution. EPCs that continue to rely solely on bespoke project models risk being undercut by those willing to evolve toward platform-based, service-oriented, or performance-linked roles.

Integrators sit at the fault line between technical ambition and commercial reality. Translating policy intent and industrial demand into configurations that supply chains can deliver, financiers can underwrite, and operators can run reliably is now a core competence. Where this translation is weak, projects become fragile. Where it is strong, scale becomes possible.

If I Were an Industrial Incumbent

For industrial players, hydrogen is not primarily a climate decision. It is a competitiveness and resilience decision. The question is not whether decarbonisation will matter, but whether it will be shaped early — or imposed later under less favourable conditions.

At this stage, the most rational strategy is optionality. That means engaging early enough to learn, understanding hydrogen supply chains, participating in pilot offtake arrangements, and preparing assets for integration — without committing prematurely to full conversion.

Digitalisation materially accelerates this learning. Simulation, system modelling, and data-driven optimisation allow technologies, operating modes, and supply arrangements to be tested virtually before they are locked in physically. This shortens learning cycles and reduces the cost of early mistakes. In a transition defined by sequencing and timing, the ability to learn faster than competitors can be as decisive as the technology ultimately selected.

Closer collaboration with EPCs and technology providers further compounds this advantage. Shared digital models, aligned planning assumptions, and early integration of operational constraints reduce later rework and delay. Industrial actors that treat hydrogen purely as an external input risk discovering too late that system design decisions have already been made around them.

Delay is often framed as prudence. In systems undergoing structural change, it is also a choice with consequences.

If I Were a Long-Term Investor or Infrastructure Financier

For long-term investors, hydrogen presents a familiar paradox. The underlying need is structural and durable, yet early projects appear complex, heterogeneous,

and exposed. The temptation is to wait for clarity. The risk is that clarity only emerges once platforms are built.

The most constructive role for long-term capital is therefore not to seek technological certainty, but to participate in risk removal. This means favouring platforms over isolated assets, aggregation over fragmentation, and structures that allow learning and standardisation over time.

As with renewables, the invest ability of hydrogen will ultimately depend less on the technology itself than on the emergence of stable, repeatable business models that convert technical performance into predictable cash flows. Investors who engage early can help shape those models. Those who wait may find that the most attractive risk-adjusted opportunities have already been absorbed into regulated or utility-like structures.

Patience remains an advantage. But patience should not be confused with distance.

If I Were a Young Engineer or Student

For those entering the system, the transition will not be defined by a single technology or sector. It will be defined by integration. The most valuable skills will sit at the intersections: between power and molecules, between physical systems and digital tools, between design and operation.

That begins with fundamentals. Strong grounding in mathematics, physics, chemistry, and engineering remains irreplaceable. To that, fluency with data, coding, and digital systems increasingly adds leverage. A generation comfortable moving between core science and digital capability will shape the transition more decisively than those specialising narrowly in any single technology.

I am conscious here of writing not only in the abstract, but personally. If I were advising a young engineer or student — and I mean you here, Fynn — I would not suggest trying to predict which technology "wins". I would suggest building the skills that allow you to work across systems, adapt as they evolve, and make complex things function reliably at scale.

The transition will be built less by visionaries than by people willing to work within constraints, improve imperfect systems, and keep them running day after day. That work rarely makes headlines. It does, however, endure.

What Must Be True by 2030

By 2030, the direction of the energy transition will be largely set. Not because targets will have been reached, but because the structures that determine what is possible will either be in place or foreclosed. The question is not whether progress has occurred, but whether it has become self-reinforcing.

Several conditions must therefore be true.

First, hydrogen must have moved decisively beyond pilots. Multiple large-scale hydrogen hubs must be operating, not merely announced or permitted. These hubs should integrate production, storage, transport, and demand, demonstrating that hydrogen can function as part of a system rather than as a collection of projects. Learning must be visible in falling costs, shorter delivery timelines, and increasing standardisation.

Second, credible demand must be anchored. At least one major industrial value chain in each leading region should have committed hydrogen at scale — not as an experiment, but as a core input. Long-term offtake arrangements must be sufficiently common that they are no longer treated as exceptional. Without this, supply will remain fragile and investment hesitant.

Third, hydrogen infrastructure must be built ahead of full utilisation. Pipelines, storage, ports, and grid connections should exist with surplus capacity, reflecting confidence in future demand rather than reaction to past volume. Energy systems have never scaled by waiting for perfect utilisation. If infrastructure is still being delayed pending certainty, fragmentation has already won.

Fourth, regulatory frameworks must be stable and interoperable. Definitions, certification schemes, and market rules do not need to be identical across regions, but they must be compatible. Investors and operators should be able to move capital, equipment, and molecules across borders without renegotiating fundamentals at every interface.

Fifth, capital must be flowing at infrastructure scale. Financing structures should increasingly resemble those used for other energy networks, with declining risk premiums and growing participation from long-term institutional investors. If hydrogen assets are still treated primarily as speculative or experimental, risk has not been sufficiently addressed.

Finally, failure must be visible as well as success. Some projects will not proceed. Some assumptions will prove wrong. By 2030, this should be evident — not hidden. A system that cannot tolerate failure cannot learn fast enough to succeed.

If these conditions are in place, the transition will no longer depend on belief alone. Momentum will be embedded in infrastructure, contracts, skills, and institutions. If they are not, no amount of long-term ambition will compensate. The decade will have spoken.

Responsibility Has Shifted

By this point, the contours of the challenge should be clear. The energy transition no longer hinges on discovery or proof. The technologies exist. The pathways are visible. The constraints are understood. What remains unresolved is not whether hydrogen can work, but whether institutions, markets, and societies are willing to organise themselves to let it.

Responsibility has therefore shifted. It no longer sits primarily with inventors or advocates, but with those who shape systems: policymakers who set durable direction, companies that choose how to invest and collaborate, financiers who decide where risk resides, and institutions that determine how quickly capability can accumulate. None of these actors can deliver the transition alone. Together, they determine whether progress compounds or stalls.

The most consequential decisions are rarely dramatic. They are embedded in permitting timelines, contract structures, standards, curricula, and investment mandates. Taken individually, they appear incremental. Taken together, they define the pace and shape of the transition long before outcomes are visible in emissions data or infrastructure maps.

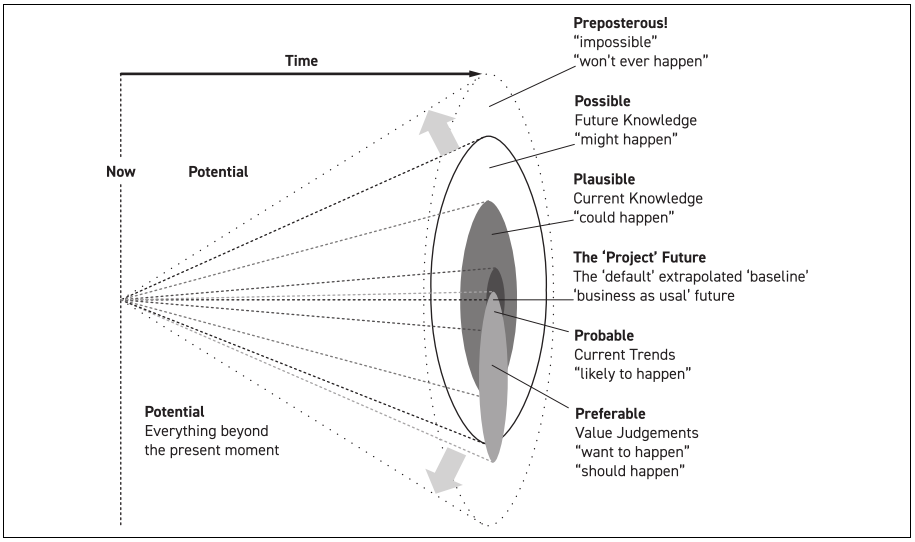
This is why action in the present decade matters disproportionately. Long-term targets express intent. Near-term decisions create reality. Once infrastructure is built, skills are trained, and capital is committed, trajectories become difficult to reverse.

History will not judge this transition by the ambition of its announcements or the sophistication of its plans. It will judge it by whether early actions were decisive enough to make progress self-sustaining. The question is no longer what we aspire to build, but whether we are prepared to organise ourselves to build it — now.

17. Hydrogen in 2050: Three Possible Futures

"The future is already here — it's just not evenly distributed."
— William Gibson

"Prediction is very difficult, especially if it's about the future."
— Niels Bohr



The futures explored in this chapter are not predictions, but trajectories. All sit within the range of what is possible given current knowledge and constraints.

Aviation's Hydrogen Turning Point

The First Quiet Flight

Some turning points announce themselves with spectacle — headlines, ceremonies, and the language of revolution. Others arrive quietly, recognised only later, once their consequences have had time to unfold. In the energy transition, it is often the second kind that matters most.

On a cold January morning in 2023, the runway at Cotswold Airport looked almost unremarkable. Mist clung to the surrounding fields, the low winter sun glinted off dewy grass, and the familiar sounds of small aircraft drifted across the airfield. It could have been any ordinary morning in the English countryside.

But on this particular day, something was different.

At the edge of the apron stood a Dornier 228 — a workhorse regional aircraft, familiar to pilots and engineers alike. Yet this aircraft carried a quiet departure from a century of convention. Mounted in its left wing was a hydrogen-electric powertrain developed by ZeroAvia, preparing to power a commercial-scale aircraft in flight for the first time.

There were no grandstands, no ceremonial speeches, no live broadcast. Just a small team of engineers, a pilot, a plane, and a shared understanding that aviation was about to cross a line it had approached for decades but never quite reached.

Inside the hangar, screens glowed with telemetry. Cooling systems hummed steadily. Pressurised hydrogen sat contained and inert, waiting to be converted into electricity. The system had been tested, simulated, and rehearsed countless times. But flight is unforgiving. Nothing substitutes for lift, vibration, and air beneath the wings.

The pilot climbed into the cockpit and began the pre-flight checks. Switches were set, gauges confirmed, routines followed — the familiar choreography of aviation. Yet for those watching, the difference was unmistakable. One engine would operate without combustion. It would draw oxygen from the air and release nothing but heat and water.

When the aircraft taxied, the left propeller turned with a controlled, almost subdued hum. At the runway threshold, the aircraft paused briefly. Brakes set. Systems aligned. One wing held the past. The other held something new.

Then the brakes were released.

Acceleration felt normal. The familiar growl of conventional propulsion blended with the quieter tone of the electric motor. Nothing about the moment demanded attention — until the aircraft lifted from the runway.

There was no dramatic transition, no visible disruption. Just the irreversible fact of flight. A real aircraft, in real airspace, powered in part by hydrogen fuel cells. Below, engineers exhaled. Some smiled. Others simply watched.

The aircraft circled the airfield, telemetry confirming what the eye already suggested: stability, consistency, control. The hydrogen system behaved not as an experiment, but as a dependable member of the propulsion system. Ten minutes later, the Dornier returned and touched down lightly on the runway.

Applause followed — restrained, sincere, and reflective. It marked not a finished solution, but a threshold crossed. Hydrogen had left the test bench and entered operational reality.

This flight did not solve aviation's challenges. Storage, certification, infrastructure, scale — all remained unresolved. But something fundamental had changed. The question was no longer whether hydrogen-electric flight was possible, but how far it could go, and how quickly.

Across the industry, attention shifted. Concepts became pathways. Timelines shortened. What had once been framed as distant ambition began to look like a ladder of progression: small aircraft first, then regional fleets, then broader integration into aviation's future architecture.

Looking back, the moment already carries historical weight — not because it was loud or dramatic, but because it was quiet and conclusive. A plane took off using hydrogen-electric propulsion. And the world, whether it realised it or not, moved with it.

The year 2050 often sits in an awkward space in discussions about the energy transition. It is close enough to feel tangible — well within the working lives of today's decision-makers — yet far enough away that it can still feel abstract. That distance allows both complacency and fantasy to creep in; the assumption that either everything will somehow work itself out, or that salvation will arrive through a single, transformative breakthrough.

This chapter takes a different approach. Rather than presenting a single vision of the future, it explores three plausible pathways to 2050. None of them is utopian, and none of them is dystopian. Each reflects a coherent outcome of choices already being made today, extrapolated forward through policy, markets, and

institutions. The differences between them are not matters of belief or optimism, but of coordination, consistency, and speed.

Crucially, the technologies required for all three futures already exist in some form today. Electrolysis, renewables, storage, hydrogen transport, fuel cells, and system integration are no longer theoretical. What separates these futures is not invention, but alignment: the ability of governments, industry, finance, and society to move in step rather than in parallel.

The scenarios that follow should not be read as predictions. They are tools — ways of visualising the consequences of different trajectories, and of understanding how small decisions made repeatedly over time compound into fundamentally different outcomes. In that sense, the question is not which future is possible, but which one is being actively chosen.

Scenario 1: Business as Usual — Fragmented Progress

In this future, hydrogen exists in 2050 — but it never quite becomes what it was meant to be. Progress has occurred, sometimes impressively so, yet it arrives unevenly and without coherence. The system moves forward, but not together. As a result, hydrogen remains an important component of the energy transition without ever becoming a defining one.

On the supply side, electrolysis has scaled, but largely within isolated industrial hubs rather than as part of a connected energy system. Projects have been developed where conditions were favourable — access to cheap power, supportive local policy, or anchor industrial demand — but these successes remain geographically contained. Blue hydrogen continues to play a significant role, not because it is optimal in the long term, but because decisions on infrastructure, permitting, and long-distance transport were repeatedly deferred.

The infrastructure that does exist reflects this hesitation. Regional hydrogen pipelines have been built, but cross-border connectivity remains limited. Ports and storage facilities are under-utilised, often developed ahead of demand in some regions and lagging badly in others. Hydrogen is still treated primarily as a specialist chemical — something produced, consumed, and managed within defined industrial boundaries — rather than as a flexible energy carrier integrated into the wider system.

Policy and market frameworks mirror this fragmentation. Carbon pricing mechanisms are in place, but they remain inconsistent and politically fragile, vulnerable to economic cycles and changes in government. Subsidy schemes come and go,

frequently redesigned before they have time to deliver scale. Long-term offtake agreements, essential for project finance, remain the exception rather than the norm. As a result, financial markets continue to price hydrogen projects as high-risk infrastructure, demanding returns that further slow deployment.

By 2050, the consequences of this incrementalism are visible in the climate system itself. Global emissions have declined, but not fast enough to meet the ambitions set out in the Paris Agreement, which has been repeatedly reaffirmed but never fully delivered. Warming has not been halted, only moderated. Extreme weather events are more frequent, and the costs of adaptation — flood defences, grid hardening, insurance, and disaster recovery — have become a permanent and growing burden on public finances.

In terms of end use, hydrogen finds a role primarily where there are few credible alternatives. Refining, fertiliser production, and parts of the steel sector decarbonise incrementally using hydrogen, often driven by regulation rather than economics. Beyond these applications, electrification dominates. Transport, heating, and power systems continue along largely incremental pathways, optimised around existing grids and assets. Aviation and shipping make progress, but struggle with fuel availability, cost, and scale.

The system-level outcome is mixed. Emissions fall, but not far or fast enough to consistently meet net-zero targets. Energy security improves in some regions, yet dependence on legacy fuels persists elsewhere, exposing economies to continued price volatility and geopolitical risk. Climate risk becomes a background condition of economic planning rather than an emergency response, shaping investment decisions even as it constrains them.

Hydrogen, in this future, never truly fails — but it never leads. Its potential is acknowledged, its role accepted, and yet its impact remains constrained.

This is not a future of collapse or crisis. It is a future defined by caution, delay, and partial commitment. Above all, it is a missed opportunity.

Scenario 2: Coordinated Transition — Hydrogen as System Infrastructure

In this future, hydrogen becomes boring — and that is precisely its greatest success. It is no longer discussed as a "fuel of the future" or framed as a technological bet. Instead, it fades into the background architecture of the energy system, doing its job quietly and reliably, much like electricity networks or gas grids once did.

At the infrastructure level, hydrogen has scaled as part of an integrated system rather than as a collection of standalone projects. Large-scale electrolysis is routinely co-located with offshore wind, nuclear baseload, and grid-balancing assets, absorbing surplus power and providing flexibility to the wider energy network. Hydrogen transmission and distribution increasingly resemble mature gas and electricity systems, with shared planning, regulated access, and interoperability across regions.

Storage plays a central role in this stability. Salt caverns, liquid hydrogen carriers, and hydrogen-derived fuels such as ammonia provide seasonal resilience, smoothing variability across weeks and months rather than hours alone. Crucially, standards for hydrogen quality, certification, and safety have converged internationally, allowing equipment, fuels, and capital to move across borders with minimal friction.

Policy and market design play a decisive — if often underappreciated — role in this future. The foundations were laid not in the 2040s, but two decades earlier, when governments recognised that the primary barrier to hydrogen deployment was not technology, but revenue certainty. Rather than attempting to pick winners or dictate specific technical solutions, policy in the late 2020s focused on creating conditions that allowed markets to function.

Stable carbon frameworks were established early and maintained consistently, providing a credible long-term signal rather than a shifting political compromise. Crucially, mechanisms to guarantee and stabilise offtake were introduced at scale. Contracts for Difference, long-term offtake guarantees, and clear certification regimes reduced demand risk and gave investors confidence that hydrogen projects would have predictable revenue over decades, not just election cycles.

As these frameworks bedded in, the role of government gradually evolved. Policymakers acted as market shapers rather than micromanagers, setting boundaries, standards, and guardrails while allowing industry and finance to determine optimal configurations. Guarantees of origin, regulated returns for strategic infrastructure, and interoperable standards further reduced risk, allowing hydrogen assets to be financed on terms increasingly comparable to regulated utilities.

The impact on capital markets was profound. As revenue uncertainty fell, so too did the cost of capital. Projects that would once have required heavy subsidy became financeable at scale, attracting institutional investors seeking long-term, stable returns. By 2050, hydrogen infrastructure was no longer perceived as

experimental or speculative, but as a mature component of the energy system — a shift made possible by policy decisions taken a generation earlier.

By 2050, the impact of this coordination is visible across end-use sectors. Steel, chemicals, and refining are largely decarbonised, with hydrogen embedded into core industrial processes rather than bolted on as an afterthought. Hydrogen also supports dispatchable power generation, long-haul transport, and the production of synthetic fuels for aviation and shipping. For most consumers, however, hydrogen is rarely noticed. It works behind the scenes, enabling reliability and flexibility without demanding attention.

At the system level, net-zero targets are broadly met. Emissions reductions are sustained rather than episodic, and climate risks, while not eliminated, are stabilised and increasingly manageable. Energy systems are more resilient, better balanced, and less geopolitically brittle, with fewer choke points and more diversified supply chains. Hydrogen fulfils its intended role not as a silver bullet, but as the connector that allows multiple decarbonisation pathways to function together at scale.

My Vision

'This is not a prediction. It is one plausible outcome among many — contingent on choices made decades earlier.

The year is 2048. I am 69 years old, and I find myself sitting on a tram from the ever-expanding KVV (Karlsruhe Verkehrsverbund), travelling from my house deep in the Black Forest to Baden Airpark. Today I am flying to London to speak at a conference — not about the energy transition and what we still need to do, but a retrospective, looking back at how we got it right.

As I travel on the tram, I find myself reflecting on everything that has happened over the past decades. Where I live, near Rothaus, the energy transition isn't happening — it's done. I don't notice it. The nearby brewery has long since installed a hybrid energy system across its facilities, using solar power in the summer and an agreement with the grid to purchase green wind power from northern Germany in the winter. This energy is used to produce fuel for all the forklifts operating across the site.

My own house has a very modern heat pump, designed to keep my wife and I warm and to cope with the increasingly variable weather conditions of the region. I must admit, I still have a wood burner. Not because I need it, but for effect —

especially on those cold days when it snows, when I'm reading a book and looking out of the window.

The state of Baden-Württemberg as a whole has taken great steps forward. It is still run by the Green Party, but without much of the baggage of yesteryear. In fact, following several breakthroughs in nuclear fusion in recent years — one of them led by a former PhD student of Kate Lancaster in York — Baden-Württemberg is now set to be one of the first states to install a pilot small-scale fusion reactor. It is planned for KIT Nord, returning the site in some ways to its roots as a nuclear research facility. I find this particularly exciting, as it means that, hopefully before I am gone, one of the next times I ride this tram it may be powered, at least in part, by that reactor.

I have other reasons to be excited today. I will be flying into London City Airport, which is ideal as it sits right next to the ExCeL Centre where the conference is being held. More importantly, I will be travelling on one of the new Embraer hydrogen aircraft, developed in partnership with ZeroAvia after an agreement signed many years ago. Even more impressive is that the fuel cells powering the aircraft were developed by a local team in Brazil, a reminder that local manufacturing can — and does — work.

After the conference, the trip becomes more personal. I travel across London to visit my brother in Ealing, where we meet for a pint in the pub we drank in as youngsters. Being younger than me, he has only recently retired. As an electrician in London, he is in a far stronger financial position than his older brother — not just because of his profession, but because he spotted an opportunity early on. I like to think it was after reading his big brother's book. Already specialised in providing power solutions for events and festivals, he was among the early adopters of hydrogen fuel cell technology to power temporary sites. Being first gave him a clear advantage, helped build his reputation, and allowed him to grow his business.

Later, my wife joins us for dinner after spending the day at an art gallery, rather than listening to her husband drone on at a conference. Late that night, we head back towards the station near our hotel. Tomorrow the journey continues — we will be travelling north to Stirling to visit my other sisters. To get there, we'll take the new HS 5 line (and yes, if the UK really does make it to HS 5, I will be impressed), and after changing in Edinburgh, we'll continue on one of the latest-generation Alstom hydrogen trains, delivered just three years earlier.

All in all, I am really looking forward to the journey — simply seeing, first-hand, how much has changed over the past twenty-plus years.

And as for my son — well, he didn't listen to me and get directly involved in the energy transition. Instead, he found his own path, developing a love for the natural sciences and becoming a chemical engineer. He may not work on the transition explicitly, but he is very much part of it, working for a large chemical company and holding a patent for a next-generation liquid hydrogen organic carrier. We may not see him as often as we would like, but whenever there is snow on the Feldberg, he is always there with the family, skis on, as ever.'

This is not a future defined by spectacle or radical disruption. It is defined by competence, coordination, and follow-through. The transition works — and because it works, it largely disappears from view.

Scenario 3: Accelerated Breakthrough — A Rewritten Energy Map

In this future, progress does not follow a smooth or linear path. A combination of urgency, steep learning curves, and systemic shock compresses decades of change into a matter of years. Hydrogen does not simply integrate into existing energy systems; it reshapes them, altering where energy is produced, how it is traded, and who benefits from it.

What Drives the Acceleration

The speed of this transition is not the result of optimism or collective enlightenment. It is driven by pressure. By the early 2030s, the cost of delay begins to exceed the cost of action across multiple dimensions at once.

Climate impacts, long treated as future risk, become present and cumulative. Extreme weather events, supply-chain disruptions, and infrastructure failures no longer register as anomalies, but as systemic features. At the same time, energy price volatility exposes the fragility of fuel-dependent economies, while insurers and financial institutions begin to withdraw from high-risk assets and regions. Adaptation costs rise faster than mitigation costs, forcing governments and markets to reassess their priorities.

In parallel, technology learning curves steepen. Electrolysis, power electronics, storage, and digital optimisation benefit from scale effects and cross-sector spillovers. Costs fall faster than expected, performance improves, and integration becomes easier. What initially required support becomes increasingly competitive, and hesitation begins to look irrational rather than prudent.

Geopolitics further accelerates this shift. Energy security is reframed as industrial strategy. Nations recognise that clean power, hydrogen production, and conversion capacity are not just tools of decarbonisation, but sources of economic resilience and influence. Competition for leadership in clean energy supply chains intensifies, compressing timelines that would once have stretched across decades.

Under this combined pressure, institutional barriers give way. Planning, permitting, and grid expansion processes are reformed not out of idealism, but necessity. Emergency measures introduced to address short-term crises persist and evolve into new norms. Once these constraints are removed, deployment accelerates rapidly, driven as much by momentum as by policy.

The result is not a single breakthrough moment, but a cascade. Each improvement lowers risk, attracts capital, and reinforces the next. Acceleration becomes self-sustaining.

System Transformation at Speed

On the technology side, cost reductions arrive faster than expected. Electrolysis benefits from rapid manufacturing scale-up, materials innovation, and tighter integration with surplus renewable and nuclear power. As costs fall and deployment accelerates, hydrogen production becomes increasingly locationally flexible. Rather than being tied to legacy industrial centres, production follows clean energy availability, water access, and infrastructure connectivity.

This flexibility enables a global hydrogen trade that begins to rival liquefied natural gas at its peak. New shipping routes emerge around hydrogen and hydrogen-derived carriers, such as ammonia and synthetic fuels, unlocking markets that were previously constrained by geography or fuel availability. Energy trade becomes less about moving fossil resources and more about moving value added through conversion, storage, and transport.

Policy and markets respond to this acceleration by reframing energy as strategic infrastructure. Governments streamline planning and permitting processes, accelerate grid expansion, and treat energy networks with the same urgency once reserved for transport or defence. As regulatory frameworks stabilise and scale becomes visible, capital flows rapidly. Hydrogen assets increasingly resemble regulated utilities, drawing in institutional investors and reducing financing costs. The energy system shifts away from managing scarcity towards managing flexibility, resilience, and optimisation.

The effects are most visible at the point of use. Heavy industry begins to relocate or expand in regions with abundant clean energy and hydrogen infrastructure, reshaping industrial geographies. Aviation and shipping decarbonise faster than anticipated through the widespread availability of synthetic fuels. Hydrogen also enables negative-emissions systems at scale, combining clean energy, carbon capture, and long-duration storage. As flexibility increases, energy poverty declines, particularly in regions that were previously exposed to volatile fuel imports.

At the system level, emissions fall faster than planned, not because targets were lowered, but because deployment outpaced expectations. Geopolitical energy dependencies weaken as energy systems diversify and decentralise. The traditional boundaries between electricity, molecules, and digital systems blur, replaced by integrated energy platforms optimised in real time.

This future is not guaranteed. It depends on alignment, speed, and the willingness to act decisively under pressure. But it is not speculative. Every component required to reach it already exists in some form today. What distinguishes this future is not technological miracle, but the compound effect of learning, coordination, and urgency.

From Futures to Choices

The three futures described in this chapter are not separated by technology or by physics. In each case, the same core components are available: clean power, electrolysis, hydrogen infrastructure, storage, and digital optimisation. The underlying science does not change. What changes is how these elements are combined — and how quickly.

In the Business-as-Usual pathway, progress is real but fragmented. Decisions are deferred, risks are left unaddressed, and deployment remains cautious. Hydrogen finds a role but never escapes its industrial silos. The system moves forward, but not together.

In the Coordinated Transition, those same technologies are aligned. Policy reduces risk rather than adding complexity, markets provide long-term signals, and infrastructure is built with system integration in mind. The transition succeeds not through speed alone, but through consistency.

In the Accelerated Breakthrough, alignment meets urgency. Pressure removes optionality, learning curves steepen, and institutional barriers fall away. The system does not merely integrate — it reconfigures.

The difference between these futures is not belief, ambition, or innovation. It is speed of decision-making, willingness to build ahead of certainty, and the ability to think in systems rather than silos.

None of these futures begins in 2050. All of them are already taking shape today — in permitting offices, balance sheets, grid codes, and boardroom decisions. The trajectory is not set by what is imagined, but by what is authorised, financed, and built.

Hydrogen, as this book has argued throughout, is not the destination. It is the connector that determines whether the energy transition functions as a system or fragments into parts. The question is no longer whether the tools exist. It is whether they are being used deliberately.

The chapters that preceded this one stepped away from scenarios and towards agency — from understanding what is possible to examining what is required. This chapter has shown how those choices compound over time, shaping not just energy systems, but the world they support.

18. Conclusions

"If everything seems under control, you're not going fast enough." — Mario Andretti

By most objective measures, the energy transition has already crossed thresholds that were once considered decisive. The economics have shifted. Cost — long treated as the ultimate barrier — is no longer the sole, and often not the primary, constraint holding progress back.

Yet, decision-making remains hesitant. Projects stall. Policies lag. Capital waits. The dissonance between what is technically and economically possible, and what is actually being built has become one of the defining features of the transition.

That gap is not explained by a lack of technology, scale, or viable pathways. It is explained by something far less tangible — how expertise is weighed, how interests are defended, and how delay quietly becomes the default response to uncertainty.

It is here, rather than in chemistry or engineering, that the real friction begins.

ACT I — We Are Not Where People Think We Are

The Economics Have Already Turned

One of the most striking indicators that we are no longer waiting on a technology or cost breakthrough comes from the latest report by the International Renewable Energy Agency (IRENA). According to Renewable Power Generation Costs in 2024, an astonishing 91 % of all new utility-scale renewable power projects commissioned in 2024 were more cost-effective than any new fossil fuel alternative, underscoring a dramatic shift in the economics of energy. Solar photovoltaic projects were, on average, 41 % cheaper than the lowest-cost fossil alternatives, while onshore wind projects were 53 % cheaper, making renewables not just environmentally preferable but financially compelling.

Over 582 GW of new renewable capacity was added in 2024 — nearly a 20 % increase over the previous year — helping to avoid billions of dollars in fossil fuel expenditures and reinforcing that clean energy is often the least expensive source of new generation.

This evidence dismantles a persistent argument used to delay action: the notion that clean technologies, including hydrogen, are prohibitively expensive. It reveals that the broader energy system is already experiencing a cost parity — and in many cases, a cost advantage — for clean technologies. The fact that renewables now dominate new capacity additions globally signals that the economic conditions for decarbonisation are not only emerging, but they are also already here; the remaining questions are about policy coherence, infrastructure deployment and social will, rather than about whether the technologies themselves can compete.

Adoption Is Happening Faster Than We Admit — Even If the Headlines Suggest Otherwise

There is a persistent narrative surrounding hydrogen that it is slow, lagging, or perpetually five years away. That progress is tentative, fragile, and constantly on the verge of disappointment. It is a narrative repeated so often that it has begun to feel like common sense.

History suggests otherwise.

When measured honestly, hydrogen is being adopted faster than many previous energy technologies at comparable stages of maturity — faster than wind, faster than solar, and certainly faster than nuclear. The difference is not pace; it is visibility. Wind turbines and solar panels announce themselves on skylines and fields. Hydrogen infrastructure often operates quietly in the background — inside industrial sites, behind factory walls, within ports, depots, and pipelines. Its progress is structural rather than spectacular.

This distinction matters. Perception shapes policy, capital allocation, and institutional confidence. If hydrogen is assumed to be slow, fragile, or perpetually delayed, decisions are deferred and ambition diluted. In reality, the pace of adoption is not the problem. The lack of recognition of that pace is.

This disconnect was articulated clearly at the hydrogen expo in Hamburg in October 2025. In one discussion, Ivana Jemelkova, CEO of the Hydrogen Council, made an observation that resonated strongly across the room: there is, in fact, more positive news in hydrogen than negative — but negative news sells papers.

This is not a conspiracy, nor is it unique to hydrogen. It is simply how attention works. Delays, setbacks, and controversy generate headlines. Quiet progress does not. A project that slips a year becomes a story; dozens that move steadily from concept to construction rarely do. The result is a narrative skewed toward exception rather than trend.

When viewed systemically, however, a different picture emerges. Hydrogen is advancing simultaneously across multiple fronts: industrial decarbonisation, mobility, power balancing, synthetic fuels, export markets, and microgrids. Manufacturing capacity has expanded. Policy frameworks have hardened. Pilot projects have become first-of-a-kind deployments, and those deployments are now being replicated.

Data from organisations such as IRENA reinforces this broader trajectory, as does the growing depth of supply chains and institutional commitment across regions.

None of this is to suggest that the transition is easy, linear, or guaranteed. Challenges remain — technical, economic, and political. But the dominant story is not one of stagnation. It is one of uneven but accelerating progress.

There is, therefore, real cause for optimism here — not naïve optimism, and not hype, but data-driven confidence. Hydrogen is not waiting to begin. It is already moving. The task now is not to invent momentum, but to recognise it, reinforce it, and ensure that perception finally catches up with reality.

What This Book Has Been Arguing All Along

At its core, this book has not argued that hydrogen is the answer to the energy transition. It has argued something more specific — and more important. Hydrogen is the connector that allows the transition to function at scale. It links sectors that otherwise remain siloed. It carries energy across time and distance where electrons alone cannot. It enables industry, mobility, storage, resilience, and system balancing to move together rather than in isolation.

Without hydrogen, decarbonisation remains fragmented: successful in parts, constrained in others, and fragile at the edges. With it, the transition becomes coherent — not simpler, but workable. The challenge, therefore, has never been whether hydrogen matters, but whether we are prepared to organise our systems, markets, and decisions around the role it is uniquely suited to play.

ACT II — What the Market Actually Needs

Selling the Journey, Not the Destination (and the Fallacy of the Corporate Product)

One of the quiet mistakes repeatedly made in emerging energy markets is the belief that success comes from selling the end state. Too often, companies present hydrogen as a finished vision: fully decarbonised systems, optimised infrastructure, perfect economics, and seamless integration. In doing so, they aim to maximise future market share and position themselves as the eventual winner. The problem is that this is not what the market needs today.

The market does not need a Rolls-Royce.

It needs a Volkswagen Golf.

Transitions succeed when they meet people where they are, not where the roadmap says they should be in twenty years. Customers, utilities, municipalities, and operators are not looking for perfection; they are looking for reliability, familiarity, and incremental progress. A solution that works today, even if it is not optimal, will always beat a perfect solution that remains theoretical.

Hydrogen's opportunity lies not in promising a flawless destination, but in offering a credible journey — one that allows organisations to move step by step, learning, adapting, and scaling without having to bet everything upfront. The companies that understand this will not necessarily dominate the future market, but they will shape it.

This tendency to oversell the destination is closely linked to another, less discussed problem: what might be called the fallacy of the corporate product.

Recently, at a trade fair, I found myself standing beside a peer who works for one of the large global consulting firms. We were both looking at the same product on a stand — well presented, confidently marketed, and supported by an impressive narrative about how it would "unlock value" and "accelerate transition." Curious, I asked them what they thought.

They paused and said, almost dismissively, "It's just another corporate product."

Seeing my confusion, they explained. It wasn't that the product was poorly engineered or badly thought through. Quite the opposite. It had clearly absorbed a great deal of internal effort. The issue was that it had been designed to succeed inside a large organisation rather than in the market it claimed to serve.

"It's something that's been put together to help someone get promoted, or to justify a position internally," they said. "But it offers nothing to my customers. There's no real value add. There's no way I could take that into a consultancy and sell it."

I couldn't help but chuckle — because the pattern was instantly recognisable.

This is the corporate product fallacy: solutions optimised for internal alignment, executive approval, and PowerPoint coherence, rather than for customer outcomes. They look compelling at first glance — neat frameworks, elegant diagrams, ambitious language — but when you scratch the surface, there is little underneath that truly reduces risk, cost, complexity, or uncertainty for the end user.

These products are rarely created in bad faith. They are a natural consequence of large organisations operating under pressure, where visibility matters as much as impact and where success is often measured by how well an idea travels upwards, rather than how well it performs outwards. In such environments, it becomes easier to build something that can be defended internally than something that must prove itself in the field.

The hydrogen economy is particularly vulnerable to this dynamic. The challenges it faces are inherently cross-sectoral and system-level, yet the market is increasingly crowded with generic solutions that promise universality while avoiding specificity. They describe the destination in exquisite detail but offer little guidance for the first practical steps of the journey.

For customers, the test is disarmingly simple: does this help me make a decision today? Does it fit how budgets are approved, assets operated, risks managed, and systems integrated in the real world? If the answer is no, then the sophistication of the narrative is irrelevant.

This matters because every product that fails to deliver tangible value does more than waste time. It erodes trust. It reinforces scepticism. Additionally, it strengthens the perception that the energy transition is being driven by rhetoric rather than utility.

If hydrogen is to scale, it will not do so through solutions designed primarily to impress committees or justify internal strategies. It will scale through offerings that are almost unremarkable — practical, modular, and immediately useful. Not a Rolls-Royce vision of the future, but a dependable vehicle that gets people moving now, and can be upgraded as the road unfolds.

It is also worth being clear about what this argument is not saying. This is not a rejection of sophistication, scale, or integrated solutions. Nor is it an assertion that so-called "corporate products" will never have a role. Their day may well come.

But it will come later, not first.

Large, highly integrated solutions make sense only once customers have already travelled part of the journey — once they have experience, confidence, organisational capability, and a clear understanding of how hydrogen fits into their operations. Only then does the case for a Rolls-Royce solution become credible, justifiable, and economically defensible.

Trying to sell that level of complexity too early does not accelerate the transition; it slows it. It asks customers to absorb risk they are not yet equipped to carry, and to commit to architectures they have not yet learned to operate. In contrast, simpler, modular solutions allow organisations to move, learn, and adapt — building the very capability that will eventually make more sophisticated offerings viable.

In that sense, the journey is not a compromise on ambition; it is a prerequisite for it. The destination remains important, but it can only be reached if customers are allowed — and enabled — to progress there step by step. The Rolls-Royce will make sense in time. First, the road needs to be travelled.

What This Means in Practice

If hydrogen's role is to connect systems rather than replace them, then the implications differ depending on where one sits in the transition. The challenge is shared, but the levers are not.

For policymakers, the priority is not perfection, but consistency. Long-term frameworks matter more than short-term optimisation. Clear standards, predictable permitting, and demand-side signals reduce uncertainty far more effectively than incremental subsidies or ever-more complex schemes. Markets move when direction is stable.

For investors and lenders, the implication is temporal rather than ideological. Hydrogen does not fail because it lacks returns; it struggles because its returns arrive later. Risk-sharing mechanisms, blended finance, and structures that reward learning-by-doing matter more than precision forecasting. Capital does not need certainty — it needs confidence that rules will not change midstream.

For industry and OEMs, scale will not come from proprietary completeness, but from deployability. Modular solutions, interoperability, and systems that integrate cleanly into existing operations lower barriers to adoption far more effectively than vertically perfect architectures. The winners will not be those with the most elegant end-state designs, but those whose products get used, iterated, and trusted early.

For customers and the wider public, the test is simpler and more human. Energy transitions succeed when they deliver reliability, safety, and tangible everyday benefits — quieter streets, cleaner air, resilient power, and secure jobs. Acceptance grows not through persuasion, but through familiarity. What works, stays.

None of this requires unanimity, and none of it requires belief in a single pathway. It requires alignment — not around ideology, but around function. When each part of the system plays to its role, momentum stops being something we wait for and becomes something we create.

ACT III — Why Arguments About “Winners” Miss the Point

Which Technology Wins — and Why It Doesn’t Matter

One of the most unproductive debates in the hydrogen economy centres on technology: PEM versus alkaline, AEM versus emerging alternatives; electrolysis versus thermochemical routes; this catalyst versus that membrane. These debates often carry an intensity that suggests a single, decisive winner must eventually emerge.

The truth is simpler — and far less dramatic.

It does not matter which single technology wins because none of them will win alone.

The scale of the global energy system is simply too large, too diverse, and too constrained by local realities for a single technological pathway to dominate. Different applications demand different characteristics. Heavy industry prioritises availability and throughput. Mobility values responsiveness and footprint. Grid balancing rewards flexibility and integration. Remote systems optimise for robustness and simplicity. No single hydrogen technology excels across all of these dimensions simultaneously.

Geography adds another layer. Regions with abundant renewable power optimise differently from those with constrained grids. Water availability, land use, temperature, regulatory frameworks, and existing infrastructure all shape which technologies make sense in practice. Cost curves will diverge. Supply chains will evolve unevenly. Material availability will impose real, physical limits on how far any one solution can scale on its own. As rare earths, catalysts, membranes, and power electronics remind us, supply-chain dependency does not disappear in a transition — it simply changes form, making diversity across technologies a form of resilience rather than indecision.

What matters, therefore, is not technological purity — it is deployment volume.

The energy transition is neither a beauty contest, nor an academic exercise in optimisation. It is a logistics challenge of unprecedented scale. Success will be determined not by which technology looks best on a slide, but by which combinations of technologies can be manufactured, deployed, operated, and maintained fast enough to matter.

This reality becomes even clearer when viewed through a geopolitical lens. China is, and will remain, a major player in hydrogen technologies — particularly in manufacturing scale, supply chains, and cost reduction. But it will not be the only player, nor can it be. The global hydrogen market is far too large, and far too strategically important, to be supplied by any single country or region.

Global energy demand is vast — and still growing. Even under the most aggressive decarbonisation scenarios, the absolute scale of hydrogen required leaves room for multiple technologies, multiple manufacturers, and multiple national ecosystems to coexist. There is no credible pathway in which one electrolyser type, one fuel cell architecture, or one country supplies a meaningful share of global demand on its own.

The elephant is simply too big.

This is particularly true for electrolyser and fuel cell OEMs. The volumes implied by national hydrogen strategies, industrial decarbonisation, mobility, power balancing, and export markets are such that there is room — and need — for many manufacturers to operate successfully in parallel. From a system perspective, diversity is not inefficiency; it is resilience.

For this reason, every serious hydrogen nation should be thinking not only about deployment, but about home-grown OEM capability. Not to exclude international suppliers, but to ensure security of supply, industrial learning, skilled employment, and the ability to adapt technology to local needs. Hydrogen infrastructure

will underpin critical sectors of future economies. Relying entirely on external manufacturing for that backbone introduces vulnerabilities that few countries are comfortable accepting.

This does not imply fragmentation or protectionism. On the contrary, global collaboration, shared standards, and open markets will be essential to scale. But collaboration works best when it is built on multiple strong participants, not on dependency. A healthy hydrogen ecosystem is one where companies compete, cooperate, and co-evolve across borders — not where success for one implies failure for others.

Scarcity thinking — the fear that another technology's success comes at the expense of one's own — drives defensive behaviour: lobbying against alternatives, fragmenting standards, slowing deployment to protect position. In hydrogen, this instinct is not just misplaced; it is counterproductive.

Scale will not be achieved by waiting for rivals to fail. It will be achieved by building capacity together — even if that capacity initially exceeds demand. Overcapacity in the early phases is not a flaw; it is how learning curves are accelerated, costs are driven down, and markets are formed. Every major energy technology that now appears inevitable passed through a phase of apparent excess.

The hydrogen economy will be no different.

The question, then, is not which technology wins, but how quickly enough of them scale in parallel to reshape the system. Collaboration does not require uniformity, and competition does not require exclusion. There is space — and need — for multiple pathways, multiple OEMs, and multiple countries to advance at once.

In a transition measured in gigawatts, millions of tonnes, and decades, the real risk is not backing the "wrong" technology. The real risk is failing to build fast enough, broadly enough, and together.

ACT IV — The Real Friction Is Human and Structural

Expertise, Interests, and the Cost of Delay

Unfortunately, the uncomfortable truth is that too often the future is being sacrificed for today's profit. We live in a world where personal belief, immediate interest, or political convenience frequently carries more weight than informed expert opinion. The consequences of this are not abstract; they play out in

delayed decisions, compromised policies, and transitions that move far more slowly than the underlying realities demand.

Hailing from the UK, I cannot ignore one moment that crystallised this shift in public discourse. During the Brexit referendum campaign, Michael Gove remarked, "I think the people of this country have had enough of experts." Whether that was his intention or not, the statement became a touchstone — a signal that expert analysis could be dismissed outright if it conflicted with personal conviction or popular sentiment. Similar dynamics have since played out elsewhere. Donald Trump's repeated use of the phrase "Fake News" demonstrated how easily public trust in expertise can be undermined when facts become inconvenient.

This erosion of trust in expertise is not confined to politics. Many of us will recognise it from professional life. Detailed analysis is presented, assumptions are made explicit, data is carefully interpreted — only for the management response to be, "I have other facts than that." Not because the work is flawed, but because the conclusions do not align with a preferred outcome. Equally familiar are meetings where a complex issue is presented with nuance and balance, yet a single throwaway line — convenient to an individual agenda — becomes the sole takeaway.

The hydrogen economy is not immune to these dynamics. Companies with deep, long-standing investments in oil and gas will naturally seek to protect their positions. Politicians with personal or financial ties to the status quo may hesitate to advance legislation that risks disrupting it. These behaviours are rational at the individual level, but collectively they slow progress, fragment decision-making, and dilute ambition.

If the transition is to accelerate, the industry must not only make the case for hydrogen — it must also get its own structures right. Too often, organisations focus narrowly on specific verticals without accounting for wider system requirements. Hydrogen, by definition, does not sit neatly within a single silo. Production, storage, transport, and use cases span multiple sectors, each requiring different skill sets, time horizons, and integration challenges.

Bundling activities simply because they belong to the same industry is rarely efficient. In practice, there is often more in common between teams developing electrolyser stack manufacturing and those building battery factories than between those teams and others working on full hydrogen system integration. The same applies across use cases: electric vehicle charging infrastructure, hydrogen refuelling solutions, and energy storage share strikingly similar design, safety, and

operational challenges. Treating them as entirely separate domains duplicates effort and slows learning.

A similar pattern emerges on the supply side. Modular electrolysis, wind, battery energy storage systems, solar, and — in the future — small scale nuclear all share comparable modular design philosophies and integration requirements. Yet too often they are still organised, evaluated, and regulated as if they were extensions of traditional, centralised chemical or thermal plants. Aligning these technologies around the operational realities of a modular, distributed energy system — rather than the structures of the past — would significantly streamline deployment, integration, and scale.

There is one further dimension to this journey that deserves explicit acknowledgement: the role of microgrids and islanded systems. As discussed earlier in this book, the energy transition will not unfold as a single, perfectly integrated system from the outset. It will emerge through a patchwork of semi-autonomous "islands" — ports, industrial clusters, remote communities, campuses, data centres, and critical infrastructure — each developing at its own pace, under its own constraints.

Building these islands requires a specific kind of expertise. Microgrid design, integration, and operation sit at the intersection of power systems, storage, controls, and increasingly hydrogen. This is not simply a scaled-down version of the centralised grid; it is a discipline in its own right, demanding a deep understanding of resilience, intermittency, system boundaries, and failure modes. These skills are essential to making early hydrogen systems work in practice.

In many cases, it will be these islanded solutions that allow customers to take their first meaningful steps. A hydrogen-enabled microgrid can reduce emissions, improve resilience, and create local value without waiting for national infrastructure, market reform, or full system integration. In doing so, it builds operational experience, institutional confidence, and technical capability — the foundations required for later scale.

Seen this way, microgrids are not a detour from the destination; they are part of the journey itself. They create the proving grounds in which technologies mature, teams learn, and customers become ready for more complex, more integrated solutions. Without this phase, the leap to fully interconnected, system-wide architectures is not bold — it is brittle.

The irony is that many of the competencies developed in these early islands — system integration, digital coordination, modular design — are precisely those

needed at scale. The path from islands to continents is not a reset; it is an accumulation. Without those islands being built well, the wider system will have nothing solid to connect to.

The digital layer reinforces this point further. As discussed earlier in the book, many of the digital solutions required to coordinate hydrogen systems — forecasting, optimisation, asset management, and market integration — are largely technology-agnostic. One well-designed digital architecture can serve multiple parts of the system simultaneously, yet too often these solutions are developed in isolation.

None of this requires a technological breakthrough. What it requires is a willingness to prioritise long-term system value over short-term positional advantage, and to organise expertise around how the future energy system must function, rather than how today's industries are structured.

The people capable of doing this already exist. What remains uncertain is whether they will be empowered — politically, organisationally, and culturally — to act before delay becomes the most expensive choice of all.

Shareholder Value, Time Horizons, and the Cost of Waiting

It would be incomplete to discuss the pace of the hydrogen transition without acknowledging the role of shareholder value and the way modern capitalism evaluates success.

Public companies are structurally incentivised to optimise for short-term performance. Quarterly earnings, near-term margins, and predictable returns dominate decision-making, not because executives lack vision, but because the system rewards certainty over transformation. Investments that carry higher upfront risk, longer payback periods, or uncertain market timing are often deprioritised — even when their long-term strategic value is clear.

Hydrogen exposes this tension more starkly than most technologies. Its benefits accrue over decades, not quarters. Infrastructure must often be built ahead of demand. Early utilisation is, by definition, inefficient. These characteristics sit uncomfortably within financial models designed for incremental improvement rather than systemic change.

There is a harder truth beneath the time-horizon problem: markets do not innovate towards what is needed; they innovate towards what is measurable and profitable. That is not a moral claim — it is simply how the system is designed. If the benefits of an innovation are social, ecological, or preventative, they are often

weakly priced or not captured at all. Avoided disasters do not show up as revenue. Clean air does not generate a cash flow to the people who fund it. A stable climate is the ultimate public good — and public goods are systematically underprovided when left to private optimisation.

This creates a structural drag on the kinds of innovation the transition most requires; long-horizon, infrastructure-heavy, coordination-dependent investments whose returns are real but diffuse. Capital follows the scoreboard it is given. If the scoreboard rewards near-term monetizable value more reliably than long-term system integrity, innovation will cluster accordingly— even when the latter is what society cannot afford to delay.

This does not mean capitalism is incompatible with the energy transition. On the contrary, markets have already demonstrated an extraordinary ability to scale renewables, reduce costs, and mobilise capital once trajectories become clear. But capitalism is at its most effective after direction has been set, not while it is still being debated.

The risk is not that companies are acting irrationally, but that collectively rational decisions produce a sub-optimal outcome. Each firm waits for clearer signals; each board hesitates to move first; each investor prefers optionality over commitment. The result is delay — not because the future is unattractive, but because the present still pays.

Breaking this dynamic requires actors willing to stretch time horizons: investors who accept longer payback periods, executives who frame early hydrogen investments as strategic positioning rather than immediate profit centres, and policymakers who reduce uncertainty so that capital can move with confidence. Without this shift, the system remains locked into defending existing value rather than creating future value.

In that sense, the challenge is not ideological. It is temporal.

The energy transition asks a short-term optimisation machine to solve a long-term problem. Progress depends on how effectively those time horizons can be aligned — and on how many decision-makers are prepared to accept temporary discomfort in exchange for durable advantage.

The Heroes Are Already On the Field

There is a moment in the rugby documentary *Living with Lions* that has stayed with me for years. As an avid rugby enthusiast, I have watched it more times than I can count, and yet it never loses its force. The British & Irish Lions are losing.

Half-time has arrived. The dressing room is heavy with frustration and expectation. And then Jim Telfer — in that unmistakable Scottish brogue — addresses the team.

“We need some heroes out there,” he says. “You can’t expect it to happen. You need to make it happen.”

It is not a technical speech. There is no discussion of tactics, formations, or marginal gains. It is a challenge about responsibility. About agency. About the difference between waiting for momentum and creating it.

I find myself returning to that moment often when thinking about the hydrogen economy — and the energy transition more broadly.

Much of the current discourse assumes that progress will emerge organically once the conditions are right: once costs fall a little further, once regulations align more neatly, once markets mature, once someone else moves first. But change of this scale does not work that way. Momentum is not something that arrives fully formed. It is built, deliberately, by people willing to step into uncertainty and act before success is guaranteed.

The truth is that the heroes we are waiting for already exist.

They are visible at every level of the system. In politics, there are policymakers pushing for long-term frameworks in environments dominated by short-term incentives. In industry associations, there are individuals trying to align competitors around shared standards and infrastructure. In business, there are executives committing capital to projects that will not deliver easy returns within a single reporting cycle. In engineering teams, regulators, financiers, and operators, there are countless professionals doing the hard, often uncelebrated work of turning ambition into reality.

What is missing is not capability, intelligence, or intent. It is the collective permission — and expectation — to lead rather than wait.

Hydrogen, by its nature, exposes this dynamic more starkly than most technologies. It does not fit neatly into existing systems. It requires coordination across sectors, time horizons, and political boundaries. It demands early movers to absorb risk so that others can follow. That is not a failure of the system; it is the defining feature of every major infrastructure transition in history.

In rugby, heroes are not always the ones who score the tries. Often, they are the players who make the first dominant carry, who take the hit, who change the emotional direction of the game. Their impact is felt before it is measured.

The same is true here. Leadership in the hydrogen economy is not about grand declarations or perfect strategies. It is about visible commitment. About acting as if the transition is inevitable — and behaving accordingly — even when the surrounding system hesitates.

Waiting for urgency to materialise on its own is in itself a decision, qnd history suggests it is rarely the one that changes outcomes.

If the energy transition is to accelerate meaningfully, it will not be because we finally found the missing technology or the flawless policy design. It will be because enough people, across enough parts of the system, decided to stop expecting it to happen — and chose instead to make it happen.

ACT V — Reclaiming Agency

The Future Is Still Open — But Not For Long

One of the most persistent misunderstandings about the energy transition is the belief that the end state is fixed — that there is a single, inevitable version of 2050 waiting to arrive once the technology matures. In reality, the future is path-dependent. What matters is not only where we are trying to go, but how — and when — we move.

Hydrogen sits at the centre of this uncertainty. It can become a narrow, marginal solution, confined to niche applications and permanent pilot status. It can become a contested battleground, slowed by fragmentation, politics, and defensive positioning. Or it can evolve into a connective layer — linking power, industry, transport, and resilience in ways that make the wider transition workable at scale.

Which of these futures emerges is not determined by chemistry or engineering limits. It is determined by decisions made in the present decade: how infrastructure is sequenced, how markets are shaped, how risk is shared, and how much learning is allowed to happen in the real world rather than on paper.

This is why the timing matters. Choices made in the 2020s do not merely accelerate or delay progress; they narrow the range of futures that remain possible. Delay does not preserve optionality. It quietly removes it.

This is already visible in climate reality. By most credible assessments, the targets set out in the Paris Agreement — particularly the ambition of limiting warming to 1.5 °C — are no longer achievable through gradual adjustment or incremental policy correction. Emissions trajectories, infrastructure lock-in, and political delays have narrowed the feasible pathways far more than was assumed when those goals were agreed.

That does not mean the Paris Agreement was meaningless, nor that action is futile. It means that the window for comfortable transitions has closed. What remains is a narrower, steeper set of options — where outcomes depend less on distant targets and more on near-term decisions about infrastructure, coordination, and learning speed. The question is no longer whether we can follow the original roadmap, but whether we are willing to act decisively within the constraints created by delay.

Seen this way, the energy transition is not a problem of foresight, but of commitment. The future will arrive regardless. The question is whether it arrives by design — or by default.

What We Don't Yet Know — and Why That's Not a Weakness

It would be misleading to suggest that every pathway described in this book will succeed, or that the transition will unfold cleanly. Some technologies will underperform. Some projects will fail. Some assumptions — technical, economic, or political — will prove wrong.

That is not a flaw in the transition. It is how large systems learn.

Every energy system now taken for granted was shaped through trial, error, and correction. Gas grids, power markets, aviation, digital networks — none emerged fully formed or optimised. They evolved because societies were willing to deploy imperfect solutions, observe how they behaved in the real world, and improve them over time.

The real risk, therefore, is not backing the wrong technology or making the wrong bet. It is insisting on certainty before experience. Learning only happens at scale, and scale only arrives once action begins.

Hydrogen will not succeed because it is flawless. It will succeed if we allow it to be used, stressed, adapted, and improved — and if we judge progress by learning speed rather than by early perfection.

The Missing Ingredient: Urgency

By this point, one conclusion should be unavoidable: the energy transition — and hydrogen's role within it — is not being held back by a lack of technology. Electrolysers work. Pipelines can be built. Storage is understood. Use cases exist across industry, mobility, and power. The engineering challenges are real, but they are solvable. The economics are difficult but improving. Policy frameworks are imperfect, but not absent.

What is missing is something less tangible, and far more powerful: a shared sense of urgency.

In theory, climate change should be the ultimate burning platform. Yet climate is not the only driver of urgency. For many countries—particularly those that import a large share of their energy—the transition is also about resilience and sovereignty: reducing exposure to volatile global fuel markets, strengthening domestic energy systems, and insulating economies from geopolitical disruption. In that sense, hydrogen and other new energy infrastructures are not only climate tools, but instruments of long-term energy security. In practice, it remains abstract for most people — distant in time, diffuse in responsibility, and easily displaced by more immediate concerns. Energy systems fail quietly. Progress is slow, incremental, and invisible. The result is a paradox: we face a structural transformation on the scale of industrial revolutions, yet we are trying to deliver it with the emotional intensity of a regulatory update.

If the public were truly demanding this transition — if voters were loudly insisting on faster action, clearer direction, and long-term consistency — political systems would respond very differently. Democracies move quickly when urgency is unmistakable. The fact that progress remains hesitant is itself evidence that society has not yet crossed that threshold.

This is not a moral judgement; it is a systems observation.

John Kotter's well-known framework for organisational change, first set out in *Leading Change* (1996) and later popularised for wider audiences in *My Iceberg Is Melting* (2006, with Holger Rathgeber), begins with a deceptively simple first step: create a sense of urgency. Without it, none of the subsequent stages matter. Coalitions stall. Visions remain theoretical. Resistance hardens. Change initiatives decay into box-ticking exercises rather than transformations.

Viewed through this lens, the energy transition is behaving exactly as one would expect of a system that has skipped its first step.

There are guiding coalitions — across industry, finance, and policy — but they remain fragmented. There are visions, but they are often technical, contested, or regionally inconsistent. Communication is constant, yet rarely resonant. Barriers are well understood but rarely removed at pace. Short-term wins are celebrated but not felt by society at large. And without public ownership, change never truly anchors.

Urgency, in this context, does not mean panic. It does not require fear-based messaging or apocalyptic narratives. In fact, those approaches often backfire, leading to disengagement rather than mobilisation. Effective urgency is grounded in relevance: making the future feel personal, immediate, and consequential to everyday life.

There is another, subtler risk here. Many of the people working most closely on hydrogen are already convinced. When you meet them — in projects, in labs, in control rooms, in policy teams — their commitment is rarely in doubt. They are motivated, informed, and often deeply passionate about the transition that they are helping to build. That passion has been essential in getting hydrogen this far.

However, passion can become counterproductive if it turns into evangelism. The wider public does not need to be converted; it needs to be reassured. Urgency is not created by preaching the future, but by making it feel relevant, credible, and aligned with everyday priorities. When messaging becomes ideological, it risks triggering resistance rather than momentum — especially among those who feel they are being asked to change before they understand why.

The challenge, then, is not to amplify belief within the system, but to translate it outward — calmly, pragmatically, and without moral pressure. Urgency grows when people recognise themselves in the story, not when they feel spoken at by it.

In some regions, hydrogen has also been pulled into political identity — framed as a “liberal” solution rather than an energy one. It echoes the nuclear debate, where ideology can overwhelm engineering reality. Energy systems do not care about identity; they care about physics, reliability, and economics. Hydrogen scales fastest when it is treated as infrastructure, not symbolism.

At present, much of the hydrogen conversation lives in conference halls, strategy documents, and pilot projects. It is discussed in terms of megawatts, subsidies, and infrastructure corridors. Necessary language — but insufficient on its own. People do not mobilise around levelized costs or regulatory acts. They mobilise around stories they recognise themselves in.

This is where social perception becomes decisive. New energy systems must not only be cleaner and more efficient; they must feel normal, reliable, and inevitable. They must be framed not as sacrifice, but as continuity — an evolution of how societies power themselves, not a rupture with modern life. Until that perception shifts, urgency will remain confined to specialists rather than shared broadly.

History suggests that urgency often emerges only when change is no longer optional. Energy shocks, geopolitical disruptions, and price volatility have periodically accelerated action — but always reactively. The challenge now is to generate urgency before crisis dictates the terms.

This requires leadership that understands change as a human process, not just a technical one. It requires consistency over electoral cycles, clarity over complexity, and narratives that connect long-term systems thinking with short-term lived experience. Most of all, it requires acknowledgement that building the future energy system is as much about earning belief as it is about deploying assets.

The tools are ready. The pathways are visible. Capital is waiting for certainty. What remains unresolved is whether society as a whole decides that waiting is no longer acceptable.

Until that decision is made — collectively and emotionally, not just analytically — progress will continue to feel slower than it should. Not because we cannot move faster, but because we have not yet convinced ourselves that we must.

Hydrogen Is Already Everyday — and That's the Point

Throughout this book, I have deliberately focused on practical examples: trains running quietly through Bavaria, homes heated without carbon, aircraft taking off in near silence. These stories span sectors, geographies, and scales — and together they describe a system that is no longer hypothetical. The people who appear in these stories are mostly real. They are not symbolic figures or fictional stand-ins, but individuals working in industry, research, operations, and infrastructure today.

This matters, as by the time a technology can be demonstrated across such a wide range of real-world contexts, the question is no longer whether it works. The question becomes whether it can be trusted when conditions are hardest, visibility is highest, and failure is not an option. That trust is built not only through design and engineering, but through the judgement, experience, and responsibility of the people involved.

That is where this final story belongs.

Extreme H is not included here because it is dramatic, fast, or visually compelling — although it is all of those things. It is included because it brings together, in one place, many of the themes explored throughout this book: safety, systems integration, digital coordination, logistics, public scrutiny, and the need to learn at speed rather than wait for certainty. It also brings together real teams, operating in real conditions, with real consequences.

In Extreme H, hydrogen is not protected by pilots, caveats, or controlled environments. It is asked to perform repeatedly, under pressure, in public, with no room for abstraction. Every race weekend compresses years of design assumptions, operational decisions, and risk management into a matter of days.

When hydrogen functions reliably in that setting, it resolves a different kind of question — not whether the chemistry works, but whether the system holds together.

Extreme H shows what happens when technology, infrastructure, teams, and processes are aligned and then tested without the comfort of perfect conditions. It demonstrates resilience, not just performance; behaviour, not just capability. In doing so, it offers something more valuable than a proof of concept: it offers confidence.

This is why Extreme H sits at the end of this book.

Not as a spectacle, and not as an exception — but as a synthesis. A place where many strands of the hydrogen transition converge, are stressed, and still hold. A reminder that progress is not proven in theory, but in motion — and that the people behind these systems are already doing the work.

Extreme H — Making the Invisible Visible

It is early October 2025 in the Saudi Arabian desert outside Qiddiya City — a municipality rising from the sand with the ambition of a new Las Vegas, a place designed for spectacle, adrenaline, and reinvention. The light here is harsh and clean. Heat presses down on the paddock like a physical weight. Beyond the temporary fencing, the desert stretches out in pale waves, interrupted only by the dark ribbons of a racing course carved into sand and gravel.

Twenty-seven-year-old Swedish rallycross driver Kevin Hansen sits strapped into the cockpit of a Spark Racing Pioneer 25, hands steady on the wheel, eyes fixed forward as the starting sequence runs. From the outside, the car looks like what it is: an off-road racing machine designed to be punished. Wide stance. Aggressive aero. The kind of posture that suggests it wants to fight the ground

itself. But underneath the bodywork, the logic is different to anything motorsport has fielded at this level before.

This is not internal combustion chasing its final refinement. It is a hydrogen-electric vehicle — a fuel cell system producing electricity from chemistry, feeding motors that deliver torque instantly. What exits the exhaust is not smoke or carbon, but water vapour.

The lights flick from red to green. Kevin's foot drops. The Pioneer 25 lunges forward and the desert detonates behind it — sand exploding in a fan, traction arriving and leaving in sharp pulses as the tyres bite and release. To spectators it is thrilling sport in its purest form: speed, noise, risk, judgement. To the engineers and organisers watching the data, it is something else as well: a full-scale experiment, running at the limit.

Hydrogen has lived for decades in places most people never see — behind refinery fences, inside industrial plants, in pipelines and tank farms where the public imagination rarely goes. What Extreme H attempts is the opposite. It drags hydrogen into daylight. It puts it under pressure, under scrutiny, and then broadcasts it to an audience that does not care about policy, or levelized costs, or strategy documents. People care if it works. People care if it feels real.

That is precisely why the championship exists.

It was conceived by Alejandro Agag, who had already done something similar with Extreme E: proving that electric off-road racing could capture attention and legitimacy at the same time. With Extreme H, the ambition is bolder. It is not simply to compete; it is to demonstrate — to compress years of technical argument into something visible, kinetic, and emotionally understood.

But what makes Extreme H meaningful is not just the car. It is the ecosystem required to make it run.

Hydrogen at an event like this is not a single technology. It is production, logistics, storage, power generation, safety systems, refuelling, digital coordination, and operational discipline — all functioning together, under conditions that do not forgive improvisation. The leadership team reflects that. Ali Russell has overseen the commercial, media, partner, and sustainability dimensions — ensuring that what happens here is not just executed, but communicated in a way the world can absorb. James Taylor has carried the sporting and operational load — course, officials, race control, safety, the unglamorous machinery of making sure everything works when the clock starts.

Around them sits a network of partners who are, in effect, staking their reputations on hydrogen's credibility.

GeoPura brings its Hydrogen Power Unit technology to power the event site. Power Logistics integrates that power into operations. ENOWA — part of NEOM — supplies the green hydrogen that fuels both cars and infrastructure. Siemens contributes industrial software, digital twin capability, and safety expertise. PRF (Hydrogen Refuelling Solutions) provides the fuelling systems. Symbio supplies the 75 kW fuel cell that sits at the heart of the Pioneer 25. Aurora Media Worldwide ensures that the story is not lost inside the paddock but translated into something legible beyond it.

Taken together, this is more than a supply chain. It is a miniature hydrogen economy running in real time: molecules arriving, being stored and converted, energy being delivered, vehicles being refuelled, safety being monitored and performance being measured. It is hydrogen not as an idea, but as a working system.

Then — as it always does in motorsport — the system meets its moment of truth.

During one high-speed run, a Pioneer 25 hits a pocket of sand at speed. The car lifts, rolls, tumbles. For a second, instinct takes over in everyone watching: the collective flinch that comes when you see mass and momentum lose control. The car comes to rest in a haze of dust.

Then, the driver moves.

He climbs out, shaken but unhurt. The marshals move in. Engineers watch for the signals that matter to them — the ones the public rarely thinks about. Tanks remain intact. The fuel cell contains its energy. The safety systems hold. There is no fire, no explosion, no hazard — only dust settling back onto the desert.

In one tumble, a fear that has haunted hydrogen in the public mind for decades quietly loses its power.

For engineers, it is not drama. It is validation. The design assumptions were correct. The safety architecture worked. Hydrogen behaved as intended under stress. A single incident achieves what thousands of slides often fail to: it makes safety feel real.

Racing has always served this purpose. It accelerates learning by making failure expensive and immediate. One year of competition can compress a decade of traditional development because every lap creates data. Every jump stresses tanks, cooling, electronics, and control systems. Every acceleration cycle asks questions

about durability that laboratory conditions rarely replicate. The desert becomes a classroom. The race becomes an experiment with immediate feedback.

Yet, perhaps the most important thing about Extreme H is how ordinary it can feel while you are watching it.

The Pioneer 25 looks, to most eyes, like any other rallycross car. Teams run familiar routines. Drivers push, slide, jump, and drift as they always have. There is the same urgency, the same problem-solving under time pressure, the same culture of engineering competence. Hydrogen's promise is not that it will replace the world with something alien. It is that it will flow into the world we already recognise — into vehicles, depots, logistics chains, and infrastructure — and allow them to operate cleanly without demanding a wholesale reinvention of daily life.

Extreme H is a public rehearsal of that idea.

It turns hydrogen from something discussed into something witnessed. It makes the invisible visible, not by explaining it, but by showing it under pressure. It does not ask for belief. It earns it, one lap at a time.

The effect is subtle but important. People do not leave an event like this thinking about electrochemistry. They leave thinking: it worked. It was fast. It was safe. It looked like the future and felt like the present.

That is how technologies cross thresholds — not when experts agree they are ready, but when the wider world begins to treat them as normal.

Extreme H is not the hydrogen economy. But it is a concentrated proof of what the hydrogen economy requires: collaboration across disciplines, integrated systems rather than isolated components, and the willingness to put hardware in the real world before the world feels fully prepared.

By the end of the weekend, the cars return to their tents, data downloads complete, dust settles over the dunes. The spectacle passes. What remains is less theatrical and far more valuable: the quiet accumulation of credibility.

Hydrogen, long discussed and often doubted, has been made to perform — safely, repeatedly, publicly. Not as a promise, but as a system.

And that, in the end, is how transitions start to feel inevitable.



Extreme H — hydrogen under pressure, in motion, and in public.
Image courtesy of Extreme H. © 2026 Extreme E Ltd. All rights reserved. Used with permission.

At the beginning of this book, we borrowed a line from *The Hitchhiker's Guide to the Galaxy* — a reminder that complex systems rarely yield to simple answers.

But the lesson was never that the answer does not exist. It is that knowing the answer is not the same as acting on it.

This book has argued that hydrogen is not a silver bullet, nor a distraction, nor a future technology waiting for its moment. It is a connective tissue — between sectors, time horizons, and energy vectors — that makes the wider transition physically and operationally possible.

The question, then, is no longer whether the system can work. It is whether we are willing to move before everything feels comfortable.

Progress Is Supposed to Feel Uncomfortable

The energy transition will never feel tidy. If it does, it is almost certainly moving too slowly.

A line often attributed to racing driver Mario Andretti captures this tension perfectly: "If everything seems under control, you're not going fast enough."

Throughout this book, a recurring pattern has emerged: the technology exists, the pathways are visible, and the capabilities are already in place. What repeatedly holds progress back is not a lack of solutions, but an insistence on certainty before action. We wait for perfect policy alignment, flawless economics, complete public consensus, and fully de-risked investment cases. By the time all of those conditions are met, the moment for leadership has usually passed.

Real transitions do not unfold from a position of comfort. They begin with unease — with the recognition that the status quo is no longer sufficient, even if the alternative is not yet fully defined. That unease is not a failure of planning; it is a signal that the system is stretching beyond familiar boundaries.

In complex systems, speed and control exist in tension. Move too slowly, and momentum evaporates. Move too cautiously, and incumbency hardens. Progress requires operating in a space where not everything is fully known, where trade-offs are real, and where decisions must be made with incomplete information. The absence of discomfort is not reassurance; it is inertia.

Extreme H illustrated this dynamic in its most concentrated form. There, hydrogen was not perfected behind closed doors or proven through endless modelling. It was placed under pressure, exposed to risk, and asked to perform in public. The value did not come from everything being controlled, but from the system proving it could cope when it was not. Learning happened at speed, not at rest.

The same principle applies to the wider energy transition — and particularly to hydrogen. It sits at the intersection of energy, industry, geopolitics, and society. It cannot be deployed sequentially or incrementally in isolation. Infrastructure must be built ahead of demand. Markets must be shaped before they exist. Capital must be committed before certainty arrives. None of this feels controlled in the conventional sense — and it was never going to.

The temptation, especially in large organisations and political systems, is to mistake caution for responsibility. But there is a difference between managing risk and avoiding it altogether. History rewards those who act early, learn fast, and adjust course — not those who wait for the road to be fully paved before stepping forward.

As explored throughout these chapters, the heroes required for this transition are already on the field. The expertise exists. The intent exists. What remains is the collective willingness to accept that progress will feel uncomfortable — and to move anyway.

If this transition is to succeed, we must stop asking whether everything is under control and start asking whether we are moving fast enough to matter.

If everything feels safe, predictable, and perfectly managed, the truth may be simpler — and more dangerous.

We are not yet going fast enough.

Afterword

I never, in my wildest imagination, envisaged a world in which I would write a book. I have always enjoyed reading, but I never considered myself capable of actually producing one.

A series of moments brought me here. A conversation with my son, mentioned at the very beginning. The death of a close personal friend. A discussion in a bar in Guiyang, China, with the head of my business's technical support team. A senior manager remarking that we did not have enough experts within the company. In addition, and perhaps most importantly, a persistent feeling that began to take hold in the spring and summer of 2025 — a sense that ideas were forming, chapters were taking shape in my head, and that eventually I would have to sit down and begin.

I remember spending much of a summer holiday with blocks of text running through my head — even walking through Japan, trying to take in the scenery and the culture, I still had hydrogen molecules bouncing around somewhere in the background.

Once I did, the process took on a life of its own. What started as a structured attempt quickly became something far less linear. I began writing sequentially, only to hit an early block. From there, I found myself jumping — almost like a grasshopper — from section to section, following wherever the ideas led. In doing so, I realised that the problem was not the writing, but the structure. I had the order of the book wrong. That realisation was both helpful and uncomfortable, as it meant going back and rewriting the entire opening.

That same "grasshopper" approach brought its own challenges. Writing in bursts of inspiration inevitably led to repetition. I have done my best to remove it and place ideas where they belong, but if some repetition remains, I ask for your patience — and attribute it to a first-time author learning the craft as he goes.

This process largely unfolded in the margins of a full-time job that demanded much of my time. For around six months, most of the writing happened in the

evenings and at weekends — often built on hurried notes taken during the day as ideas surfaced, to be translated later into something more coherent. It was, in many ways, an exercise in fitting the work in whenever space allowed.

If there is one thing that helped me write this book, it was reading. I have always loved books, but like many, I had increasingly drifted toward the distractions of modern technology. Last year, however, I underwent major surgery, which took me away from work for several weeks and into a longer period of rehabilitation. That time gave me something unexpected: space. Space away from work, and space away from constant digital noise. I found myself reading extensively again, across a wide range of topics.

I would, however, not recommend major surgery as a strategy for writing a book.

If there is one lesson I would share, it is this — if you want to write, you must read.

To anyone reading this who feels they could never do something similar: you can. Start with a topic you care about. Try to understand it deeply. And begin.

When I set out to write this book, I was aware that there were areas where my knowledge was not as strong as I wanted it to be. Addressing that required time, effort, and a willingness to explore beyond my comfort zone. The perspective presented here is inevitably shaped by my own professional lens — particularly through automation and systems — and by a European viewpoint, as this is where I am based and where most of my direct experience and touchpoints lie. Others, closer to different parts of the value chain or operating in different regions, may see things differently. I hope, however, that this viewpoint still provides a useful contribution to the broader discussion.

That said, I do not claim to be an expert in everything. I may understand the game of rugby union reasonably well, but even after many years in the energy sector, I would never claim to know it all. I try, instead, to learn something new each day.

So, if you read this book and find points where you disagree — or even think I am wrong — that is entirely fair. In fact, it is part of the intention. I hope this book sparks discussion — because across the energy transition there remain many valid, and often conflicting, perspectives, and these pages reflect my own interpretation of that landscape. As I noted in the closing chapters, we are not lacking ideas in the energy transition; we are often lacking urgency. If this book contributes, even in a small way, to moving that conversation forward, then for me it has served its purpose.

I am also under no illusion that some colleagues and friends in the industry will pick up the phone for what they might call a "healthy debate." If we are fortunate, those conversations will happen over a beer rather than in a meeting room.

As for whether I would write another book — I genuinely do not know. But having gone through this process, I can already see several ideas that could be explored. Perhaps revisiting the energy transition in five or ten years to see how we have progressed. Perhaps a broader look at the energy transition beyond hydrogen. Or something entirely different — stories from my time as a commissioning engineer, in the spirit of Paul Carter's *Don't Tell My Mum…*, which would almost certainly sit in a different section of the bookshop. Maybe even, one day, a reflection on life in the corporate world.

For now, however, I am simply glad to have finished this one.

I hope you have enjoyed it, perhaps learned something, and most of all, I look forward to the conversations that may follow — wherever our paths may cross.

Thanks

I could not have written this book without the support of my family — Christina and Fynn — who gave me the freedom and license to embark on this project. Writing in my free time meant evenings and weekends were often lost, with me either in the office or 'present' on the sofa while not truly present as *The Rookie* or *Stranger Things* played in the background. Too often a Sunday walk became "just five minutes more," which turned into hours — always without complaint. Many nights also ended with me saying I'd be upstairs shortly, only to finally make it there at 2 a.m. Thank you both for the patience, understanding, and support to follow this path.

Writing the book itself turned out to be the easier part. What followed was far less predictable and often outside my control. For helping me navigate that, I would like to thank Stuart Cooper, Eva Augsten, and Vincent Brocard for their guidance, insights, and support.

I would also like to thank Neonbold for transforming rough sketches into the illustrations that appear throughout these pages. My thanks also go to Beecom, whose creativity, professionalism, and support helped bring the original *Industry Edition* to life and provided the foundation from which this published edition has grown. My sincere thanks also go to Alex Grossmann and the team at Grossmann Verlag for their guidance, editorial expertise, and commitment in bringing this first edition to publication.

I would also like to thank everyone who received and read the Industry Edition and the pre-publication manuscript. The volume of feedback, encouragement, and positive response that followed was both unexpected and deeply appreciated. Knowing that the ideas in these pages resonated with people working across different parts of the hydrogen industry gave me considerable confidence as the book moved towards publication.

A few people went further, taking the time to read the manuscript critically and share detailed comments, technical insights, corrections, and suggestions. Those contributions helped me challenge some of my own assumptions, clarify argu-

ments, and tighten the final manuscript. I will not name everyone individually here, but you know who you are, and I am genuinely grateful for the time, expertise, and honesty you gave me.

More broadly, thank you to everyone working across the hydrogen community. Over the years, I have had the privilege of meeting many of you at conferences, trade fairs, project meetings, and panel discussions around the world. Your ideas, challenges, questions, and willingness to share your experiences have shaped my own thinking far more than you probably realise. This book is, in many ways, a reflection of those countless conversations.

Whether you agree with every conclusion or not, I hope these pages encourage discussion, challenge assumptions, and contribute in some small way to the collective effort of building a more sustainable energy future.

Thank you for reading.

About the Author



Robert Atkinson is an industrial professional working at the forefront of the energy transition, with over two decades of experience spanning engineering, sales, and business development in electrification, automation, and large-scale energy systems.

Born in Ashford, Middlesex, and raised in West London, he studied Electronics and Communications Engineering at Nottingham Trent University, including a year in industry at Forschungszentrum Karlsruhe. He began his career in the UK as a test and commissioning engineer, working on bespoke PLC automation systems and medium-voltage drives for applications across marine and oil and gas applications, including pipelines, storage, and offshore systems.

Since relocating to Germany in 2007, he has built an international career focused primarily on commercial and business development roles, working to bring complex energy and industrial solutions into practical, deployable projects. His experience sits at the intersection of technology, markets, and execution.

Drawing on this combined technical and commercial perspective, *The Hydrogen Age* reflects a pragmatic view of the energy transition — grounded not only in how systems work, but in what it takes to make them viable and deliverable in the real world.

He is based in Germany with his family. Outside of his professional work, he has been deeply involved in rugby as a player, coach, and referee, including coaching

at 2nd Bundesliga level and contributing to a German vice-championship in 7s rugby with Karlsruhe in 2003. He is also a keen reader and traveller.

Illustration Sources and Rights Information

All illustrations contained in *The Hydrogen Age* were created specifically for the book. The majority are original illustrations developed by the author. Where an illustration draws upon an identifiable external source, this is acknowledged directly in the caption beneath the relevant illustration.

The following illustrations contain adapted or externally sourced information:

Page	Source / Rights Information
p. 4	Adapted from a climate visualisation originally attributed to Robert Corell, Heinz Center. The illustration has been redesigned and updated specifically for this book. Attribution is included in the caption.
p. 8	Adapted from Mikael Höök, *Historical Energy Use of the United States from 1635 to 2000*. The illustration has been redesigned specifically for this book. Attribution is included in the caption.
p. 111	Adapted from the International Association of Oil & Gas Producers (IOGP), *Oil in Everyday Life*. The underlying concept and information were used as a reference; the illustration has been substantially redesigned specifically for this book. Attribution is included in the caption.
p. 199	Adapted from World Economic Forum analysis concerning employment impacts of the energy transition. The illustration has been redesigned specifically for this book. Attribution is included in the caption.
pp. 217–218	NOW GmbH – German/EU regulatory framework. Adapted from NOW GmbH. The regulatory overview has been recreated/redesigned specifically for this book. Attribution is included directly in the caption.
p. 326	Adapted from Xu et al. (2023), *Rational Design of Hydrogen Evolution Reaction Electrocatalysts for Commercial Alkaline Water Electrolysis*. Additional information has been added by the author and the illustration has been redesigned specifically for this book. Attribution is included in the caption.
p. 388	Extreme H illustration/material. Used with permission. Permission has been obtained directly from the relevant rights holder.

Further illustration: One illustration was created specifically for the book and substantially develops a simple conceptual categorisation encountered in an unattributed online graphic. No identifiable original source could be established. The final illustration uses a new visual structure and contains substantially developed content.

All other illustrations: Original illustrations created by the author/specifically for *The Hydrogen Age* and do not reproduce third-party artwork.