Outdoor Connect
management consulting

The First Serious Fusion Strategy

Why First Light Fusion treats energy as a product with a unit, a cost curve, and a control point – not as a science project.

By Reinout Schotman

Nuclear Fusion at Human Scale (foto: First Light Fusion Ltd.)

Every serious board conversation today circles the same question in different language: Can we keep growing without breaking the climate or our political stability? So far, the honest answer has been: *not really*. Our prosperity still runs on combustion. Fusion is not a science project; it is the only credible technology that could decouple long-term prosperity from combustion at planetary scale.

This article argues that nuclear fusion could change that equation structurally, and that First Light Fusion (FLF) is one of the few players treating fusion as an engineered product rather than a science project.

If you are a policymaker, investor, or business leader, you do not need another physics explainer. You need to know:

- Why prosperity and energy are inseparable – and why that has become an existential problem.
- Why fusion is structurally superior to fission, yet strategically underused in policy and investment.
- How the traditional "big is beautiful" fusion model trapped the field in complexity.
- How impact-driven inertial fusion and fast-ignition architectures – and FLF in particular – reframe the challenge as an engineering and product problem.

- Which challenges remain, and why they are now mostly about choice and commitment, not new laws of nature.
- What a working fusion unit would mean for economies, geopolitics, and climate risk.

That is the journey this article takes for boards, policymakers, and investors.

1. How prosperity and energy locked us into a trap

Over the past century, every major rise in prosperity has ridden on a matching rise in energy use. Countries with higher GDP per capita almost always consume more energy per person. That correlation is not an accident; it is the operating logic of the modern economy.

The problem is what sits underneath it.

Today, the energy system accounts for more than three-quarters of global greenhouse gas emissions. Most of that comes from burning coal, oil, and gas for power, heat, transport, and industry. At the same time, our dependence on fossil fuels and critical materials anchors geopolitics: who controls pipelines, straits, LNG terminals, and mining supply chains controls leverage.

In structural terms, we have built a three-way constraint:

- We want continued economic growth.
- We need deep, rapid emissions cuts.
- We seek greater energy security and less dependence on unstable actors.

Under the current energy architecture, these three aims conflict. We can optimise at the margins – more renewables, more efficiency, better grids – but prosperity is still chained to combustion and contested supply.

A different kind of energy source would not just be a climate measure. It would redraw the economic and geopolitical playing field.

That is the context in which nuclear fusion should be discussed: not as a scientific curiosity, but as a potential escape route from this structural trap – and the rest of this article follows that line from physics to product.

2. Fission versus fusion – structurally, fusion is the better bet

Today's low-carbon baseload is dominated by **nuclear fission** and **hydropower**, with wind and solar growing fast but still dependent on storage, interconnectors, and demand management.

Fission has delivered decades of carbon-free electricity, but from a strategic and societal risk lens it has three enduring liabilities:

1. Long-lived high-level waste that must be secured for tens of thousands of years.
2. Non-zero meltdown risk, requiring multiple safety layers and constant public reassurance.
3. A fuel and technology set associated with weapons programmes, making both politics and public opinion fragile.

Nuclear fusion, in principle, solves most of these issues:

- It produces no CO_2 at the point of generation.
- It generates far less and shorter-lived radioactive waste; most radioactivity resides in reactor structures that can be replaced and managed on human timescales.
- It has no runaway chain reaction: if conditions are disturbed, the fusion process simply stops.
- It relies on fuel derived from water and lithium salts, not on enriched uranium or plutonium.
- If you were designing a planetary energy system from first principles, you would pick fusion over fission.

And yet fusion has remained marginal in real economic and policy strategy. Not because the physics is uninteresting, but because the way we chose to pursue fusion has been strategically misaligned with the goal of creating an industrial product.

Fusion only becomes strategically relevant once it is engineered as a repeatable, bankable product.

3. How fusion got trapped in the "big is beautiful" mindset

For roughly six decades, mainstream fusion programmes have followed one dominant line of thought:

If we build bigger, more complex machines, commercial fusion will eventually emerge.

That instinct was not irrational. The amount of energy that must be delivered to fusion fuel in a very short time is enormous. Most of the field concluded that the only way to do this was with ever larger magnetic systems or laser arrays: if the driver had to be huge, the machine around it would be huge as well.

That logic produced two major architecture families:

- Magnetic confinement – tokamaks like ITER in France, and stellarators with extremely intricate magnet geometries.
- Laser-driven inertial confinement, exemplified by the National Ignition Facility (NIF) in the US.

These are engineering marvels. Strategically, they are deeply flawed as routes to deployable energy technology.

3.1 Scale without a product

ITER is designed as the world's largest experimental tokamak. Its projected cost has escalated far beyond initial estimates, and first full-power operation has slipped by decades. That might still yield valuable science. But as a template for a commercial fleet, it fails a much simpler test: there is no clearly defined unit you can replicate and sell.

A technology becomes economically transformative only when there is a standard unit – a turbine, a reactor design, a module – that you can build again and again. Most large fusion projects still cannot describe that unit in concrete, financeable terms. By contrast, combined-cycle gas plants and modern wind turbines only became dominant once there was a clear reference unit that engineers, financiers and regulators could all recognise, certify and repeat.

3.2 Institutions before markets

Mainstream fusion projects are governed as international research collaborations. Their success metrics are plasma parameters, publications, and diplomatic continuity – not levelised cost of energy, uptime, or manufacturability.

This is not a criticism of the scientists. It is a governance choice. We built institutions optimised for knowledge, not for product.

3.3 Complexity mistaken for progress

As systems grew more complex, so did the models that described them. Over time, the field drifted into what Outdoor Connect calls the illusion of knowing: more data, more parameters, more simulation – without a proportional increase in true information about how to build a working product.

In strategy terms, complexity became a substitute for clarity.

The question fusion should be answering is simple: *What is the smallest, simplest machine that can produce net energy in a way we can manufacture, operate, and maintain at scale?*

The traditional fusion ecosystem rarely framed the problem that way. First Light Fusion does.

4. Inertial fusion and First Light Fusion – a different line of thought

Before looking at the physics, it is worth understanding where First Light Fusion (FLF) comes from and why its logic is different.

4.1 Origin: from impact physics, not plasma orthodoxy

First Light Fusion is a British fusion company based near Oxford, founded in 2011 by Dr. Nicholas Hawker and Professor Yiannis Ventikos as a spin-out from the University of Oxford. Its roots lie in hypervelocity impact physics – the study of how materials behave under extreme compression when struck at very high speed.

That starting point matters. FLF did not emerge from the tokamak or laser-fusion mainstream. Its founding question was not:

"How do we out-ITER ITER?"

but:

"Can we reach fusion conditions using mechanisms that are radically simpler and more buildable than conventional fusion machines?"

Where the mainstream treated complexity as inevitable, FLF treated complexity as a design flaw.

4.2 Philosophy: let engineering lead

FLF's philosophy can be summarised in four principles:

1. **Engineering first, physics second.**

   The company starts from what can be manufactured, operated, and maintained as an industrial product. Fusion is then pursued along paths that fit those constraints.

2. **Known physics over exotic regimes.**

   FLF's impact-driven inertial fusion and fast-ignition concepts rely mostly on hydrodynamics, shockwaves, and material science – domains where theory and experiment are already robust, and where industry has long experience.

3. **The reactor must be a product, not a monument.**

   The goal is not to build the largest experiment, but a standardised reactor concept that can be deployed as a fleet, similar to families of gas turbines or aircraft.

4. **Use geometry, not brute force.**

   Instead of betting on ever-more-powerful lasers or magnets, FLF invests in target design – the disposable component that amplifies the driver pulse into the extreme pressures and temperatures needed for fusion.

Strategically, this is exactly the kind of reframing that Clayton Christensen, Peter Drucker and Bruce Henderson each, in different ways, described as the essence of advantage: change the game, not just your performance within it.

In private-equity language, FLF behaves less like a national lab and more like a focused mid-market disruptor: pick an architecture you can actually own, scale, and defend.

4.3 What FLF actually does

In its early years, FLF used hypervelocity projectiles as the fastest and cheapest way to test its core idea: that an intermediate structure – an amplifier – can take a relatively modest input pulse and concentrate it into the extreme pressures needed for fusion.

Those experiments demonstrated that the amplifier physics works. But that does not mean the commercial reactor will be a "projectile fusion" machine. In the FLARE concept, the projectile experiments are a proving ground for the amplifiers and the fast-ignition architecture, not a blueprint for the final driver.

In simplified form, the FLARE architecture does three things:

It uses an impact or other pulsed driver to create a strong but controlled pressure pulse on an amplifier structure.

The amplifier performs largely isentropic compression of the fuel for only a small fraction of the total energy needed – on the order of a few percent.

A separate, concentrated hotspot ignition then starts a burn wave through the fuel, so that most of the output energy comes from the fusion burn itself, not from the driver.

The critical insight is that the amplifier and the burn physics do most of the work, not the driver.

The firing machinery can therefore remain relatively simple and robust. The complexity – and thus the innovation – is concentrated in a replaceable component and in the burn design that can be simulated, tested, and iterated in large numbers.

Strategically, that matters more than the physics detail. It means FLF is placing the difficulty where industry knows how to learn quickly: in component design, manufacturing, and burn optimisation – not in building ever larger, irreplaceable infrastructure.

4.4 In practical terms

FLARE defines a standard reactor unit rather than an open-ended experimental facility and uses hardware assumptions that are grounded in existing industrial capabilities instead of hypothetical future breakthroughs. It also positions FLF's technology as the "fuel" of a wider fusion ecosystem: plants could be developed, financed, and operated with partners, while FLF owns the core know-how and the supply of amplifiers and targets.

This is classic control-point logic. If fusion becomes a real market, the entity that controls amplifier and target design, IP, and high-volume manufacturing will sit at the structural chokepoint that defines performance and cost.

The other crucial shift is **gain**. In the FLARE architecture, FLF works with design gains on the order of **1,000** – meaning that the fusion burn releases roughly a thousand times more energy than is delivered by the driver. From a physics perspective there is no hard technical ceiling on that gain; the practical limit is what a reactor can reasonably contain and cycle.

High gain changes everything. It allows the driver to be much smaller and cheaper than traditional fusion programmes assumed, while still delivering a very large, fast pulse of energy into the fuel via the amplifier. Internally, FLF estimates that this could translate, in realistic scenarios, into levelised costs of electricity on the order of **$50/MWh** – more conservative than earlier, more aggressive $25/MWh modelling, but still highly competitive.

The precise number will ultimately depend on engineering outcomes and finance. The directional point is clear: if FLARE-like concepts can combine very high gain with modest driver energy, fusion can become not just clean and secure but one of the lowest-cost scalable sources of dispatchable power – able to run as firm baseload and flex output to complement variable renewables.

5. The remaining challenges – hard work, not new physics

None of this makes fusion inevitable. Significant challenges remain. What has changed is their nature: they are now predominantly engineering, industrial, and governance problems, not missing pieces of physics.

5.1 Target manufacturing and cost

A commercial reactor will require vast numbers of targets each year. These must be produced with:

- high precision,
- high consistency,
- and low unit cost.

This is non-trivial. But it sits in the same universe as advanced battery manufacturing, semiconductor packaging, and precision automotive components – industries that already function at comparable levels of volume and quality control. The learning curves and investment logic are familiar.

5.2 Shot rate and machine robustness

For fusion to be economic, the system must fire frequently and reliably. That, in turn, requires:

- accelerators and chambers that can withstand continuous operation,
- components that can survive repeated shock and radiation cycles,
- maintenance philosophies that minimise downtime and cost.

This looks less like speculative science and more like the world of jet engines, gas turbines, and high-temperature chemical plants: difficult, but in scope for established engineering disciplines.

5.3 Materials and neutron activation

Neutrons produced by fusion will activate structural materials. This means there will be radioactive waste. However:

- the half-lives involved are typically far shorter than those of high-level fission waste,
- components can be designed for replacement and eventual recycling,
- ongoing advances in materials science and radiation-resistant alloys can be directly applied.

Again, this is a design challenge, not a fundamental showstopper.

5.4 System integration, regulation, and finance

Perhaps the toughest challenges sit outside the reactor:

- How will fusion plants be permitted, regulated, and insured?
- How will they connect to grids already grappling with intermittency and congestion?
- What financing structures will support high upfront CAPEX with long asset lives?

Those questions cannot be answered by FLF alone. They require policy design, market structure, and political will. The strategic point is clear:

First Light's fast-ignition fusion concept is no longer waiting on a miracle; it is waiting on decisions.

Timelines remain uncertain. The point of this article is not to predict a commercial start date, but to highlight that the bottleneck has shifted from pure physics to strategy, capital, and governance.

6. What a working fusion unit would mean for economies and the world

If First Light Fusion – or any similar approach – succeeds in delivering a replicable, commercially competitive fusion unit, the consequences would be structural.

6.1 Economic structure

At projected LCOE levels in the region of $50/MWh, FLARE-like fusion plants could compete with today's cheapest renewables and cost only a fraction of new nuclear – while providing fully dispatchable, controllable power. That combination of low cost and dispatchability would make fusion not just another clean source, but a structural reset of power-sector economics.

Energy scarcity would ease as a constraint. Long-term growth models would no longer assume ever-rising fossil fuel inputs or chronic price shocks.

Heavy industry could decarbonise without offshoring. Steel, cement, chemicals, and data-intensive sectors like AI could be located based on talent and infrastructure rather than proximity to cheap fossil fuel – especially for energy-importing economies like the EU, Japan, or the Netherlands.

New industrial clusters would emerge. Countries capable of building, operating, and exporting fusion plants – and their key components – would develop durable industrial ecosystems, similar to aerospace or semiconductor leaders today.

6.2 Geopolitics

Power would shift from resource owners to technology owners. Oil and gas reserves would lose strategic gravity. States leading in engineering, materials, and systems integration would gain influence.

Energy security would become programmable. Instead of competing over finite reserves, countries could choose to deploy domestic fusion capacity as a matter of industrial policy.

6.3 Climate and planetary stability

Given that the energy system produces more than 75% of global greenhouse gas emissions, a viable fusion technology would fundamentally change decarbonisation arithmetic. Ambitious climate policy would face lower economic and political resistance, because deep emissions cuts could be achieved without threatening energy security or industrial competitiveness.

Electrification of hard-to-abate sectors would become more plausible at scale, and emerging economies could grow without repeating the high-emission paths of Europe, North America, and parts of Asia.

Fusion would not replace the need for renewables, efficiency, or behaviour change. But it would shift the debate from "how do we manage permanent scarcity?" to "how do we design responsibly around abundance?" – a very different strategic conversation.

Given that the energy system produces more than 75% of global greenhouse gas emissions, a viable fusion technology would fundamentally change decarbonisation arithmetic:

- Ambitious climate policy would face lower economic and political resistance.
- Electrification of hard-to-abate sectors would become more plausible at scale.
- Emerging economies could grow without repeating the high-emission paths of Europe, North America, and parts of Asia.

Fusion would not replace the need for renewables, efficiency, or behaviour change. But it would shift the debate from "how do we manage permanent scarcity?" to "how do we design responsibly around abundance?" – a very different strategic conversation.

7. Closing: what First Light Fusion really represents

First Light Fusion is not interesting because it promises magic. It is interesting because it does three things that are rare in a field dominated by legacy assumptions:

- It starts from engineering and economics, not from monumental physics experiments.
- It deliberately simplifies where others have added complexity.
- It designs for a future in which fusion is not a single facility, but a product with a fleet and a supply chain.

For governments, that means treating fusion as strategic infrastructure and designing permitting, regulation, and public–private models early, not reactively. For investors and capital allocators, it means tracking which fusion concepts are built as products with a clear unit, supply chain, and control points, rather than as open-ended experiments.

For corporates and industrials, it means asking where assured, low-carbon baseload could fundamentally change your economics – and starting to shape partnerships and options around that possibility.

If fusion becomes the backbone of our energy system, it will not be because we discovered new laws of nature. The physics has been known for decades.

Fusion only matters when it becomes a product you can repeat, not an experiment you visit once.

It will be because someone finally asked the right strategic question – and had the discipline to build around the answer.

Acknowledgement

The author would like to thank **Bart Markus**, Chairman of the Supervisory Board of First Light Fusion, for his thoughtful review and technical clarifications. His comments on driver energy, fast-ignition architecture, gain, and realistic cost ranges have significantly sharpened the arguments in this paper. Any remaining errors or interpretations are the author's alone.

About Outdoor Connect

Outdoor Connect is an independent strategy advisory platform focused on board-level value creation for mid-sized, growth-driven companies (€50–€1B). We bring direct senior engagement—without the traditional consulting pyramid—to help founders, CEOs and boards set direction, make sharper capital allocation choices, and embed an execution rhythm. Core areas include growth strategy in technology and the energy transition, strategic repositioning in fragmented markets, and board-level sparring on value creation and M&A preparation.