CUSTODIA WEALTH MANAGEMENT SA
RESEARCH AND ANALYSIS

TECHNOLOGY AND ENERGY: THE DEEP-TECH FRONTIER

NUCLEAR FUSION

From history to commercialization: technology, industrial competition and investment opportunities

Updated as of September 2026

Thematic report — Deep-tech asset class

Document intended exclusively for professional and institutional investors

Table of Contents

1. Executive Summary

For nearly a century, nuclear fusion was above all a scientific problem: proving that it was possible to reproduce on Earth the process that powers the stars. Today the problem is turning into a problem of engineering, industrialization, capital and market. This distinction is fundamental, because it radically changes how a professional investor should read the sector.

> *"Is fusion achievable?"*

The question has become:

> *"Can we build a plant that produces net electricity, continuously, economically and at an acceptable cost of capital?"*

In 2026 a real industrial race for fusion exists. The sector's four key figures are summarized below.

$4.48 bn	$14.24 bn	56	>16,000
Capital raised in the last 12 months (as of Jul. 2026)	Cumulative capital raised since 2021	Companies counted worldwide	Employees working in the sector

Figure 1. Key indicators of the fusion industry, September 2026. Source: Fusion Industry Association, The Global Fusion Industry in 2026.

Capital, moreover, is becoming qualitatively different: it no longer comes only from venture capital and tech funds, but also from pension funds, sovereign wealth funds, industrial companies, energy companies and hyperscalers (operators of large data centers, e.g. AWS, Azure). The real turning point is the shift from fusion science to fusion industry — but that does not mean commercial fusion has been demonstrated.

An update from the Financial Times confirms the trend: in the first eight months of 2026 fusion startups attracted a record $3.8 billion in private capital according to PitchBook, already surpassing the whole of the previous year ($3.3 billion, itself a historic record). A 2026 Fusion Industry Association survey, however, also shows how divided expectations are within the sector itself: 32 of 45 companies surveyed say they could supply commercial energy by 2040 — a significant minority disagrees even among industry insiders (Financial Times, September 2026).

As of September 2026 several fundamental problems remain unresolved:

- plasma confinement and stability
- net energy at the whole-plant level
- materials subjected to neutron bombardment
- tritium breeding
- fuel cycle
- component maintenance
- operating lifetime
- machine availability
- plant cost

- licensing and grid connection
- final cost of electricity

From a financial standpoint, fusion should therefore still be regarded as an extremely high-risk deep-tech asset class, but with an extraordinarily high potential payoff. For a professional European investor there is not yet a convincing reason to treat it as a traditional energy investment: there is instead a thesis under which fusion represents a long-duration technology option.

The most interesting investment might also not necessarily be the company that builds the first reactor, but rather the suppliers of technologies that are hard to replace: HTS magnets, advanced materials, cryogenic systems, power electronics, gyrotrons, vacuum technology, fuel-cycle systems and high-temperature components.

We will explain each of these points and concepts in more detail in the report.

2. From the Stars to the Earth: the Scientific Fundamentals

In 1926 Arthur Eddington understood that the enormous energy of stars could be explained by the transformation of hydrogen into heavier elements: fusion thus became one of the great questions of 20th-century physics. The principle is relatively simple: two light nuclei (deuterium) are brought together until they overcome electrostatic repulsion; when they enter the regime of the **strong nuclear force**, they can fuse, producing a heavier nucleus and releasing energy.

Deep-Dive Box: the strong nuclear force

The strong nuclear force (or simply "strong interaction") is one of the four fundamental forces of nature, together with gravity, electromagnetism and the weak nuclear force. It is the force that holds protons and neutrons together inside the atomic nucleus. It is needed because otherwise protons, all carrying positive charge, would repel one another through electrostatic repulsion (like charges repel) and the nucleus could not exist.

Why is it "strong"? It is the most intense of the four fundamental forces — far stronger than the electromagnetic force — but it has an extremely short range: it acts only at distances on the order of the size of the atomic nucleus (femtometers, i.e. millionths of a billionth of a metre). Beyond that distance it practically vanishes. This very short range is precisely what explains the fundamental difficulty of nuclear fusion: two light nuclei (such as deuterium and tritium) repel each other through electrostatic repulsion as long as they are at "normal" distance; only if they are brought close enough to enter the range of the strong nuclear force does it overcome the electrostatic repulsion and let them fuse, releasing energy. Bringing the nuclei to that minimum distance requires overcoming the electrostatic repulsion — and this requires enormous temperatures (hundreds of millions of degrees): only at those temperatures do the nuclei move fast enough to "collide" closely enough for the "strong" force to come into play.

The **physical state** needed to create the conditions (velocity, density, temperature) that let nuclei get close enough to "feel" the strong nuclear force is called **plasma**. Technically, plasma is the fourth state of matter, beyond solid, liquid and gas. As a gas is heated further, at some point the energy becomes so high that electrons are "torn" from the atoms. The result is a set of charged particles — positive ions (nuclei stripped of their electrons) and free electrons — moving chaotically but remaining, overall, electrically neutral.

Being made of charged particles, plasma responds to electric and magnetic fields in a way that a normal gas does not. This is crucial for fusion: precisely because it is made of electric charges, plasma can be "manipulated" and contained by magnetic fields, something impossible with a neutral gas. Plasma is, ultimately, the state in which the fusion fuel (deuterium and tritium) exists at the hundreds-of-millions-of-degrees temperatures the reaction requires. At those temperatures no solid material could physically contain it (it would melt or evaporate instantly); it must therefore be confined "without touching it", through powerful magnetic fields (the tokamak approach used by CFS and STEP) or by compressing it very rapidly before it disperses (inertial confinement, as at NIF, or hybrid approaches such as those of Helion and General Fusion).

The practical problem is that plasma is **intrinsically unstable**: it tends to develop turbulence, oscillations and instabilities that make it "escape" confinement. This is exactly why every scientific advance so far tested has revealed — for decades — a new obstacle: keeping a plasma hot, dense and stable long enough to produce net energy is one of the most difficult engineering problems ever tackled.

In short: the strong nuclear force is the "glue" that makes fusion possible once the electrostatic repulsion barrier has been overcome — and it is precisely why (literally) stellar temperatures are needed to make it happen on Earth. Extreme temperatures allow matter to become plasma; the nuclei in the plasma move so fast that, upon colliding, they manage to overcome electrostatic repulsion and enter the range of the strong nuclear force; nuclear fusion is triggered, releasing an enormous amount of energy.

Scheme 1. The strong nuclear force.

The reaction most used in current projects is:

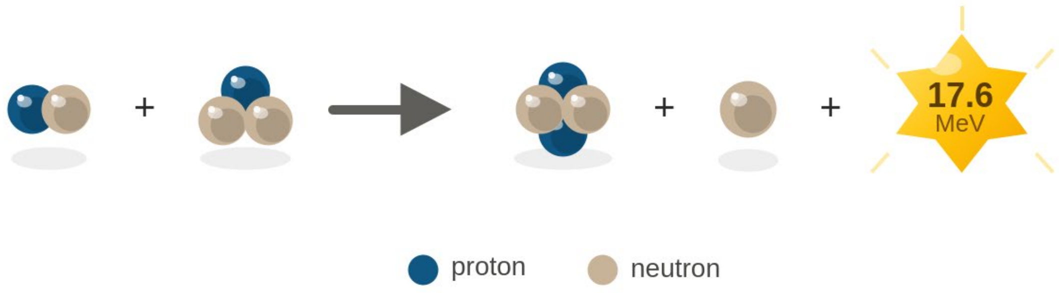

$$D + T \Longrightarrow He4 + n + 17.6MeV$$

Equation 1. Explanatory formulation of industrial fusion.

where

D = deuterium, a hydrogen isotope with one proton and one neutron in the nucleus (ordinary hydrogen has only the proton);

T = tritium, another hydrogen isotope, with one proton and two neutrons. It is radioactive and unstable;

$He4$ = a helium-4 nucleus (two protons and two neutrons) — the “heavy” nucleus formed by the fusion;

n = a free neutron, ejected from the reaction;

17,6 MeV = the energy released by the reaction (MeV = million electronvolts, a unit of energy used in nuclear physics).

When a nucleus of deuterium and one of tritium collide with sufficient energy (thanks to the plasma's high temperature, as explained earlier), they overcome electrostatic repulsion and enter the range of the strong nuclear force. At that point they fuse into a heavier nucleus (helium-4), but the new nucleus is slightly lighter than the sum of the two starting nuclei. That “missing mass” does not disappear: it turns into energy (following Einstein's famous equation $E = m \cdot c^2$). This energy is split between the ejected neutron (which carries off about 80% of the energy, roughly 14.1 MeV) and the helium nucleus (about 20%, 3.5 MeV).

Among all possible fusion reactions (including the one that constantly occurs at the stellar level), the one shown in Equation 1 occurs most easily at the “lowest” temperatures achievable with current technology — which does not mean low in absolute terms (we are still talking about hundreds of millions of degrees), but it is the most practicable trade-off today. The price to pay is that it requires tritium.

Deuterium is a relatively abundant element and can be extracted from seawater; tritium, by contrast, is rare, radioactive and not available in sufficient quantities to fuel a future global fusion industry. This apparent contradiction contains one of the technology's central problems: a commercial fusion plant will probably have to produce internally most of its own tritium through a lithium-containing blanket.

The blanket (literally "blanket" or "mantle") is a structure that surrounds the plasma inside the reactor, positioned just behind the first wall (the one closest to the plasma). It contains lithium and is designed to do two things at once: produce tritium and recover heat.

When deuterium-tritium fusion occurs, a high-energy neutron is released (the n in Equation 1). This neutron leaves the plasma and passes through the blanket, where it strikes the lithium atoms it contains. Lithium has two natural isotopes that interact with the neutron differently:

- Lithium-6 + neutron → Tritium + Helium-4 (energy-releasing reaction)
- Lithium-7 + neutron → Tritium + Helium-4 + another neutron (energy-absorbing reaction, but one that "returns" a neutron that can trigger further tritium-breeding reactions)

In practice, the neutron the plasma produces is "recycled" to generate exactly the fuel (tritium) the plasma itself consumes. It is not enough for the reaction to be theoretically possible: the reactor must produce more tritium than it consumes — that is, the ***tritium breeding ratio*** (TBR) must be greater than 1, to compensate for:

- the natural radioactive decay of tritium (it has a half-life of about 12 years, so part of it "vanishes" over time regardless);
- losses during handling and storage, and the inventory needed for continuous operation;
- losses in separation and purification processes

Tritium breeding, however, does not need to be demonstrated in a small-scale laboratory, but in a real industrial plant running uninterruptedly. This is why we maintain that fusion is not simply a matter of "putting two atoms together": it is an entire industrial system.

Why Fusion Has Been "Always Twenty Years Away" for Decades

The fundamental difficulty is temperature: to achieve a sufficiently high probability of fusion, the plasma must reach temperatures on the order of hundreds of millions of degrees, at which matter cannot be contained by a physical wall. The plasma must therefore be confined through magnetic fields (magnetic confinement) or compressed very rapidly (inertial confinement).

The tokamak, one example of a magnetic-confinement plant (see Figure), traps the plasma inside a toroidal chamber using powerful magnetic fields — but the plasma is extremely unstable, and every experimental improvement tends to reveal new problems. This is why the joke that "fusion is always twenty years away" has held up for so long: every time science got closer to the result, another piece was missing — first temperature, then confinement, then confinement time, then *plasma burning*, then materials, then tritium, then the thermodynamic cycle, then economics.

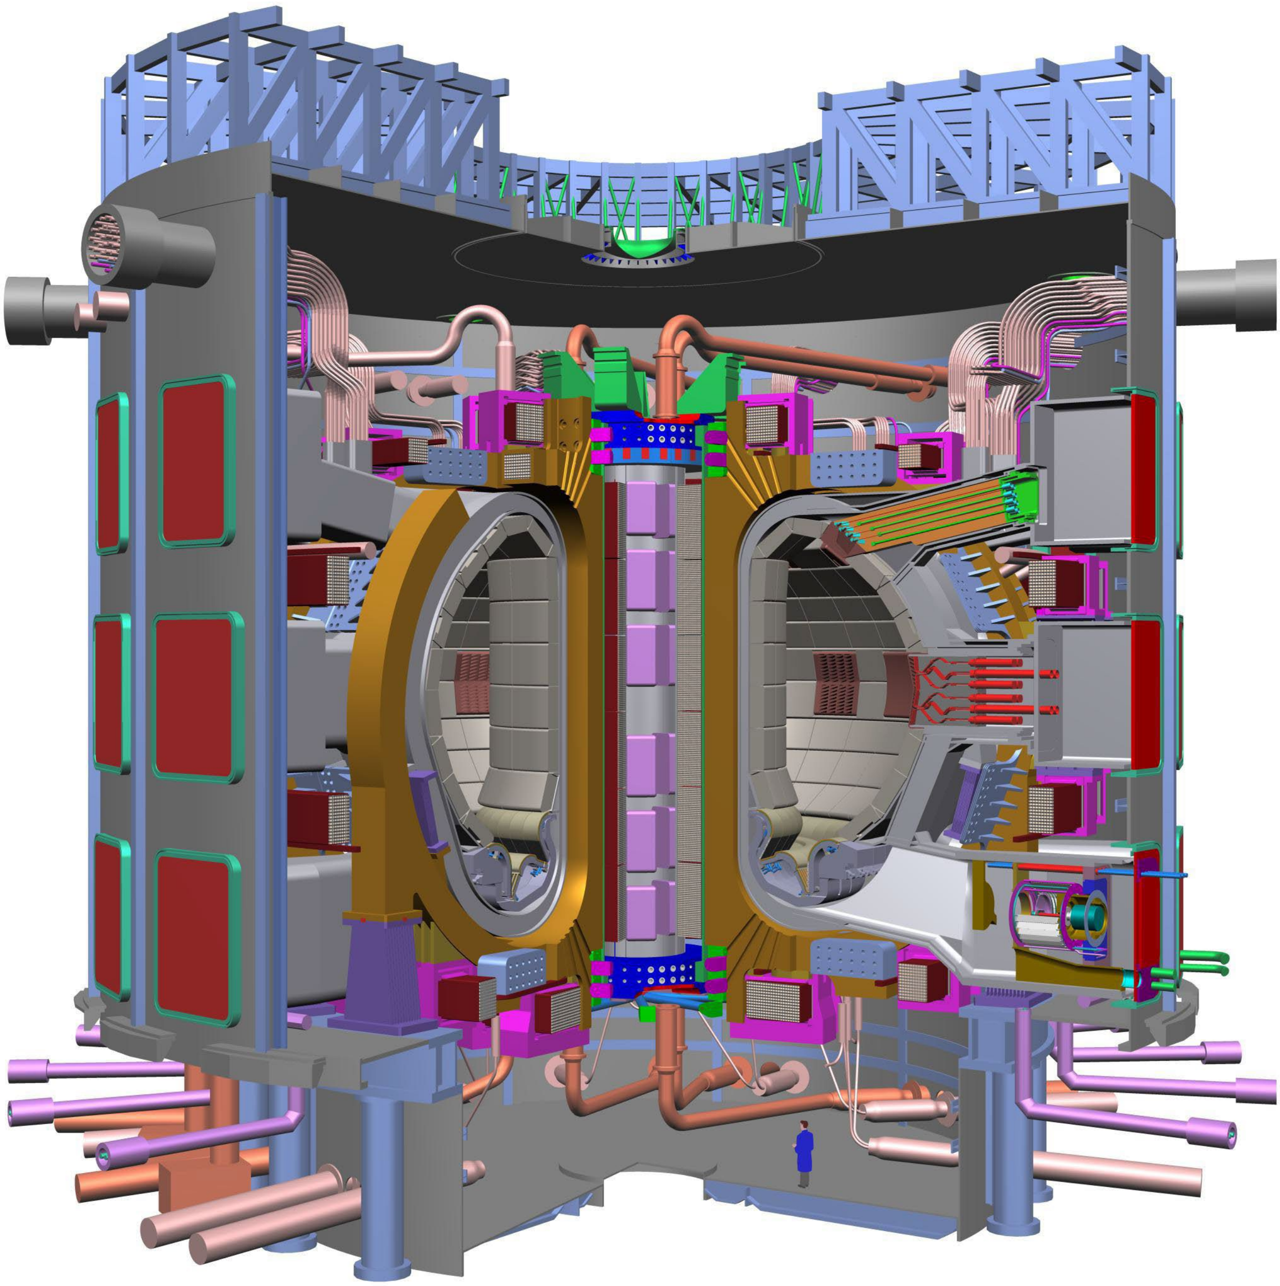

Figure 2. Cross-section diagram of the International Thermonuclear Experimental Reactor (ITER), the world's largest tokamak, whose construction began in 2013 and which is expected to enter its initial operating phase in 2034/2035. The project aims to demonstrate the feasibility of a working fusion reactor and will produce 500 megawatts of power. The blue human figure at the bottom is included for scale.

Deep-Dive Box: plasma burning

Plasma burning (or "burning plasma") is a specific and highly sought-after condition in fusion physics: it is the state in which the plasma becomes energetically self-sustaining.

To start fusion, the plasma must be heated from outside — with gyrotrons, with particle-beam heating, and so on. At the start, therefore, the energy that keeps the plasma hot is supplied by external systems.

But the fusion reaction itself produces energy. The helium-4 nucleus produced (which stays “trapped” in the plasma, unlike the neutron, which escapes) carries with it part of that energy (about 3.5 MeV) as kinetic energy, which further heats the surrounding plasma.

When the amount of heat generated by the helium particles produced by fusion becomes enough to keep the plasma at the required temperature — with little or no further need for external heating — the plasma is said to be “burning”: it self-sustains, somewhat like a flame that, once lit, keeps itself going as long as there is fuel.

Achieving a plasma that “burns” on its own is a fundamental intermediate step between simply “we produced fusion for an instant” and “we have a plant that runs continuously”. It is also linked to the concept of Q discussed further on: burning plasma is typically associated with very high Q values (ideally tending to infinity, so-called “ignition”, when external heating is no longer needed at all).

But — as noted further on regarding ITER — even achieving burning plasma is still not equivalent to a commercial reactor: one must still account for all the energy consumed by magnets, cryogenics, pumps and auxiliary systems (i.e. the so-called “***engineering Q***”), not just the plasma's own energy balance.

Scheme 2. Plasma burning.

The history of fusion is therefore not a succession of failures, but a succession of problems being moved from a scientific level to a higher engineering level. And it is precisely this that is changing today.

3. Two Industrial Paradigms Compared

ITER: the traditional scientific paradigm

ITER (*International Thermonuclear Experimental Reactor*) is a huge international project built in southern France and funded by a collaboration among major industrial powers. Its purpose is not to sell electricity: it is not a power plant, but an experimental machine meant to demonstrate that a fusion plasma can reach sufficiently advanced burning-plasma conditions, with a historic goal of $Q \approx 10$ (see Figure).

In simplified terms, Q = energy produced by fusion / energy supplied to the plasma. A plasma Q above 1 does not mean, however, that a plant produces more electricity than it consumes: magnets, pumps, cryogenics, heating, energy conversion (power electronics), control, vacuum systems, fuel cycle and other auxiliary loads must all be taken into account.

Deep-Dive Box: *picks and shovels*

“Picks and shovels” is an English idiom that metaphorically refers to all the integrated, supporting components needed to develop a project. It comes from the 19th-century American gold rush. Many gold prospectors went broke, but those who sold them tools — picks, shovels, sieves, boots — made money reliably, regardless of which prospector struck gold. The tool merchant did not have to bet on which prospector would succeed: he earned money regardless, because everyone needed his tools.

The main “picks and shovels” for this particular sector are:

- **Power electronics**: this is the branch of electronics dealing with the conversion and control of electrical energy — static transformers, converters, inverters, systems that convert direct current into alternating current (or vice versa), that regulate voltage and frequency, and so on. In the context of nuclear fusion, it is one of the critical technological components for fusion plants: it manages and converts the energy produced (or the energy needed to power magnets, plasma-heating systems, etc.).
- **Gyrotron**: the gyrotron is a device that generates extremely high-power, high-frequency microwaves, used to heat the plasma in magnetic-confinement reactors (tokamaks). Without this component it is not possible to bring the plasma to the hundreds-of-millions-of-degrees temperatures fusion requires.
- **Vacuum technology**: vacuum technology (or “vacuum technologies”) covers the systems that create and maintain the vacuum inside the chamber where fusion takes place — vacuum pumps, sealed chambers, ultra-high-vacuum components. It is essential because the plasma must be isolated from any atmospheric contamination.
- **Fuel-cycle systems**: the fuel cycle refers to the entire system that manages the reaction's fuel (deuterium and tritium) — extraction, production (tritium breeding via lithium), storage, purification and reinjection into the reactor.

Scheme 3. The components needed to build nuclear fusion reactors.

ITER has also suffered significant delays: the new baseline envisages an initial research phase in the 2030s and the start of deuterium-tritium operations in 2039. This does not mean ITER is a failure — it means it is demonstrating just how difficult it is to turn fusion physics into an industrial machine. And it is precisely this difficulty that has created room for the second paradigm.

The Private Startup Paradigm

The startup fusion model is radically different: instead of building a giant machine to scientifically demonstrate the principle, these companies try to simultaneously optimize size, cost, magnets, materials, lifetime, maintenance, energy conversion, industrial production and construction times.

The question is no longer “how can we build the largest possible experiment”, but “how can we build a plant that can be replicated”. This difference is fundamental from an investor's point of view: a scientifically superior technology is not necessarily the economically winning one. The winner will likely be whoever manages to combine plasma physics, materials science, industrial manufacturing, power electronics, fuel cycle and project finance.

It is important to correct a simplistic reading: the main players are not playing the same game. ITER is above all an international scientific program; Helion, CFS and General Fusion are commercial companies; STEP is a national industrial program. These projects occupy different points on the technology-maturity curve. The real competition concerns which technological architecture will first become economically replicable — not necessarily who produces the first watt. The next section will go into a bit more detail on each of these projects.

4. The Industrial Protagonists

Helion Energy — the boldest bet

Helion uses a configuration called Field-Reversed Configuration (FRC): plasma is produced in compact configurations and then compressed through magnetic fields. Unlike the classic path fusion → heat → steam → turbine → electricity, Helion is aiming for direct conversion of plasma energy into electricity via electromagnetic induction — an economically very important feature, since it could eliminate much of the complexity of a conventional thermal power plant.

The seventh prototype, Polaris, was designed specifically to demonstrate the conversion of fusion energy into electricity; the system can use D-D, D-T and D-He-3 and incorporates direct inductive energy recovery. The company is also building Orion in the state of Washington. In 2023 Helion signed a PPA with Microsoft for 50 MW of fusion electricity, with operations expected to start in 2028 — a feature that makes Helion particularly interesting financially: it is not simply selling a technology, it is trying to sell electricity.

A governance factor worth monitoring: according to the Financial Times, Helion is “notoriously secretive” and reluctant to publish peer-reviewed results, unlike General Fusion and TAE, which are more open about sharing their data (both good and bad). CEO David Kirtley justifies the choice as normal protection of intellectual property, arguing that publishing scientific papers “takes time and resources” away from the team. It is a specific opacity factor to bear in mind when assessing the company's risk profile relative to its competitors.

An indirect but independent confirmation comes, somewhat paradoxically, from a competitor: Guo Houyang, founder of China's NovaFusionX (which uses the same FRC configuration), stated that Helion's February 2026 announcement — fusion achieved with plasma at 150 million degrees Celsius — “fully confirmed the scientific soundness, feasibility and commercialisation potential” of this specific technological approach (Financial Times, September 2026). It is a rare case of cross-validation between two competing companies in different markets.

Deep-Dive Box: the types of fusion that can be implemented

These can be obtained by combining hydrogen isotopes with each other or with helium-3. There are essentially two types of helium:

- He4 (2 protons + 2 neutrons) = the clean, stable “waste product” of almost all fusion reactions (non-radioactive, non-hazardous)
- He3 (2 protons + 1 neutron) = a rare, valuable reactant, used in the **D-He3** reaction precisely because it produces very few neutrons — an enormous advantage for the reactor's materials, but its scarcity is why this path remains more theoretical than industrially practicable in the short term

The difference lies in the number of neutrons in the nucleus. Both are isotopes of helium (so they always have 2 protons, which chemically makes them helium), but their nuclear composition differs. In every reaction described below, the end product is always He4 — the “ultimate stable” nucleus, the one toward which nearly all fusion reactions in the universe tend (also why stars burn mainly by producing helium-4, as Eddington wrote in 1926). He3 is rare on Earth because it forms mainly through tritium decay (T → He3 + electron, with a half-life of about 12 years) or is produced in nuclear reactions — it does not exist in significant natural quantities on our planet. It is relatively more abundant on celestial bodies such as the Moon, where it has accumulated over time thanks to the solar wind (the periodic discussions about “lunar mining of helium-3” for fusion energy are more speculative than practical).

With that established, here is a brief description of the three practicable (or nearly practicable) types of fusion:

1) D-T (deuterium-tritium)

This is Equation 1.

Advantages: it is the “easiest” reaction to achieve, because it requires the lowest temperature relative to the others (still enormous, though — hundreds of millions of degrees). This is why it is the choice of almost all current projects, including ITER and CFS's SPARC/ARC.

Disadvantages: tritium is rare and radioactive — hence the breeding problem via the lithium blanket discussed earlier. It also releases a very energetic neutron, which damages the reactor's materials over time (the materials problem is one of the main obstacles).

2) D-D (deuterium-deuterium)

Two deuterium nuclei fuse together. This reaction actually has two possible outcomes:

$$D + D \rightarrow T + p$$

$$D + D \rightarrow He3 + n$$

Where p = proton and n = neutron.

Advantages: deuterium is abundant (it can be extracted from seawater) and there is no need to produce tritium externally.

Disadvantage: requires higher temperatures than **D-T** and is less efficient — it produces less energy per reaction. It still produces neutrons.

We have encountered two types of helium:

- He4 = the clean, stable “waste product” of almost all fusion reactions (non-radioactive, non-hazardous)
- He3 = a rare, valuable reactant, used in the **D-He3** reaction precisely because it produces very few neutrons — an enormous advantage for the reactor's materials, but its scarcity is why this path remains more theoretical than industrially practicable in the short term

The difference lies in the number of neutrons in the nucleus. Both are isotopes of helium (so they always have 2 protons, which chemically makes them helium), but their nuclear composition differs.

3) D-He3 (deuterium-helium-3)

It produces the following sequence:

$$D + He3 \rightarrow He4 + p + 18{,}3\ MeV$$

Advantages: almost no neutrons produced (unlike **D-T** and **D-D**). This is very interesting because neutrons are the main cause of material damage and radioactive activation of the reactor's structures.

Disadvantages: requires much higher temperatures than **D-T** (it is harder to ignite), and above all helium-3 is extremely rare on Earth — it is more easily found on the Moon or other celestial bodies, so it is not conceivable for a large-scale industry in the short-to-medium term.

Helion's seventh prototype, Polaris, "can use **D-D**, **D-T** and **D-He3**": this means the machine is designed to be flexible across different fuels. This is an interesting feature because, in the short term, it can run on **D-T** or **D-D** (more practicable today), while in the more distant future it could aim for **D-He3**, which produces few neutrons and therefore simpler materials, lower radioactive activation and less component wear.

It is the same principle behind TAE Technologies' long-term bet on proton-boron fusion (p-B11): this too produces very few neutrons, but requires even more extreme conditions than **D-He3**.

Scheme 4. Types of nuclear fusion compared.

The most delicate point, however, remains valid in 2026: reaching fusion conditions is not the same as producing commercial electricity. The sequence that must be demonstrated is as follows: hot plasma → confinement → fusion → energy gain → energy recovery → repeated operation → component lifetime → reliability → plant → grid connection — at every step there is a potential point of failure. 2028 is therefore an extremely aggressive deadline: if met, Helion's competitive advantage would be enormous; if it slips to 2030-2032 or beyond, that would not mean technological failure, but it would require a drastic downward revision of the valuation. Risk, as is typical in deep-tech, is not binary: there is a distribution of intermediate outcomes.

Commonwealth Fusion Systems (CFS) — the most credible strategy

CFS comes out of the MIT world and uses a tokamak, differing from ITER mainly through its use of high-temperature superconducting (HTS) magnets based on REBCO. A higher magnetic field allows a more tightly confined plasma, a more compact machine, a smaller volume and a potential cost reduction.

Deep-Dive Box: rare earths — what do they have to do with fusion?

We discussed rare earths and strategic minerals in our Deep-Dive of July 20, 2026. Their importance lies in the production of permanent magnets. As noted, a higher magnetic field keeps the plasma more tightly confined, which in turn requires machines with smaller volumes and therefore lower costs.

HTS/REBCO magnets make it possible to generate much more intense magnetic fields than the traditional superconducting magnets used in ITER, while keeping the machine's dimensions far smaller. This is exactly why CFS can aim to build ARC (its commercial reactor) much smaller than ITER, while still targeting comparable or superior performance.

The **HTS** (*High Temperature Superconductor*) high-temperature superconductors are materials that, when cooled below a certain critical temperature, conduct electricity with zero electrical resistance — no energy loss, extremely high currents without overheating. The term "high temperature" is a bit misleading: it does not mean "hot" in absolute terms, but "less cold" than traditional superconductors. Classic (low-temperature, LTS) superconductors must be cooled close to absolute zero with liquid helium (about -269°C). HTS materials, by contrast, become superconducting at much higher temperatures — still frigid (typically below about -180°C), but reachable with liquid nitrogen, which is far cheaper and simpler to handle than liquid helium.

The **REBCO** (*Rare-Earth Barium Copper Oxide*) — an oxide of copper, barium and a rare-earth element (often yttrium or gadolinium) — is the specific material used today to build the best commercially available HTS superconducting tapes.

In practice, REBCO is a type of HTS — HTS being the general category of superconductors; REBCO is the specific material used to make them in the form of thin tapes that can be wound into magnetic coils. One of the most qualified producers of this type of magnet is Fujikura (see Section 6).

Scheme 5. The rare earths needed for the fusion process.

The SPARC project represents the experimental stage, ARC the commercial project. In June 2026 CFS published five peer-reviewed papers on ARC's physics: according to the company's model, ARC should produce about 1.1 GW of fusion power and about 400 MW of continuous net electricity — an important distinction, because CFS is not talking only about energy produced in the plasma, but about net energy delivered to the grid. ARC is expected to be operational in the early 2030s.

In 2026 CFS raised a further $1 billion, bringing total capital raised to about $4 billion: a signal that the capital markets are starting to finance not just research, but the industrialization phase.

The Financial Times provides a concrete detail about this round: about $1 billion came from investors including, not explicitly named, pension funds and sovereign wealth funds — the first specific public confirmation that this kind of “patient” capital is actually entering the sector, not just in theory.

Eni — probably the best indirect way to invest in fusion

Eni has been a strategic shareholder of CFS since 2018 and in September 2025 signed a PPA worth more than $1 billion to purchase electricity from the ARC plant. This changes the nature of the relationship: Eni is no longer just a financial investor, but is trying to become part of fusion's future industrial supply chain, bringing skills a plasma-physics startup does not necessarily possess — project development, large-plant management, procurement, safety, supply chain, fluid handling, fuel cycle, hydrogen technology, energy-infrastructure management and project finance.

In particular, Eni has certainly developed the skills to provide the services listed below.

- ***Procurement***. This is the business function responsible for purchasing all the goods, services, materials and equipment an organization needs to operate — from supplier selection, to contract negotiation, to managing quality and delivery times. For a complex industrial project (such as building a power plant), procurement means managing hundreds or thousands of different suppliers: who supplies the steel, who the electronic components, who the safety systems, and making sure everything arrives on time, at the required quality and at the agreed price.

- ***Project finance***. This is a specific financing technique, different from the “normal” way a company finances itself. Instead of borrowing against its own general resources (its overall balance sheet), a dedicated special-purpose vehicle is created (an SPV, “*) that exists solely to build and operate that single project. Financing is secured primarily by the future cash flows the project itself will generate (for example, revenue from electricity sales), not by the parent company's general assets. It is the technique typically used for major infrastructure: power plants, highways, pipelines,*) that exists solely to build and operate that single project. Financing is secured primarily by the future cash flows the project itself will generate (for example, revenue from electricity sales), not by the parent company's general assets. It is the technique typically used for major infrastructure: power plants, highways, pipelines, offshore wind farms — very costly, long-term projects requiring enormous capital repaid over decades.

- ***Hydrogen technology***. This is the set of technologies related to producing, storing, transporting and using hydrogen as an energy carrier — that is, as a means of storing and transporting energy, not as a primary energy source in itself. Its main components are:
    - Production: mainly through electrolysis (water is split into hydrogen and oxygen using electricity) or through natural-gas reforming. If the electricity used comes from renewable sources, it is called "green hydrogen" — a central theme in the energy transition.
    - Storage: hydrogen is a very light, low-density gas, so it must be compressed at very high pressure, liquefied at extremely low temperatures (about -253°C), or chemically bonded to other compounds — all technically complex solutions.
    - Transport: via dedicated pipelines, cryogenic tanker trucks or specialized ships.
    - Use: as fuel for fuel cells (which convert it back into electricity), for modified combustion engines, or as an industrial feedstock (for example in the production of "green" steel or fertilizers).

    The real commercial problem of fusion may in fact no longer be the physics, but the fuel cycle — precisely one of the areas where Eni is building expertise. For a European investor, Eni can therefore be read as a kind of real option on fusion: if fusion does not work out, Eni remains a large, diversified energy company; if it becomes industrially viable, Eni could hold expertise and strategic stakes in a potentially transformative technology. This asymmetry is very interesting.

    - hydrogen (and its isotopes, deuterium and tritium) are small, elusive molecules, hard to contain without leaks;
    - they require compression systems, cryogenics and special materials to be handled safely;
    - raise "embrittlement" issues in metals (hydrogen can infiltrate structural materials and make them brittle over time — a phenomenon called "*hydrogen embrittlement*")

    An energy company like Eni, which has already developed expertise in handling hydrogen for other industrial uses (refining, future green-hydrogen projects), therefore already possesses much of the technical know-how needed to handle deuterium and tritium in a fusion plant's fuel cycle — even though the applications differ, the underlying engineering problems (containment of light gases, cryogenics, materials) are quite similar.

The real commercial problem of fusion may in fact no longer be the physics, but the fuel cycle — precisely one of the areas where Eni is building expertise. For a European investor, Eni can therefore be read as a kind of real option on fusion: if fusion does not work out, Eni remains a large, diversified energy company; if it becomes industrially viable, Eni could hold expertise and strategic stakes in a potentially transformative technology. This asymmetry is very interesting.

General Fusion — the third way, now listed

General Fusion uses Magnetized Target Fusion: plasma is mechanically compressed through a system based on a liquid-lithium wall and pistons. The theoretical advantage is the ability to use relatively conventional industrial components; the disadvantage is the complexity of mechanical synchronization and the difficulty of simultaneously achieving sufficient temperature, density and confinement.

In June 2026 the company announced progress with LM26, reaching about 8.4 million °C of electron temperature — an interesting result, but still far from the conditions needed for a D-T plant. “LM” stands for “Lithium wall Machine” (referring to the liquid-lithium wall used in the design), while “26” probably indicates the generation or version of the prototype in the company's development sequence. The technology is called *Magnetized Target Fusion* (MTF): the plasma is mechanically compressed through a system based on a liquid-lithium wall and pistons. It is a hybrid approach, different from both pure magnetic confinement (tokamak) and pure inertial confinement (as at NIF — see Section 5):

- the plasma is first “magnetized” (partly confined with magnetic fields, to stabilize it);
- it is then rapidly and mechanically compressed, through pistons that compress the liquid-lithium wall around the plasma;
- the compression should bring the plasma to the temperature and density conditions needed for fusion

The theoretical advantage is the ability to use relatively conventional industrial components (pistons, mechanical systems) instead of extremely costly and complex superconducting magnets. The disadvantage is the difficulty of mechanically synchronizing everything and simultaneously achieving sufficient temperature, density and confinement.

The real financial turning point came in July 2026: General Fusion became the first pure-play fusion company listed on the Nasdaq (ticker GFUZ), entering the market with about $150 million in cash earmarked to fund the LM26 program through to the 2028 milestones. This makes GFUZ very interesting, but also extremely speculative.

A useful precedent for calibrating expectations: in 2014 founder Michel Laberge stated the goal of selling a commercial product by 2022; in January 2026 CEO Greg Twinney told investors that 2035 was the new horizon for the commercial plant — a slip of more than a decade relative to the original promise (Financial Times, September 2026).

TAE Technologies — the proton-boron bet

TAE uses a Field-Reversed Configuration and is aiming in the long term for proton-boron fusion, which produces few neutrons compared with D-T but requires far more difficult conditions to reach. In 2025 TAE announced a merger agreement with Trump Media & Technology Group for a stated value of more than $6 billion; as of September 2026 the deal has not yet closed, with completion expected by year end.

For investors, this means the DJT ticker cannot be considered a pure fusion proxy: it is a hybrid investment combining media, technology, crypto/fintech, fusion and the deal's corporate structure — with a specific risk enormously higher than that of a direct investment in an industrial company.

TAE shows a similar pattern: in 2010 the company pointed to 2020 as its target for commercial energy; today the company's website speaks of the “early 2030s” (Financial Times, September 2026). These two cases — together with ITER's delay toward 2039 — are the best empirical evidence available in support of Scenario 2 (“technical success, economic delay”) described further on: in the fusion sector, slipping deadlines are the historical norm, not the exception.

United Kingdom and STEP — the national industrial path

The STEP project (Spherical Tokamak for Energy Production) will be built at West Burton, on the site of a former coal-fired power station. The United Kingdom does not simply want to build a reactor, but to create a national fusion industrial supply chain, with start-up expected around 2040. In 2026 the British government awarded major industrial contracts and raised funding for the next phase of STEP to about £1.3 billion.

It is a much slower project than Helion, but this very feature is its advantage: STEP is less exposed to the risk of a single startup and more oriented toward building an industrial ecosystem. For the investor, however, direct exposure is less straightforward: the economic value of the British strategy may emerge above all through supply-chain suppliers.

Other Players: Pacific Fusion and Inertia Enterprises

Two less-established names, but ones that emerged strongly in 2026, deserve a mention. Pacific Fusion is a startup founded and funded by high-profile scientists and tech investors (including Eric Schmidt, former Google CEO); in August 2026 it began construction of an experimental site in Albuquerque, New Mexico, with the declared goal of representing a crucial step toward commercial fusion electricity.

Inertia Enterprises, by contrast, aims to bring to market the inertial-confinement technology developed at the Lawrence Livermore National Laboratory (the same one behind NIF), representing a possible commercial path for an approach that has so far remained within public research. It raised $450 million in its first major funding round in February 2026. CEO Jeff Lawson argues that, thanks to decades of publicly funded research and $30 billion of US government investment, the company does not need to “solve a physics problem”, only engineer a technology that has already been demonstrated. This is a claim to take with caution — the same kind of declared optimism that, as seen above, has already led to large timeline slips at other companies in the sector.

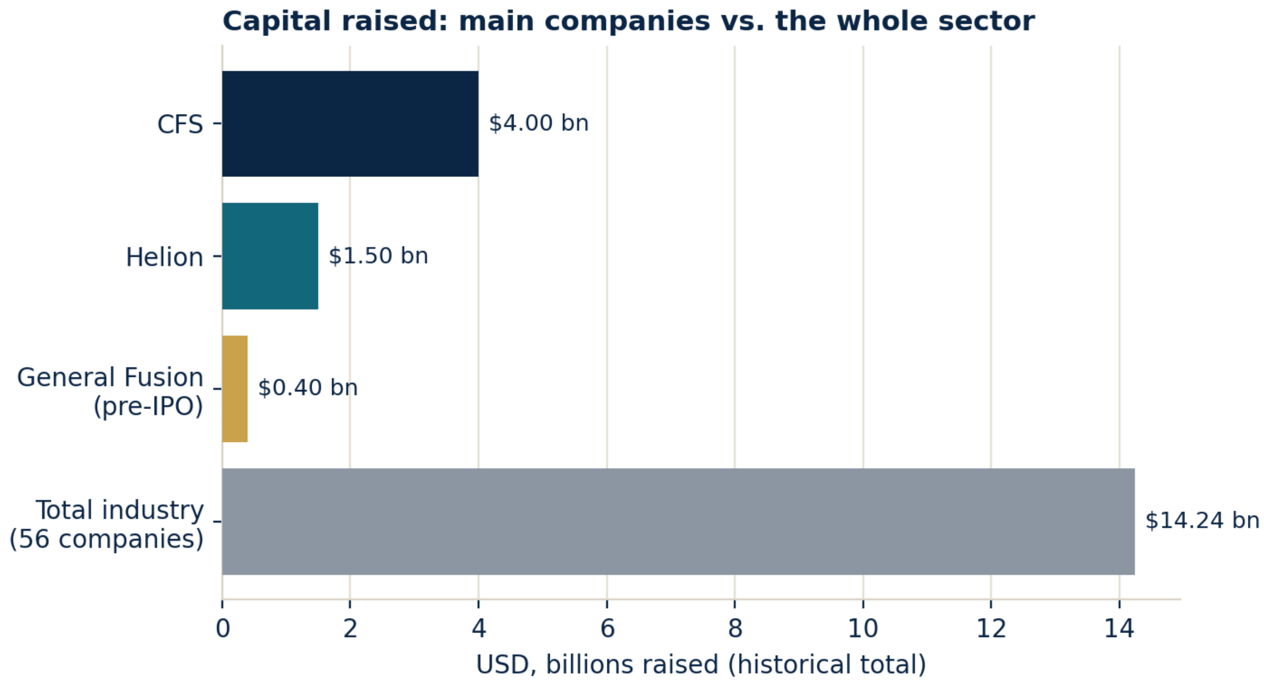

Figure 3. Capital raised: main companies vs. the whole sector (historical total, USD billions). Source: Fusion Industry Association; company press releases.

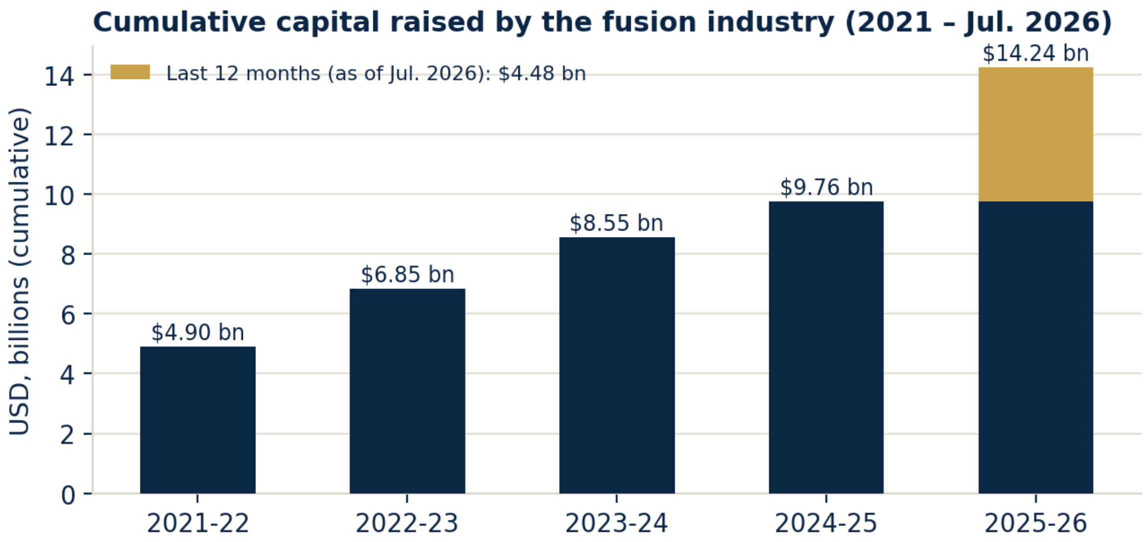

Source: Fusion Industry Association, The Global Fusion Industry in 2026. Illustrative annual trajectory built from the two published figures (total cumulative and last 12 months).

Figure 4. Cumulative capital raised by the fusion industry, 2021 – July 2026. Illustrative annual trajectory built from the two figures published by the FIA (total cumulative and last 12 months); does not represent official annual data.

5. The Real Obstacle: From Plasma to Net Electricity

National Ignition Facility (NIF): a historic result, not a power plant

The National Ignition Facility (NIF) is the largest and most powerful laser facility in the world, located at Lawrence Livermore National Laboratory in California (United States). It uses an approach different from the tokamak: it is called inertial confinement. Instead of confining the plasma with magnetic fields for a prolonged time (as ITER, CFS or STEP do), NIF does the opposite:

- it focuses 192 extremely high-power laser beams onto a tiny target containing deuterium and tritium, about the size of a peppercorn;
- the lasers strike the target simultaneously from all directions, compressing and heating it in an extremely brief fraction of a second
- the compression is so rapid and intense that the fuel reaches fusion conditions before it can disperse — hence the term "inertial": it exploits the inertia of the compressed matter, not prolonged confinement over time

In December 2022 the National Ignition Facility produced 3.15 MJ of fusion energy against 2.05 MJ of laser energy delivered to the target: a historic result[1], but not a power plant. The result demonstrates a form of breakeven at the target level, not a positive energy balance for the whole plant — a fundamental distinction for investment analysis.

The Financial Times later provided the overall figure: the 192-laser system required about 422 MJ of electricity to operate — more than 100 times the fusion energy produced (3.15 MJ). This is the number that makes the distinction between the various types of Q, described further on, more tangible: a scientific milestone at the target level, but with a wall-plug Q still deeply negative (Financial Times, September 2026).

This is the crucial point of this experiment (shown in Figure 5): that result is a positive energy balance only at the level of the target itself (target gain), not of the whole plant. If one considers the total energy consumed by the whole system — including the electrical energy needed to generate those 192 laser beams, which is enormously greater than the 2.05 MJ delivered to the target — the plant's overall energy balance remains deeply negative.

[1] A scientific milestone often called "ignition".

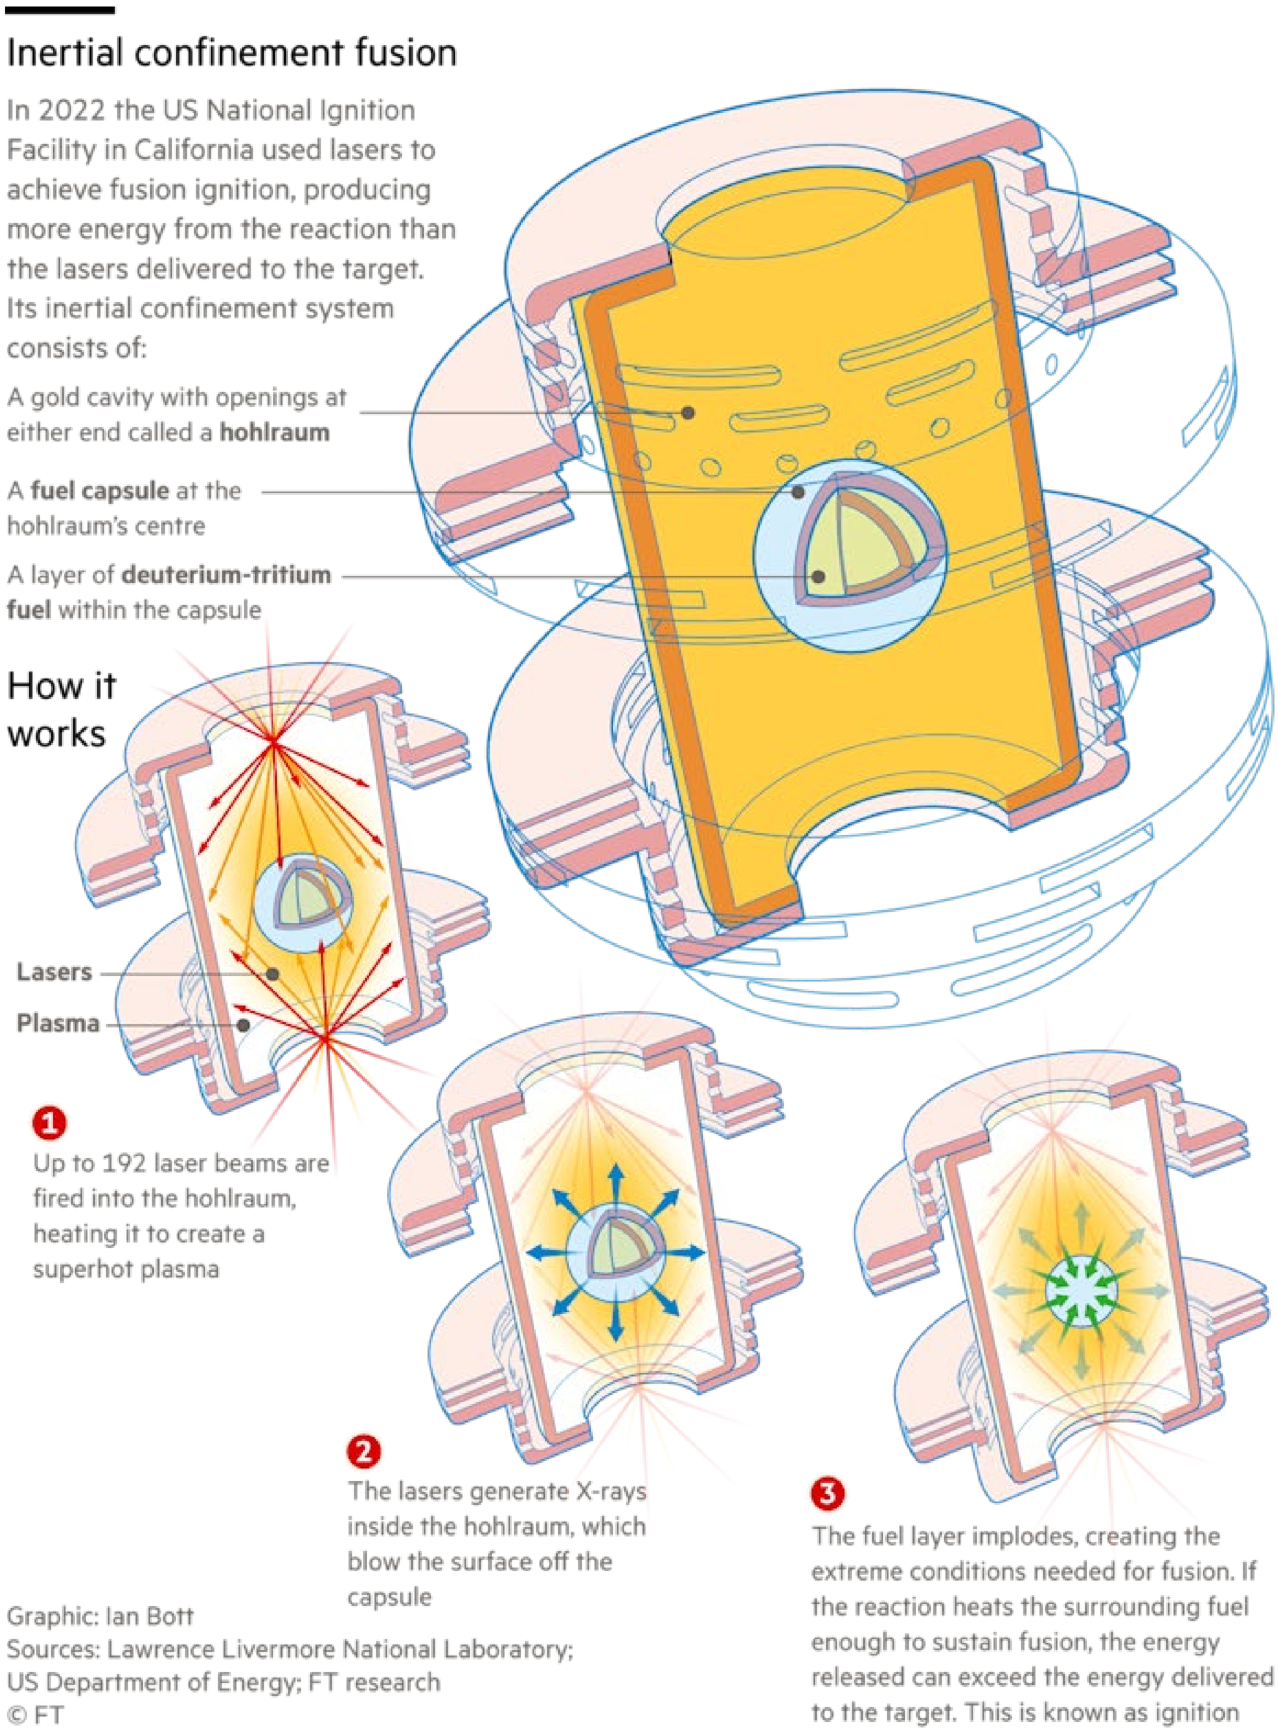

Figure 5. Experiment carried out at NIF (2022).

When a company states “we have reached $Q > 1$”, the investor should always ask: Q what, exactly? Plasma Q, target gain, wall-plug Q, engineering Q and net electric Q are very different quantities.

Deep-Dive Box: the different Q'sQ

There are five different ways of measuring the “energy gain” of a fusion reaction.

1. Plasma Q (or plasma Q)

The ratio between the energy produced by fusion inside the plasma and the energy used to heat the plasma itself. It is the most “generous” measure: it only considers the energy entering and leaving the plasma, ignoring the rest of the plant (magnets, pumps, cryogenics). It is the type of Q referred to, for example, by ITER's historic goal (Q ≈ 10).

2. Target gain

This is the concept specific to inertial confinement (NIF, discussed earlier): the ratio between the fusion energy produced and the energy actually delivered to the target by the lasers. It is precisely the measure NIF exceeded in 2022 (3.15 MJ produced against 2.05 MJ delivered to the target). But note: it only counts the energy “delivered to the target”, not what is needed to generate the lasers themselves.

3. Wall-plug Q

Here the view widens: the ratio between the fusion energy produced and all the electrical energy drawn from the power outlet (“wall plug”) to run the entire heating/laser system. For NIF this number is far lower than the target gain, because the lasers are extremely inefficient: enormous amounts of electrical energy are needed to generate those 2.05 MJ of laser energy.

4. Engineering Q

This is the “engineering breakeven” concept: it considers the entire plant, not just plasma heating. It is much closer to a realistic energy balance for a power plant, because it includes all the auxiliary systems needed to run it.

5. Net electric Q

The strictest and most commercially relevant measure: the ratio between the net electricity fed into the grid and the electricity consumed by the plant to produce it. Here conversion losses are also taken into account (turning heat or plasma energy into usable electricity is never 100% efficient). This is exactly the number that matters for knowing whether a fusion plant makes economic sense: a value greater than 1 is not on its own enough to guarantee it is worthwhile, but it is the minimum threshold for the plant to make sense as an energy source.

A company can legitimately announce “$Q > 1$” referring to *plasma Q* or target gain — a scientifically true and important result — while its *wall-plug Q* or its *net electric Q* remain deeply negative. This is not dishonesty: these are simply different measures, answering different questions. But to assess whether a technology is truly approaching commercialization, the only measure that really counts is the last one, *net electric Q* — the hardest to reach, and the one on which, to date, no project has yet provided a concrete demonstration at plant scale.

Scheme 6. The various concepts related to the Q indicator.

Engineering breakeven

For a commercial plant, the truly important concept is close to the ratio between net electrical energy produced and electrical energy consumed by the whole plant. This requires much more than an energetically positive fusion reaction: the plant must power magnets, cryogenics, pumps, vacuum systems, heating, diagnostics, electronics, control systems, fuel cycle, cooling and energy conversion. This is probably where the industrial contest will be decided.

The Tritium Problem

Tritium is probably one of the main "unknown unknowns" of commercial fusion. A D-T reactor consumes it, but the tritium available today is extremely limited relative to the needs of a future industry. The intended solution is breeding: the neutrons produced by fusion interact with the lithium present in the blanket (n + Li → T + ...), and the reactor should therefore produce its own fuel.

This creates an extremely strict requirement: the system must produce at least as much tritium as it consumes, also compensating for decay, losses, inventory, processing, maintenance and recovery times. The question becomes: is the tritium breeding ratio above 1 not just in theory, but in a real industrial system?

This doubt gives us the chance for a deeper look. According to the famous trilogy of stages of knowledge coined by Donald Rumsfeld, namely:

- ***Known knowns*** — things we know we know (e.g. we know tritium decays with a half-life of 12 years);
- ***Known unknowns*** — things we know we don't know, and can therefore plan to find out (e.g. "we don't yet know whether the tritium breeding ratio will exceed 1 in a real plant, but we know it's the question to test");
- ***Unknown unknowns*** — problems we cannot even imagine until we encounter them, because we have never had a real system similar enough to observe them on;

tritium breeding can, without fear of contradiction, be placed in the third category because it has never been demonstrated at real industrial scale, under continuous operating conditions, over long periods. All tests carried out so far have been:

- small-scale, in a laboratory;
- for short periods;
- without the cumulative neutron load of years of continuous operation;
- without having to manage all the real variables together at the same time (decay, process losses, corrosion of blanket materials over time, component aging, maintenance, inventory management)

When a system is so new that no one has ever been able to observe it operating at full scale and over extended periods, unforeseen interactions can emerge between phenomena that, individually, are well understood, but that no one has ever seen combine in that specific way. For example:

- lithium in the blanket, exposed for years to intense neutron flux, could degrade in ways not fully predicted by current models

- materials containing lithium could behave differently than expected when subjected simultaneously to extreme heat, radiation and mechanical stress for years
- tritium extraction and purification processes, which work in the laboratory on small quantities, could behave differently when scaled up to industrial volumes

None of these problems is “unknown” in the sense of being unknowable in principle — they are all understandable physical and chemical phenomena. But we do not yet know which of these problems will actually arise, in what combination, and with what severity, because no plant has ever run a tritium breeding cycle under realistic conditions for the time needed to find out.

This is precisely why fusion must be treated as a high-risk bet even beyond the problems already identified (materials, maintenance, costs): “unknown unknowns” cannot be priced precisely, because by definition they have not yet been discovered. This is one of the reasons a traditional financial valuation (such as a DCF) is nearly impossible for these companies, and it makes more sense instead to treat them as a real option — where value increases progressively only as each technical milestone is cleared and one of these potential “unknown unknowns” is finally tested at real scale.

Materials and Neutron Bombardment

The D-T reaction produces high-energy neutrons that strike the reactor's structures, causing crystal-lattice damage, embrittlement, activation, thermal degradation and reduced service life.

Embrittlement is the phenomenon by which a material that is normally strong and ductile (i.e. able to deform without breaking) progressively becomes brittle over time — it loses the ability to absorb stress and tends to break suddenly, even under loads it would previously have withstood without issue. As we saw when discussing the D-T reaction, fusion produces high-energy neutrons (the ones that carry off 80% of the reaction's energy, about 14.1 MeV). When these neutrons strike the reactor's structural materials:

- they displace atoms from their original position in the material's crystal lattice, creating microscopic defects (vacancies, dislocations)
- these defects accumulate over time, altering the material's internal structure at the atomic level
- the result is a material that, while looking physically identical on the outside, has become stiffer and more brittle — it loses resilience and can break under far smaller loads than it could withstand when new

This is a well-known problem in traditional fission reactors too, but in D-T fusion it is potentially more severe, because the neutrons produced have much higher energy (14.1 MeV) than those typical of fission.

If first-wall materials become embrittled over time, they must be replaced periodically — which means shutting down the plant, handling components that have become radioactive (through neutron activation), and bearing significant maintenance costs. This is exactly one of the reasons why “operating lifetime” and “component maintenance” are among the technical problems still unresolved in commercial fusion.

The first-wall and ***divertor*** material must therefore withstand extremely severe conditions. The divertor is a component specific to tokamaks (and other magnetic-confinement machines), typically positioned at the bottom of the plasma chamber. Its job is to manage the plasma's “exhaust”, made necessary by its imperfections; it in fact contains impurities (particles flaking off the walls), and above all it continuously produces helium “ash” (the product of D-T fusion) which, if allowed to build up, would dilute the fuel and cool the plasma. The divertor's job is to:

- diverting the "spent" plasma (the outer, cooler particles) toward a specific area of the chamber, separate from the plasma core
- extracting the helium produced by fusion and other impurities, preventing their build-up
- absorbing much of the heat and particle flux that the plasma discharges toward the walls

This job makes the divertor the chamber component that receives the most intense heat and particle load in the entire reactor — far more than the general first wall. It must therefore withstand extremely high temperatures concentrated on a relatively small surface, in addition to neutron bombardment and erosion caused by direct contact with the plasma.

The divertor is one of the reasons why the materials issue (one of the fundamental obstacles to industrializing the fusion process) is so critical: it must be built from materials able to simultaneously withstand extreme heat, particle erosion and neutron bombardment — often tungsten or advanced composite materials — and it must be replaceable periodically without shutting down the plant for too long. This is exactly the kind of "invisible" engineering problem (compared with plasma physics) that can determine whether a fusion plant is truly economically sustainable in the long run.

This obstacle, however, also represents a major investment opportunity: if fusion succeeds, someone will necessarily have to supply special steels, tungsten, composite materials, superconductors, ceramics, cooling systems and robotic components — a classic "picks and shovels" model.

Mohamed Abdou, professor emeritus of mechanical and aerospace engineering at UCLA, offers an even harsher estimate: with the technology currently available, first-generation plasma-facing components could fail on average every 20 minutes, and replacing a single reactor component would take "many months". Even a single failure in the plasma chamber forces the entire reactor to shut down — a further argument in favour of diversified exposure to materials suppliers rather than to a single reactor project.

6. The Supply Chain Might Be the Real Investment

The Fusion Industry Association highlighted in 2026 a 24% increase in fusion supply-chain spending in 2025: surveyed companies spent about $538 million and expect about $681 million in 2026. The most interesting figure is that 75% of suppliers have already invested to increase their production capacity — a sign the market is starting to prepare even before a commercial plant exists.

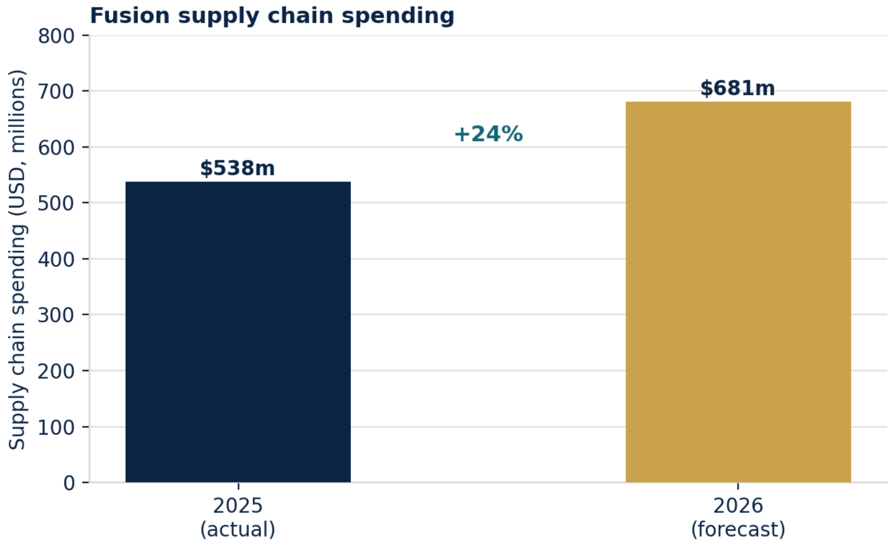

Figure 6. Fusion supply-chain spending, 2025 (actual) vs. 2026 (forecast). Source: Fusion Industry Association, Fusion Industry Supply Chain 2026.

Critical areas include power systems, heat management, vacuum, fuel cycle, first-wall materials and superconductors — exactly the kind of sector that can be attractive to invest in ahead of final commercialization.

Deep-Dive Box: the FIA

The **Fusion Industry Association** (FIA) is the trade association representing the private fusion industry worldwide.

Structure and founding

Founded in 2018 as an independent non-profit organization for the international fusion industry. It is headquartered in Washington, D.C. Its CEO is Andrew Holland, who has led the association as founding CEO since 2021, with more than twenty years of experience in science, energy and policy. It counts 42 fusion-sector companies as full members and more than 100 affiliate members, including engineering firms, suppliers, utilities, law firms, cleantech organizations and academic institutions.

It is a trade association that supports efforts to accelerate commercial fusion energy through advocacy (lobbying/representation) and education. In practice, it:

- acts as the sector's collective voice to governments, regulators and the media
- engages directly and transparently with regulators as they define regulation appropriate and proportionate to the sector's risk
- promotes international harmonization of regulation (as shown, for example, by its work on regulatory partnerships between the United States and Japan)

- publishes periodic reports on the state of the industry — including some of the very documents cited in this report: The Global Fusion Industry in 2026 and Fusion Industry Supply Chain 2026, whose data on capital raised, number of companies and supply-chain spending were used for the charts

The FIA is probably the most authoritative and up-to-date source of aggregate sector data (capital raised, number of companies, employment, supply-chain spending), precisely because it brings together almost all the world's major fusion companies as direct members — unlike individual corporate press releases, which only tell their own story. This is why our key figures ($14.24 billion cumulative, 56 companies, >16,000 employees) come from them: they are probably the best available census of the entire sector, rather than of a single company.

Scheme 7. The Fusion Industry Association: structure and functions.

The Fujikura Case

Fujikura, a Japanese manufacturer of HTS wire, is a particularly interesting example: HTS magnets are a fundamental technological element for many magnetic-fusion architectures. In 2026 Fujikura announced further investment to significantly increase its HTS production capacity, on top of investing directly in CFS and other fusion companies.

This creates a financially very interesting feature: Fujikura benefits from the growth of fusion without depending on the success of a single reactor. If Helion failed but CFS succeeded, Fujikura could benefit; if CFS failed but other tokamaks used HTS, it could still benefit. This is a far more diversified risk than direct exposure to a single technology.

7. AI and Fusion: a New Structural Demand

Artificial intelligence is rapidly increasing electricity demand. Data centers require continuous, high-density, reliable, 24/7-available energy — features that theoretically make fusion an ideal complement to intermittent renewables.

The Chinese case already offers a concrete number for this dynamic: according to forecasts by consultancy Wood Mackenzie, China's data centres will quadruple their electricity consumption, reaching 774 terawatt-hours by 2030 — a structural increase driven by AI-related workloads (Financial Times, September 2026).

This does not mean data centers will automatically make fusion economically viable, but that the market may be willing to pay a premium for stable, clean, available, high-density energy[2]. Microsoft is already the first example, through Helion; Google has entered CFS's commercial supply chain. The AI theme is therefore not merely narrative: it can become an important financing mechanism for the first generation of plants.

[2] Energy density means the amount of energy contained in (or obtainable from) a unit of mass or volume of a material or process. If we consider the reaction in Equation 1, it is easy to see that the energy released by a single fusion of nuclei, multiplied by the enormous number of possible reactions in a small amount of fuel, yields an amount of energy disproportionate to the weight and volume of the starting material — far more than any chemical reaction (such as combustion) and more even than nuclear fission, for equal mass.

8. The Financial Valuation Framework

Beyond the physical Q: the economic Q

In Scheme 6 we explored the various concepts of Q, but we limited ourselves to the engineering profile, i.e. related to electricity production.

Fusion's next great conceptual frontier will no longer simply be Q, but the economic Q, in extremely simplified form: the present value of the energy produced divided by the total cost of capital and operations.

A plant could technically produce net energy and yet be economically useless. To be competitive it would need to offer a favorable combination of cost of capital, operating cost, availability[3], capacity factor[4], lifetime, fuel-cycle cost, maintenance cost, materials cost and component-replacement cost. It is precisely on this point that many fusion-startup financial models remain highly uncertain.

The link between the technical side and the economic/financial side is direct and unforgiving. We discussed the divertor, embrittlement and neutron bombardment to explain that reactor components exposed to the plasma degrade physically over time and must be replaced periodically. Each replacement requires:

- shutting down the plant
- handling components that have become radioactive (through neutron activation)
- carrying out maintenance operations that are often complex, in an extreme environment

If these interventions are frequent or time-consuming, availability collapses — the plant spends much of the year shut down instead of producing sellable electricity. The capital invested to build a fusion plant is enormous (billions of dollars, as we saw for CFS and Helion), and that capital generates revenue only during the hours the plant is available to produce. If availability is low:

- the useful time to “pay back” the investment shrinks drastically
- the cost of the electricity produced (spread over fewer hours of production) rises
- the overall economic return on the investment worsens, even if the technology “works” from a physical standpoint

This is exactly why we introduced economic metrics such as capacity factor and availability into the list of variables that determine the “economic Q”: a reactor can technically produce net energy ($Q > 1$ in the physical sense) and still be a bad investment, if it spends half its time offline replacing components damaged by neutrons. This is one of the reasons why scenario 2 in (“technical success, economic delay”) is considered the most realistic: even

[3] Availability indicates the percentage of time a plant is technically able to operate — i.e. is not shut down for scheduled maintenance or faults — regardless of how much energy it actually produces in those moments.

[4] Capacity factor is a key metric in the energy sector: it indicates how much a plant actually produces compared with what it could theoretically produce if it always ran at maximum power, 24 hours a day, 365 days a year. It differs from availability because a plant can have high availability (it is almost always ready to run) but still a lower capacity factor, if for example it often runs at reduced power instead of at maximum. For a nuclear plant (fission, or prospectively fusion), however, the two values tend to be fairly close, because once started a nuclear plant almost always runs at full power — so the main factor that lowers both is downtime for maintenance.

solving the plasma-physics problems, these operational engineering problems remain to be solved, and they directly affect the economics.

A First Concrete Benchmark: Cost per MWh

Fusion companies generally avoid firm forecasts for the cost of their electricity, but some public benchmarks allow a first indicative comparison with competing sources (data drawn from the Nuclear Energy Agency):

Source	Estimated cost	Note
General Fusion (fusion)	\$64–\$73/MWh	Lifetime cost as estimated by the company
Commonwealth Fusion Systems (fusion)	~\$50/MWh	Target stated as indicative, not guaranteed
New nuclear fission	~\$71/MWh	Nuclear Energy Agency figure
Utility-scale solar PV	~\$38/MWh	Nuclear Energy Agency figure

Table 1. Cost of producing one MWh: comparison across different innovative production technologies.

A second data point, even more relevant to the net electric Q concept discussed above, comes from a peer-reviewed paper by Adam Jackson, a former General Fusion engineer: analysing a conceptual plant design published by the company in 2023, the study estimates that only about 40% of the electricity generated would actually be available for sale to the grid — the rest is reabsorbed by the plant's own energy needs. This means that, for the same net electricity delivered to the grid, a fusion plant would need to generate more than twice the electricity of a conventional thermal plant — a multiplier that weighs directly on capital costs and therefore on the final \$/MWh.

Valuing Fusion as a Real Option

Valuing a fusion company with a normal DCF is almost impossible: future revenues are zero today, possible in 5-15 years and potentially enormous after commercialization — but with an uncertain probability of success. Valuation should therefore be treated as a real option:

Value ≈ (probability of success × future industrial value) − capital required to reach the next milestone

This completely changes how one should invest: value rises enormously when a company clears verifiable technical milestones, which we have already discussed and list again below:

1. Plasma
2. Sustained plasma[5]
3. $Q > 1$
4. Engineering breakeven
5. Net electricity
6. Continuous operation and high plant availability
7. First commercial plant and replicability

[5] It is essentially a plasma that manages to remain in the conditions needed for fusion for a prolonged period of time, instead of existing only for a fleeting instant.

The share price or the company's valuation should be analyzed based on whether these milestones have been cleared.

Risk profile7Diversified energy companies with strategic exposure

- being permanently connected to the electricity grid
- selling electricity to a customer (utility, hyperscaler, industry) through a real contract (like the PPAs discussed for Helion and CFS)
- generating revenue that, over time, allows the investment to be paid back

We have already given some concrete examples: Helion's Polaris/Orion (with a declared 2028 target), CFS's ARC (early 2030s), or STEP in the United Kingdom (around 2040) are all aiming, if successful, to be this first commercial plant.

The next step, and the real final test of fusion as an industry: demonstrating that the first plant is not an isolated case, a one-off built at extraordinary cost and in extraordinary time, but that it can be built again, multiple times, at falling cost and in reasonable time. There is an enormous difference between building a machine that works (even at extremely high cost, with dedicated engineering teams and near-zero margin for error) and building a fleet of replicable plants, with an industrial supply chain, a mature supplier base, standardized production processes and costs that fall over time (the classic industrial "learning curve"). This is exactly why the entire supply-chain section (HTS magnets, materials, power electronics, fuel cycle) is so important to the investor: replicability does not depend only on the company that built the first reactor, but on an entire industrial ecosystem that must be ready to supply, in series, all the components needed to build the tenth, the hundredth, the thousandth plant — at ever-lower cost. A company that reaches "first commercial plant" but fails to demonstrate "replication" risks remaining, in effect, a company that built one extraordinarily expensive one-off — validating the technology without creating a truly scalable business. Only once replicability is demonstrated can the company's value truly reflect the potential of an entire market, rather than just the value (however historic) of a single plant.

The Valuation Problem: A Numerical Example

Let us assume, in purely conceptual terms, four scenarios for a pre-commercial fusion company:

Scenario	Description	Potential value (index)
A	Commercial success	100
B	Working technology, 5-year delay	40
C	Working technology, uncompetitive economics	10
D	Failure	0

Table 2. Illustrative example (not a forecast) of scenario distribution for a pre-commercial fusion startup.

The correct price is not the value of scenario A, but the probability-weighted sum ($P = \Sigma p_i \times V_i$) minus the capital needed to reach the next milestones. This is why fusion-startup valuations must be treated as venture-capital valuations, not as conventional DCFs.

The Most Important Variable: Capital

Fusion is a capital-intensive technology: CFS has raised about $4 billion, Helion $1.5 billion, General Fusion more than $400 million before listing — figures not necessarily sufficient to build a fleet of plants. The investor's real risk is therefore not only “whether the technology works”, but also “whether the company will have enough capital to get to the point where the technology works”. A technology can be valid and the company can still fail for lack of capital.

9. SWOT Analysis

SWOT of Nuclear Fusion (Sector)

STRENGTHS	WEAKNESSES
• Extremely high energy density • Very low operating emissions (no CO_2 during the reaction) • Potentially abundant fuel (deuterium from seawater) • Intrinsic safety relative to fission • Potential baseload[6], unlike solar/wind • Possibility of global industrialization once the technological problem is solved	• Technology not yet commercially demonstrated • Potentially enormous capex for first-of-a-kind plants • Complex fuel cycle (tritium is a structural problem) • Materials: neutron bombardment reduces component lifetime • Maintenance: possible frequent replacement of plasma-facing components • Availability: producing energy for a few seconds ≠ producing it 24/7
OPPORTUNITIES	**THREATS**
• AI and data centers: growing demand for stable, clean energy • Industrial decarbonization (steel, aluminium, chemicals, cement) • Energy security: reduced dependence on imported fossil fuels • New supply chain: HTS, advanced materials, power electronics, cryogenics, fuel cycle • Geopolitics: industrial advantage for the first country with competitive technology • AI-assisted engineering: simulation and ML to reduce R&D time and cost	• Alternative methods: solar+batteries, advanced fission[7], geothermal, storage • Delays: a 5-year slip can have an enormous impact on present value • Cost overruns, as is typical in large nuclear projects • Regulation: licensing and grid connection as bottlenecks • Supply chain: specialized materials and components potentially scarce • Funding risk: capital needed for many years before revenues • “Theranos moment”: some investors and experts cited by the Financial Times fear that excessive promises on timelines and costs could damage the credibility of the whole sector, drawing an explicit parallel with the Theranos affair (September 2026)

Compared SWOT: the Main Pure-Plays (Nuclear Fusion Players)

[6] Baseload indicates the minimum, constant level of electricity demand a grid must meet at all times, 24 hours a day — and, by extension, also refers to the energy sources capable of providing that constant, continuous output, regardless of the time of day or weather conditions.

[7] By advanced fission we essentially mean Small Modular Reactors (SMRs), which we discussed in our Deep-Dive of August 8, 2025.

	Helion	Commonwealth Fusion Systems	General Fusion
Strengths	Compact FRC; direct energy conversion; PPA with Microsoft; ~$1.5bn raised; ~$15.5bn valuation; Orion under construction	Established tokamak + HTS/REBCO; SPARC/ARC; 5 peer-reviewed papers in 2026; ~$4bn raised; PPA with Eni; deal with Google	>20 years of R&D; original engineering approach (LM26); public capital; first listed pure-play
Weaknesses	Technology not yet commercially proven; high execution risk; very aggressive timeline; fuel cycle still to be demonstrated	Large capital requirement; tokamak complexity; fuel cycle and materials; commercialization still years away	Technology far from commercialization; capital much lower than CFS/Helion; preliminary results; high stock volatility
Opportunities	If Orion works, an enormous first-mover advantage; Microsoft could be the first of many hyperscalers	Could become the “Tesla” of the first generation of fusion plants if SPARC confirms expectations	If LM26 validates Magnetized Target Fusion, GFUZ could become one of the sector's most explosive stocks
Threats	A delay would not destroy the technology but would drastically reduce the financial value attributed to it	Helion could get there first; other, cheaper architectures; risk of cost overruns on ARC	Public listing as a potential disadvantage; risk of further capital raises well before commercialization

10. Technology Ranking and Investable Universe

The qualitative assessment, as of September 2026, is as follows. This table should not be read as a ranking of "who will win": it represents, rather, the relationship between technological maturity and optionality.

Company/project	Technology	Maturity	Capital	Risk	Potential
CFS / ARC	Tokamak + HTS	High	Very high	High	Very high
Helion	FRC	Medium/high	High	Very high	Extreme
TAE	FRC / advanced fuel	Medium	High	Very high	Extreme
General Fusion	MTF	Medium	Medium	Extreme	Very high
STEP	Spherical tokamak	Medium	Public	High	High
ITER	Tokamak	Very high (scientific)	International public	Low scientific / high economic	Strategic
First Light	Inertial/hybrid	Medium	Medium	Very high	Very high

Table 3. Technology investment ranking, the author's qualitative assessment as of September 2026.

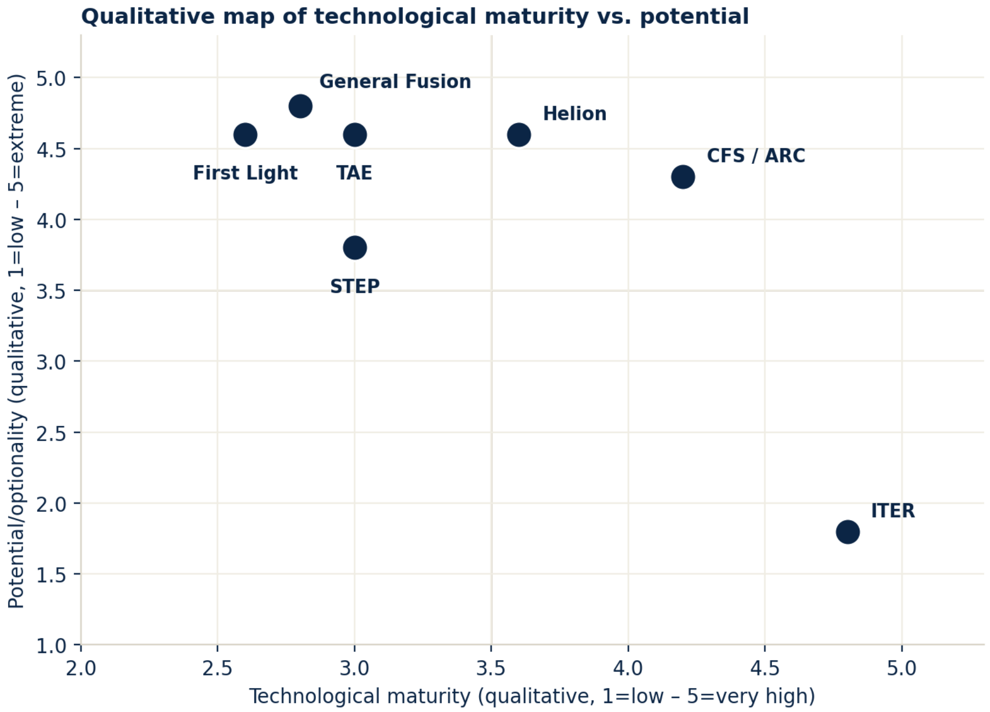

Figure 7. Qualitative map of technological maturity vs. potential/optionality, indicative positioning derived from Table .

.

The Publicly Investable Universe

The problem for the investor is that most fusion leaders are not listed, creating a paradoxical situation: the most interesting technology is mostly private. The public universe consists of companies with indirect exposure or of just-listed companies.

Exposure	Ticker	Note
General Fusion (direct)	Nasdaq: GFUZ	The purest case of listed fusion, but also one of the most speculative investments: it should be treated almost as a listed venture investment.
TAE (indirect)	Nasdaq: DJT	A potential public proxy for TAE if the Trump Media deal closes; should not be treated as a fusion pure-play, being a complex hybrid investment.
Eni (strategic)	—	Direct exposure to CFS, PPA, energy and fuel-cycle expertise, industrial capacity: fusion risk diluted within a much larger group.
Fujikura (supply chain)	—	Exposure to HTS: it doesn't matter which company wins, but how much demand for superconductors grows; a more diversified form of optionality.

Table 3 4. A few examples of investments that could be made in the nuclear fusion sector. The content of this table should not be considered an investment proposal or advice, but only an illustration. Indeed, the selection was chosen specifically to exemplify the risk/return differences of various approaches, such as direct or indirect investment, or investment in the "picks and shovels".

11. Portfolio Strategy: the Barbell Strategy

In a sector with such high risk, adopting a prudent approach seems almost obligatory to us. The best way to diversify this enormous risk is a “distributed and balanced” exposure (*barbell strategy*) to fusion: avoiding concentrating everything on a single technology, by building the portfolio like a gym barbell: weight concentrated at the two ends, almost nothing in the middle.

The basic concept is as follows: instead of spreading capital evenly across the whole risk spectrum (from very safe to very risky), investment is concentrated at two opposite extremes:

- on one side, very safe/conservative positions
- on the other, very risky but extremely high-potential positions

deliberately avoiding the “middle zone” — moderate risk/return investments which in theory offer a compromise, but which in practice provide neither the safety of the low extreme nor the explosiveness of the high extreme. The idea, made famous above all by Nassim Taleb, is that:

- the “safe” part of the portfolio protects capital and ensures that, in the worst case, not everything is lost
- the “risky” part offers exposure to rare but potentially enormous events (so-called positive “black swans”) — where even a small slice of capital, if the event occurs, can generate a disproportionate return

The “middle” is avoided because it is often the most deceptive zone: it looks balanced, but in reality it exposes one to risks that are not fully understood without offering any real exceptional upside potential.

Working across the four levels set out in Table 4,

Level	Description	Examples	Risk profile
1 — Core	Diversified energy companies with strategic exposure	Eni	Moderate
2 — Picks & shovels	Producers of HTS, materials, power electronics, cryogenics, vacuum, components	Fujikura and other suppliers	Medium-high, diversified
3 — Fusion pure-play	Reactor developers, mostly private (VC, PE, secondaries, specialized funds)	Helion, CFS, TAE, General Fusion	Very high
4 — Public venture bet	Listed venture position, extremely speculative	GFUZ	Extreme

Table 5. Four-level structure for a professional exposure to fusion (barbell strategy).

instead of investing a “medium” amount in a middle ground (for example, a generic energy company with vague exposure to fusion, neither safe enough nor explosive enough), a sound solution — in our view — is to combine a solid “core” (Eni) with small very-high-risk/return positions (the pure-plays), avoiding intermediate positions that would offer neither real protection nor real upside.

Fusion is the textbook example for applying this strategy: it is a sector where the probability of total failure of a single company is high, but where the potential payoff in case of success is enormous (a global energy market worth trillions of dollars). In such a context, putting “all your eggs” in an intermediate-risk position would probably be the worst choice: one would still be taking on significant technological risk without capturing the explosive upside that justifies that risk.

For a professional (but also institutional) European investor, the indicative hierarchy is as follows:

- 1. Eni — diversified strategic exposure
- 2. Fujikura and other HTS/materials suppliers — picks and shovels
- 3. CFS — probably the best technological profile, but not listed
- 4. Helion — greater optionality and greater technological risk
- 5. TAE — great potential but complex corporate exposure
- 6. General Fusion / GFUZ — listed venture capital
- 7. ITER/STEP — fundamental from an industrial standpoint, less directly investable

unless the investor happens to be one of these targets themselves (such as Eni, through its partnership with CFS).

Why Avoid a Single Investment in GFUZ

The fact that General Fusion is the first listed fusion company is important from a historical standpoint, but it does not mean it is the best investment. A pre-commercial fusion company's market capitalization embeds an extremely wide distribution of scenarios (see Table 5), and the right price is not the value of the best-case scenario alone, but the probability-weighted sum, net of the capital needed to reach the next milestones.

12. The New Role of the State and 2026 as a Turning Point

The sector is becoming a public-private partnership. States fund laboratories, materials, infrastructure, supply chain, regulation[8], training and research; startups fund prototypes, engineering, commercialization, IP and manufacturing. This model resembles the aerospace industry more than software venture capital, and it matters for valuation: the probability of technological success no longer depends only on private capital, but also on the state's willingness to build the necessary infrastructure.

2026 is probably the first year in which it is reasonable to speak of fusion as an emerging industry — not because it has been commercialized, but because billions in capital, industrial customers, PPAs, prototype plants, supply chains, listed companies, government roadmaps, specific regulation, investment from major energy groups and potential hyperscaler demand have all appeared at the same time.

The U.S. DOE (Department of Energy) published a roadmap in June 2026 explicitly aimed at delivering fusion pilot plants by the mid-2030s — a qualitative shift from traditional public fusion research.

The decisive question of the next decade will not be "which company will achieve fusion first", but "which company will manage to build the first plant that can be replicated at competitive cost". A machine can work once; a power plant must run 24 hours a day, 365 days a year, for decades — and must be able to be built again and again. The real test will therefore be industrial replicability.

13. The China Factor

So far, this report has described fusion almost exclusively through the lens of the United States, the United Kingdom and Europe. That is an incomplete picture: China is emerging as a leading player across the entire nuclear sector, fusion included, and is doing so with a combination of industrial scale, centralized state planning and capital that has no direct equivalent elsewhere.

The macro picture is stark: of the 58 new conventional nuclear plants currently under construction worldwide, more than half are in China, which is also investing heavily in an emerging class of small modular reactors (SMRs) while pouring R&D resources into what the article calls energy's "holy grail": nuclear fusion. With cost and environmental concerns slowing nuclear development across much of the rest of the world, China's industrial competitiveness opens up real room to dominate another industry of the future.

The most relevant signal for an investor is probably political-industrial rather than financial: nuclear fusion is one of eight technology areas identified as priorities in the Chinese Communist Party's latest five-year economic development plan. According to research from the Special Competitive Studies Project, a US think-tank, China has already mobilised roughly three times as much capital into "commercialisation-relevant" fusion projects as has gone into comparable projects backed by the US Department of Energy.

Damien Ma, a non-resident expert at the Beijing-based research group Gavekal, argues that China will have the "world's most dynamic and significant nuclear industry" over the next decade, "by a wide margin". In his view, a sustained build-out in the US runs into tight capital intensity and regulatory hurdles, while China aims to strengthen an already "formidable" supply chain, convert its domestic know-how into an export-oriented industry,

[8] Regulation is one of the fundamental building blocks for creating a mature industry: think of the recent cases of cryptocurrencies and artificial intelligence.

and reach the clean-energy "holy grail" represented by fusion power. This dynamic, moreover, is unfolding against a backdrop of rising US-China tension, including sweeping US controls on technology sales to the country — a theme that makes China's fusion trajectory relevant geopolitically as well as industrially, and that reinforces the point already made in this report's SWOT about the industrial advantage that accrues to the first country with competitive technology.

NovaFusionX and Junhe Atomic: China's Private Players

Alongside the large state-owned groups, an ecosystem of private companies is also emerging in China, often with links to state funds and public research institutes, seeking to develop new nuclear and fusion technologies.

NovaFusionX, based on Shanghai's eastern outskirts, is developing a Field-Reversed Configuration (FRC) — the same magnetic plasma-confinement technology used by Helion Energy. Its founder, Guo Houyang, is no minor figure in the FRC world: he previously led the US Department of Energy's largest FRC research programme at the University of Washington, and later became experimental lead and head of the experimental department at TAE Technologies, the US fusion developer already covered in this report. By his own account, his early research laid important technical groundwork that later fed into Helion's own approach.

NovaFusionX aims to achieve "first plasma" by the end of 2026 and a "net energy gain" (where energy output exceeds the input needed to sustain the reaction) by 2029 — a crucial step toward commercialisation. The company is in talks with local Chinese governments to site its first demonstration reactor, which it hopes to build in the early 2030s at a cost of about Rmb 3 billion ($440 million). It has already raised Rmb 2.4 billion ($360 million) from investors, with a valuation of more than Rmb 10 billion ($1.5 billion).

Junhe Atomic, also based in Shanghai, instead develops fission-based small modular reactors (SMRs). It was founded in 2025 by Tian Jiashu, previously deputy chief engineer at two of China's largest state nuclear groups, China National Nuclear Corporation and China General Nuclear Power Corporation — a profile that signals how directly the new generation of private Chinese nuclear startups draws on know-how accumulated within the big state programmes. Junhe has raised Rmb 300 million ($44.1 million) in venture funding, including backing from HongShan Capital (formerly Sequoia China, led by Neil Shen), and is targeting a demonstration project operational around 2031. An executive at the company, who asked not to be named, also raised an intriguing application for a fleet of small reactors: replacing China's ageing coal plants as they are decommissioned, reusing their sites and posing the far-from-trivial question of what happens to the plants, sites and workers of a structurally declining energy technology.

A more cautious note comes from Zheng Mingguang, chief designer at the State Power Investment Corporation, who at a Beijing conference in April 2026 acknowledged that much work is still needed to develop new materials to support the nuclear industry. He nonetheless firmly asserted China's strengths: "we have the strongest design teams, R&D teams and testing facilities. We have the strongest and most complete industrial chain, and we have the best talent pool — by contrast, not a single foreign country has a complete industrial chain".

For the investor, the main implication is not that China "will win" the fusion race — no source, Chinese or Western, offers sufficient evidence for that claim today — but that the number of serious, well-capitalised players with high-level technical expertise is significantly larger than a reading centred only on the United States and Europe would suggest. If anything, this reinforces this report's central thesis: betting on the technological ecosystem and the supply chain, rather than on a single company or a single country, remains the most rational strategy.

14. Scenario Analysis: 2030–2040

Scenario	Description	Probability (qualitative)
1. Fusion breakthrough	Helion, CFS or another player demonstrates commercial net electricity production before the mid-2030s: massive capital inflows, re-rating of fusion companies, a supply-chain boom, strong hyperscaler interest.	Low/medium
2. Technical success, economic delay	Fusion works, but the first plants are too expensive; another 5-10 years are needed to reach competitive economics. Still a positive scenario for the sector, but with highly volatile valuations.	Considered the most realistic
3. Scientific success, commercial failure	Plasma works, but materials, tritium, maintenance, availability and capex make the technology uncompetitive: the most dangerous scenario, because demonstrating fusion is not the same as creating an industry.	Not negligible
4. Fusion loses the economic race	Solar + storage, advanced fission, geothermal and grid-scale batteries[9] become cheaper; fusion remains scientifically possible but economically irrelevant. Pure-plays can lose almost everything; diversified suppliers, by contrast, continue to generate value.	Not negligible

Table 6. Scenario analysis 2030–2040 for the fusion industry. Qualitative probabilities, not numerical estimates.

[9] Grid-scale batteries (literally "grid-scale") are very large electrical energy storage systems, designed not for a single building or vehicle, but to be connected directly to the electricity grid and serve an entire area, city or region.

15. Contrasting Voices: Optimists and Skeptics

Before the conclusions, it is worth reporting some explicitly divergent positions gathered by the Financial Times (September 2026) from industry insiders, useful for balancing the picture described so far.

Alain Bécoulet, chief scientist at ITER, is among the most skeptical on timing: he argues that connecting a fusion plant to the electricity grid "will never happen" within the next 5-10 years, and that even beyond that horizon experimental facilities would likely produce only energy "bursts" lasting about 20 seconds — not enough to interest any grid operator ("you cannot sell anything out of an intermittent source", in his words). In his reading, fusion will only enter the energy mix in the "very last decade of this century".

Steven Cowley, director of the Princeton Plasma Physics Laboratory and among the world's foremost plasma physics experts, offers a more nuanced position: he says he is "confident" that toward the end of the 2030s there will be a pilot plant able to produce more electricity than it consumes, but insists on the need to keep the concepts of "pilot plant" and "commercial plant" clearly separate in communication to the public and investors: the early plants would be more like "technology experiments" to be tested for a decade, with a truly commercial plant not before the middle of the century.

An investor interviewed anonymously instead offers a more optimistic reading, criticizing the slowness of public laboratories and large international projects: in their view, these "have been used to thinking about ITER as being on the critical path to commercial fusion", pushing implicit timelines well past 2050, while private capital, public markets and well-coordinated public funding today represent "the best chance of delivering commercial fusion energy before 2040".

This plurality of views — among those working on large public projects, those leading academic research, and those investing private capital — reflects well the structural uncertainty of the sector described in this report, and indirectly confirms why a diversified investment strategy (barbell strategy) remains, in our view, the most rational choice.

16. Conclusions and Final Assessment

Nuclear fusion is not yet a traditional energy investment: it is a technological option on the future structure of the world's energy system. The potential is enormous, but the probability of success remains uncertain enough to make an "all in" approach unsuitable.

The most rational strategy therefore remains a *barbell strategy*: on one side, diversified energy companies with strategic exposure to fusion; on the other, small, very-high-potential positions in pure-plays and the supply chain.

The claim that Helion could "beat ITER" is journalistically catchy but technically misleading: ITER and Helion are not two companies racing toward the same finish line, but two different models. ITER is trying to demonstrate, scientifically and engineering-wise, that fusion works at experimental scale; Helion is trying to demonstrate that a startup can leap straight to commercial production; CFS represents a third path, tokamak physics made compact through HTS magnets; STEP represents a fourth path, an entire national industrial supply chain.

It is precisely this plurality that makes the sector interesting for investors. In 2026 we still do not know which technology will win, but we know something far more important than five or ten years ago: the market is starting to fund the attempt to find out[10].

The real turning point will not be the day an experiment sets a new temperature record, but the day a company can simultaneously declare:

> *net electricity + continuous operation + acceptable availability + positive economics + replicable plant design*

When all five of these conditions are true, fusion will stop being a technological promise and become a new industry. From the investor's standpoint, that will probably be the moment when most of the return has already been captured by the private companies that cleared the early milestones. For this reason, 2026-2030 could be the most interesting period to build optional exposure to fusion, but not necessarily the moment to bet on the eventual winner.

The most robust strategy is therefore to invest in the ecosystem — reactor developers, energy companies, HTS, materials, power electronics, fuel cycle, engineering. In other words:

> *don't necessarily try to guess which will be the first commercial reactor; try to identify which technologies will be indispensable regardless of which reactor wins.*

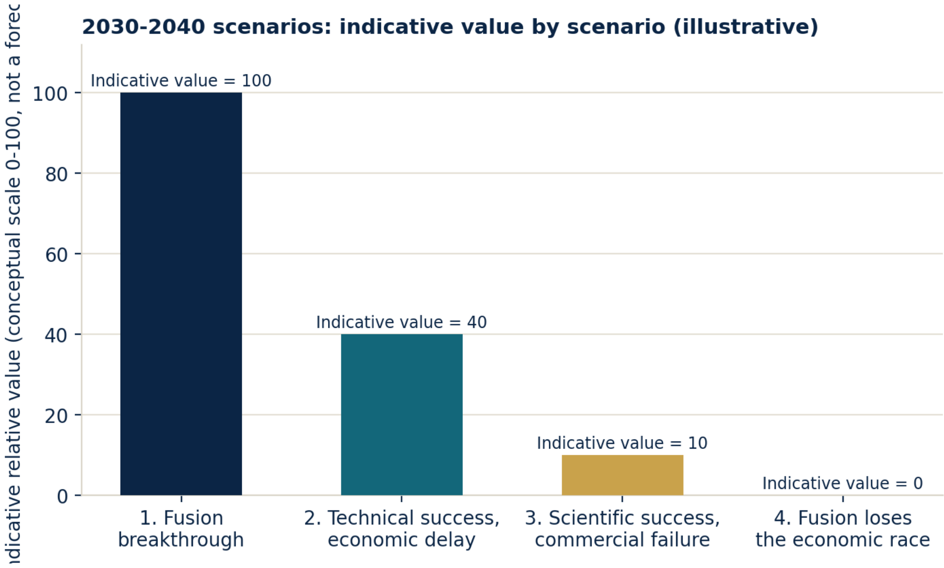

Figure 8. Conceptual representation (not a forecast) of the indicative relative value of the four 2030-2040 scenarios.

[10] For our "10 readers", this should ring a bell, because it is exactly the development paradigm that characterizes another great frontier tech industry we have covered extensively. Which one?

Main Sources

- Megaprojects, Helion Energy: Fusion Before ITER, January 2026.
- ITER Organization, new ITER baseline.
- U.S. Department of Energy, Fusion Science and Technology Roadmap, 2026.
- Fusion Industry Association, The Global Fusion Industry in 2026.
- Fusion Industry Association, Fusion Industry Supply Chain 2026.
- Commonwealth Fusion Systems, ARC Physics Basis Papers, 2026.
- Helion Energy, Polaris and Orion.
- Eni, strategic partnership with Commonwealth Fusion Systems.
- UK Government / UKAEA, STEP Fusion.
- General Fusion, LM26 and Nasdaq listing.
- TAE Technologies / TMTG, documentation on the fusion and the planned corporate transaction.