Technical Update from Geretsried:

What we built, what we learned, and what comes next

By Matt Toews, Co-founder and Chief Technology and Operating Officer at Eavor

Over the last several years, we've spent most of our time building rather than talking. As Eavor moves from technology development into large-scale commercialization as a technology licensor, we want to share a little more about our business. This extended analysis is part of that shift.

This, and future articles will explore what we've built, what we've learned, where the technology is going, the commercialization of our product, and how we think closed-loop geothermal fits into the future global energy system. The goal is to provide a technical and commercial perspective grounded in first-hand experience.

This technical update focuses on the Geretsried project: what we achieved technically, the challenges we faced, and why we believe it materially changes the commercial outlook for geothermal energy.

TLDR

Eavor's technology is what many would consider a "holy grail" in energy: cost competitive, secure, geographically scalable, dispatchable, locally sourced, and low carbon — all at the same time.

The technology works, which we proved with the first loop at Geretsried. It is cost competitive today for district heating or in certain power markets. And we have a clear line-of-sight to "geothermal anywhere" (<$75/MWh electricity cost in average geothermal gradients), by continuing down a typical learning curve and by building Eavor-Loop™ deeper.

What we do and why

The world increasingly needs energy systems that are secure, scalable, dispatchable, locally sourced, and low carbon — all at the same time. Although this is a core challenge for our civilization, very few technologies can realistically satisfy all of those constraints simultaneously.

Eavor's technology is one of them. Our mission is to enable clean geothermal energy almost everywhere.

Our approach is based on standardized closed-loop geothermal systems ("Eavor-Loops") that do not rely on rare geological hotspots, permeable reservoirs, or hydrothermal resources. Instead, Eavor-Loop functions as a large-scale subsurface heat exchanger; essentially an engineered underground radiator constructed deep within the Earth's crust.

The important distinction is that the system works in the broad range of geological conditions that exist across much of the world, not just in exceptional geothermal locales.

That creates a very different commercial opportunity. The same core technology can scale from district heating systems for small communities to multi-gigawatt projects supporting industrial facilities and data centers.

Just as importantly, the technology avoids many of the constraints facing other large-scale energy solutions. We don't have the same regulatory or stakeholder concerns as nuclear SMRs. We do not rely on fracking. Water consumption is very low. Surface land use is limited. We don't use rare earth metals. We have no ongoing fuel requirement. The supply chain is completely anchored in North America, Europe, and Japan.

The combined effect of these attributes means we own the technology high ground; it is exactly what the market wants, whether it is for critical district heating infrastructure or baseload power for German society, power for Japan, or power for global data centers. In many ways, Eavor draws comparisons to fusion energy, with one key difference -- Eavor's tech works today.

Naturally, claims like this invite skepticism, and rightly so. Several articles and analyses have attempted to assess Eavor's technology and economics, often without access to the underlying technical, commercial, or contextual information. We hope this extended update provides useful information to support future analysis.

There are three main questions that should be asked. This update will address the first in detail, while providing a brief treatment of questions two and three (more updates to come on those).

1. Does the technology work?
2. What is the cost of energy now and how will it trend in the future?
3. Can Eavor realistically execute at global scale?

In summary:

1. **Yes.** We've already demonstrated the core technology at meaningful scale. Each lateral pair we've drilled is the same scale for large commercial projects. The system works.
2. **The economics are already competitive** in European district heating, or certain power markets with high gradients, based on the performance of the last two lateral pairs drilled at Geretsried. By taking advantage of the learning curve (Wright's law), and maturing existing deep drilling technology we have clear line of sight to be competitive for power *in average geothermal gradients everywhere* (less than $75/MWh power price).
3. **Yes, by leveraging existing supply chains and a global network of partners**. One of the key enablers is Eavor's technology licensing business model. By working alongside world-class partners, we can leverage existing global capabilities, supply chains, and operational expertise rather than trying to build everything ourselves.

Does the technology work?

Geretsried Project Overview

Geretsried is our first commercial-scale Eavor-Loop project. It is situated on the site of a failed traditional geothermal project in Bavaria, Germany. Wells were previously drilled attempting to produce hot water from the subsurface and found that the rock was "hot but dry". This is a key benefit of Eavor-Loop technology in that it harvests geothermal energy in areas previously thought (or, in this case, previously shown) to be impossible.

Geretsried is our first-of-a-kind or "FOAK" project. FOAK is a term used in clean tech to describe the very difficult task of financing and building the first project with novel technology. It is capital intensive and scale brings its own difficulties with new technology.

We had previously built "Eavor-Lite", a proof-of-concept demonstration project in Alberta. Geretsried moves the system into commercial-scale, FOAK field execution.

Geretsried was originally designed for four Eavor-Loops, each with 12 lateral pairs approximately 6 km long, giving approximately 72 km of total reservoir exposure. These Eavor-Loops would provide combined heat and power to the surrounding municipalities. As previously published, the project was designed for nameplate capacities of 8.2 MW electricity (64 MW thermal), and the surface facility was designed accordingly.

The heat produced by the Eavor-Loops is directed to either the power plant or to district heating offtakers. For the first few years, all heat harvested from the earth is directed to the ORC power plant and converted to electricity. Heat from the project will be eventually sold to offtakers, including the local municipality of Geretsried. However, district heating utilities do not usually commit to critical infrastructure using new technologies that are not derisked; this is why it's so important to prove the technology and derisk the "resource" with a first-of-a-kind loop.

Geretsried serves three strategic purposes, outlined below, which have all been achieved.

Objective	Results
1. Prove the technology at meaningful scale.	• Intersected wells at >8 km MD depth using proprietary Eavor-Link™ active magnetic ranging technology • Actively cooled drilling tools using proprietary insulated drill pipe • Sealed laterals with proprietary Rock-Pipe™ system • Set and retrieved multiple whipstocks • Produced thermal output as predicted and successfully generated power
2. Demonstrate the initial point in the learning curve for the European heat market.	• Using the drilling performance/ cost of the last two lateral pairs drilled, we are competitive today in the European heat market
3. Prove the learning curve dynamics through repeated execution. This is a critical factor in the roadmap to widespread commercial viability for all new technologies (ex: shale gas/oil, wind, PV cells, batteries, etc.)	• We encountered many challenges but ultimately solved them and demonstrated a clear improvement in performance, reducing time/cost by 70% for the last lateral compared to the average of the first four pairs. • Demonstrated we can construct a lateral pair from start to finish in one run • Demonstrated the ability to go deeper/ hotter

Drilling Execution: overcoming challenges and demonstrating the learning curve

There were a number of key challenges encountered at Geretsried, all of which were solved and demonstrated the learning curve in action.

Hole cleaning and wellbore stability

The original drilling fluid system was designed based on successful traditional geothermal wells in southern Germany, but was not well suited to some of the carbonate formation behavior we encountered in the laterals, particularly weak bedding planes, hole-cleaning challenges, and the resulting risk of stuck pipe. This hurt performance early in the campaign. The fix was practical: we changed to an engineered mud system better suited to the formation and that will be a base case in future loops. After that change, the wellbore stability issue was mitigated, as illustrated by the below imaging logs.

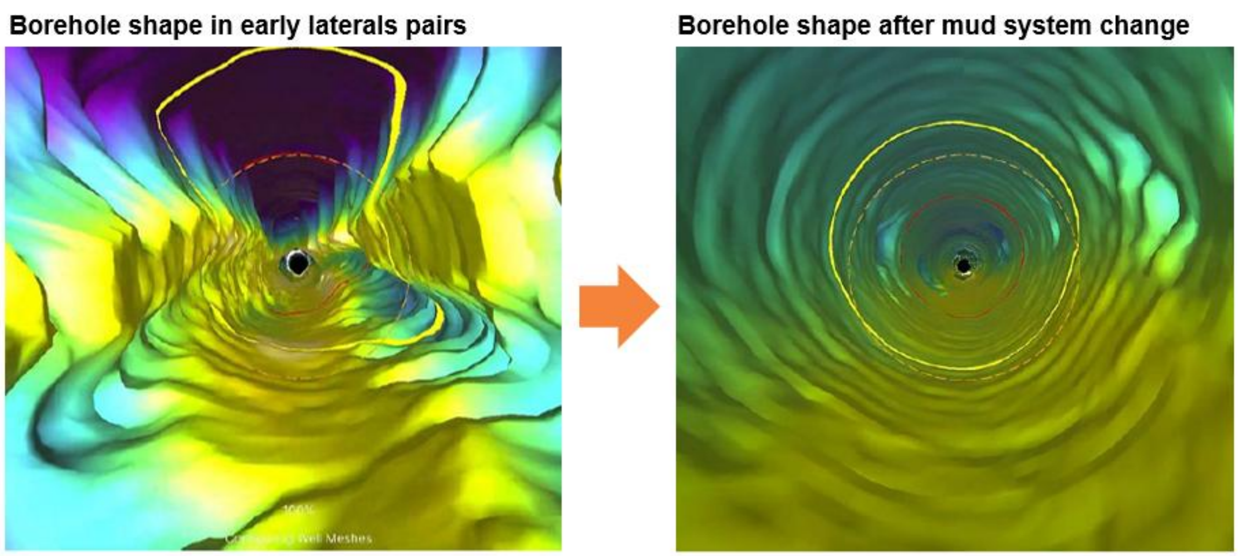

Extreme drilling dynamics and dysfunction

The combination of deep hard rock, long laterals, high friction, and difficult hole conditions created severe shocks, vibration, and torque dysfunction. This did not just slow rate of penetration; it also damaged tools, shortened bit life, and increased trips out of hole. As the campaign progressed, we improved the bit and BHA design, adjusted drilling parameters, changed the fluid system, improved lubricity and hole cleaning, and used drilling dynamics mitigation tools. Those changes translated into better ROP, longer bit runs, and fewer drilling dysfunction issues in the later wells. Ultimately, we were able to drill the last lateral pair in a single run from each rig.

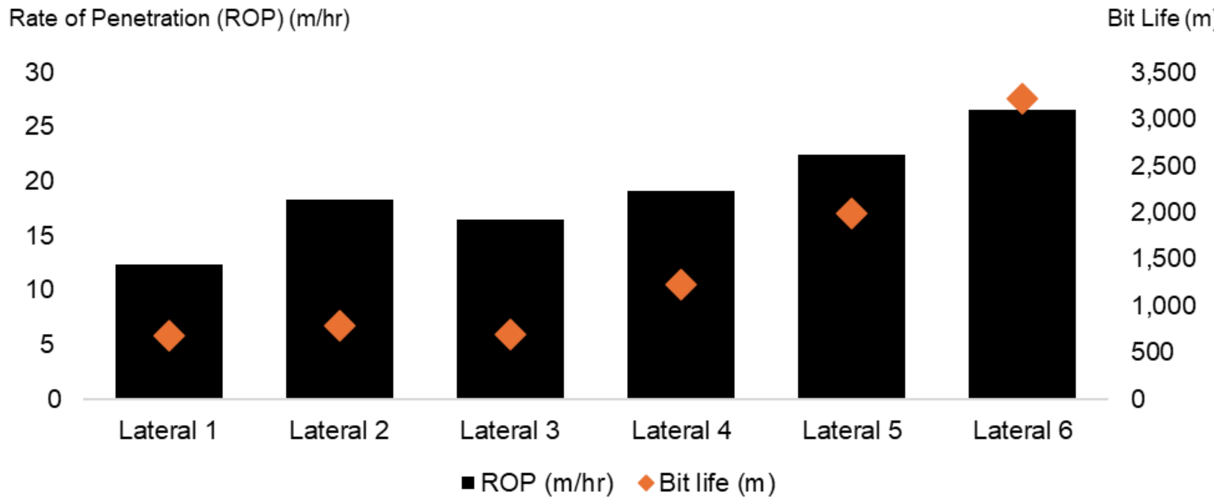

Intersection workflow

Connecting two long horizontal wells at a measured depth of 8,000 meters from surface is not a trivial exercise, especially with the positional uncertainty that comes with extended-reach drilling. Many commentators and industry insiders claimed it was impossible to do it repeatedly and with certainty. However, we intersected all six lateral pairs according to plan. Early in the campaign, we used wireline-conveyed magnetic ranging technology to guide the wells into each other. That workflow worked, but it carried time and complexity and was very expensive. As the team gained confidence, the number of ranging shots dropped significantly. We then moved further down the learning curve with Eavor-Link™ AMR, which brings active magnetic ranging into the drilling assembly and allows ranging-while-drilling rather than stopping for wireline runs. The result was a step-change reduction in ranging and intersection time ("ranging durations") while removing the cost and complexity of a wireline system, while still achieving first-attempt intersections on all six lateral pairs.

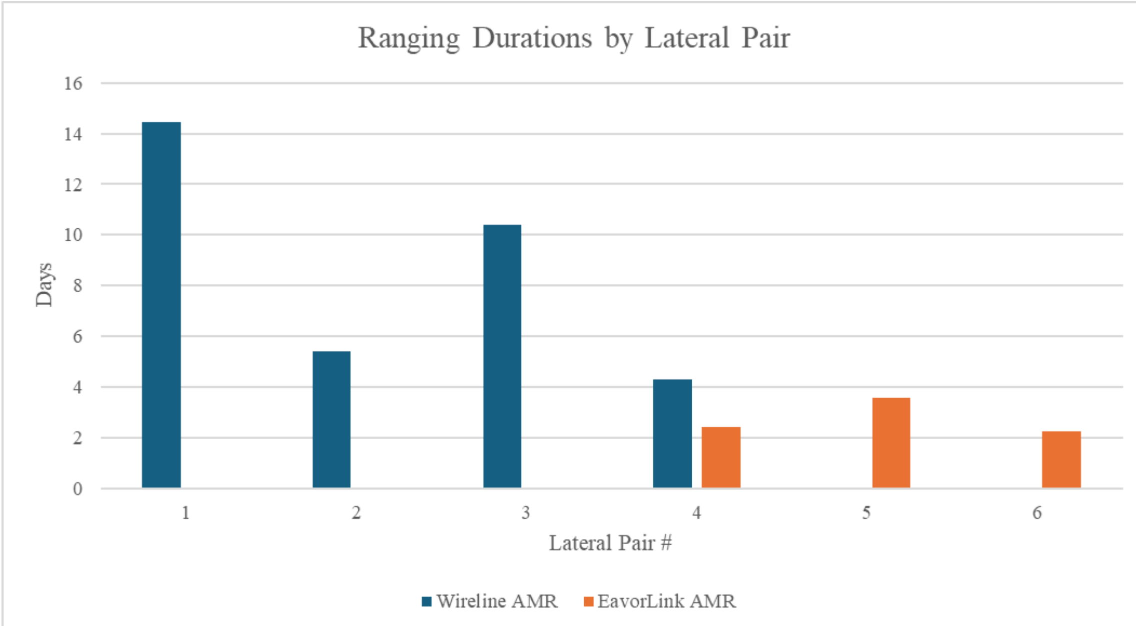

Hydraulic communication between the two rigs through previously drilled lateral pairs

This challenge was by far the most impactful. Following the sidetracking of the second lateral pair, hydraulic communication (fluid flow) between the two rigs was measured. The root cause was poor cement jobs on both cased "motherbores", enabling a significant annular fluid channel behind the casing, which connected all the lateral pairs to each other. Fluid could freely flow down previously drilled laterals and up to the other rig.

Efficient Eavor-Loop construction depends on two rigs drilling in parallel; each lateral pair is drilled by two rigs simultaneously and intersected, starting with the lowermost lateral and then repeating from the "bottom up". Each previously drilled lateral pair is supposed to be sealed (hydraulically isolated) from the next lateral pair which is drilled uphole (in the direction towards the surface within the cased motherbores). However, inadequate isolation in the cement behind the casing created hydraulic communication between the two well systems/rigs. Once that happened, pressure, mud, and cuttings could move freely between the inlet and outlet patterns through previously drilled lateral pairs.

That changed the entire execution model. Instead of drilling with two rigs simultaneously (as was done on the first lateral pair), we had to proceed with one rig at a time. This issue alone effectively doubled time and cost and also allowed cuttings-laden drilling fluid to migrate into previously drilled lateral pairs. Because of the resulting impact on drilling duration and project cost, we elected to stop Loop #1 after six drilled and intersected lateral pairs rather than continue to the original 12-pair design.

The hydraulic communication issue allowed cuttings-laden drilling fluid to backflow into previously drilled laterals, depositing cuttings in the horizontal boreholes. After we finished drilling all six lateral pairs we performed a series of flow tests combined with temperature and production logging tools. Almost all the lateral pairs (five of six) had cuttings/mud deposited except the most recently drilled pair. Based on several high certainty technical indicators, including the material recovered at surface after flushing the legs and the quality of the boreholes as measured by wellbore imaging logs, we are confident the root cause was backflow of mud/cuttings through the cement hydraulic pathway. We were able to flush three of those five compromised pairs, confirmed by production logging. So four lateral pairs are operating as expected today. However, we made the decision not to clean out the first two pairs drilled as the cost and time involved were prohibitive.

This is the most important challenge to understand and was the largest persistent project-impacting event. The root cause was unsuitable design and execution of the cementing operation (cementing between the steel casing and the wellbore) for the complexity of this application.

And the solution is not new technology but applying well-understood cementing design and industry best practices suited to this multi-lateral well architecture. Those changes have been incorporated into all future Loop designs. This critical learning will manifest itself in an inter-Loop learning curve.

Production performance matches expectations

Thermal Production: the Eavor-Loop

In closed-loop geothermal systems, like an Eavor-Loop, the heat transfer is dominated by heat conduction. Per Fourier's law, the heat transfer to the wellbores is proportional to the surface area available for heat transfer (see "A" in the equation below). As we have four lateral pairs contributing to flow, we would expect roughly 4/12 = 33% of the heat transfer and production relative to the original 12 lateral pair design.

$$\dot{Q} = \frac{\kappa \cdot A \cdot \Delta T}{L}$$

Equation 1: Simple 1D Fourier's Law

The production from Loop #1 is currently around 8.5 MWth which is exactly as expected. Recall the original nameplate capacity of the project was 64 MWth for four Eavor-Loops, each with 12 lateral pairs (48 lateral pairs in total), which equates to 1.3 MWth per lateral pair. We are currently producing over 2 MWth per lateral pair, but this will decline to 1.3 MWth after five years, which was the basis of the nameplate capacity.

The production from the flowing lateral pairs is consistent with the connected reservoir exposure - the Eavor-Loop works as intended. This is an important result because we predicted the thermal output prior to drilling (prior to spending any capital) and the output is as predicted. Loop #1 is underbuilt, but the production performance matches expectations for its size.

Electricity Production: the heat engine

For Geretsried, we built an ORC (Organic Rankine Cycle) power plant (supplied by our partner Turboden) optimized for a design point at about ¾ of full project thermal capacity (to account for the heat that will be going to the district heating network and bypassing the ORC plant). In theory, ORC efficiency depends only on the temperature of your heat source (water from the Eavor-Loop) and your heat sink (ambient air). In real life, however, these plants have significant efficiency losses when they are operated below their design point, even if there is no change to the temperature of your heat source or sink. With any ORC power plant, there is a “sweet spot” at the design point where optimal conversion efficiencies can be achieved. If we’re below this design point, conversion efficiencies will be worse. The figure below illustrates this relationship for the Geretsried ORC at “average” ambient temperatures.

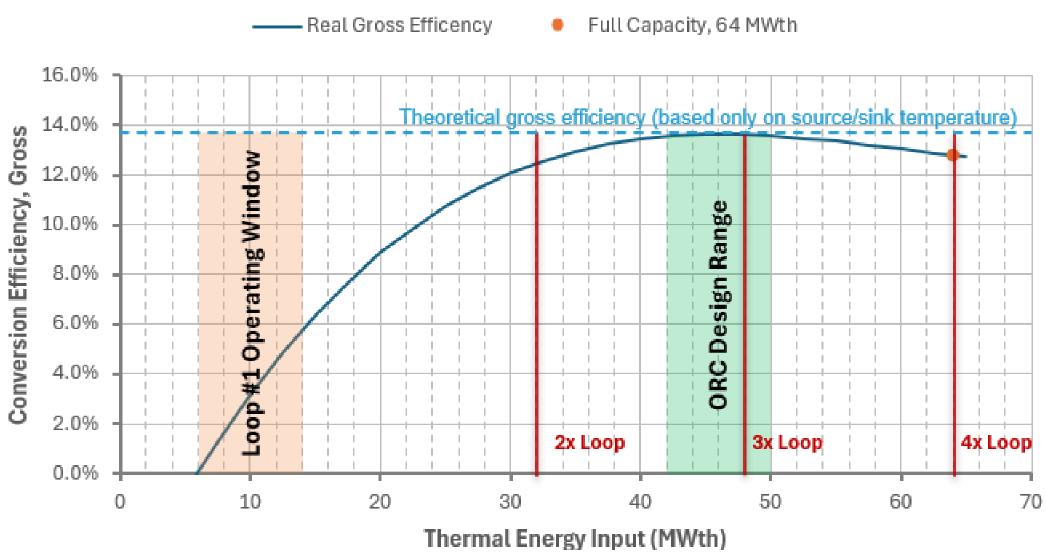

Figure 1: Illustrative Geretsried ORC Turndown Relationship.

Our current Loop #1 operations fall on the far left-hand side of this curve because the power plant is oversized for only one loop. Over the average thermal duty since startup – 11.8 MWth – the conversion efficiency is roughly 4.3%. Of course, in the winter it’s higher and in the summer it’s even lower. Although that’s a relatively small number, it’s actually slightly better than expected for the size of Loop #1 and shows the ORC will operate as intended when thermal production increases with more loops.

Leak-off

Rock-Pipe was successfully applied to all drilled lateral pairs and the technology is proven to work as intended. Field tests during execution were completed (similar to reservoir pressure transient well testing) to confirm the sealing performance of Rock-Pipe. Rock-Pipe applications were able to reduce the near-wellbore permeability by over an order of magnitude. This has yielded an operating leak-off rate that has ranged between 0.5-2% of circulating throughput, which is aligned with anticipated project performance.

As we look forward to Eavor-Loops drilled into deeper, basement rocks with significantly less matrix permeability, operating leak-off rates for the next generation Eavor-Loops will be approximately ~0.1% of throughput, another order of magnitude lower than what we are observing at Geretsried.

What is the cost of energy now and where is it going?

Now let's briefly look at the cost of energy and our line of sight to future reductions.

The key metric for energy generation technologies is the energy cost, typically referred to as the levelized-cost-of-energy (LCOE), where energy can be in heat or power units. LCOE accounts for capital expenditures, OPEX, the lifetime of the asset, and the cost of capital.

$$Levelized\ Cost\ of\ Energy = \frac{Discounted\ (CAPEX + OPEX)}{Discounted\ Energy\ Sales}$$

Equation 2: Levelized cost of electricity is the ratio of project cost to energy sales (discounted with time at cost of capital).

What is the cost of energy now?

The next loop does not start from zero – it inherits all the improvements we learned on Loop #1. That is how every learning curve works. The first units carry the cost of encountering and solving challenges. Later units show what the process can become once the major lessons have been absorbed. That has been true in shale, wind, solar, batteries, manufacturing of all sorts, and most industrial technologies that became cheap through repetition.

Without the hydraulic communication issue that forced sequential drilling, and applying all the lessons learned from Loop #1, we expect future completed lateral pairs can be delivered in approximately three weeks per pair, with significant further improvement still possible. This translates to an LCOH that is cost competitive for district heating in average gradients (30 C/km), relative to other energy systems or alternatives in Europe. (As Eavor and our partners are in commercial discussions on many projects in Europe, we won't be disclosing exact numbers). In addition, it also has the critical attributes of local energy security, 100% carbon free, no fracking, low water use, and long-term predictability.

It is critical to remember that this is all based on a globally average geothermal gradient, not a rare hotspot as can sometimes be found in the Western US, Iceland, etc.

Where is it going?

There are four fundamental ways to reduce LCOE, see equation 2: reduce capex, increase energy sales, reduce OPEX, or reduce cost of capital. When we put all of this together, we have clear line of sight to <75 $/MWhe (electric), *in average geothermal gradients*, and with a system that we've already demonstrated in the field at meaningful scale.

Capex

For an Eavor-Loop, capex is dominated by drilling performance. We know how to construct an Eavor-Loop and have demonstrated the beginnings of a steep learning curve. In sedimentary rock we have already drilled a lateral pair from start to finish in a single run, including ranging and intersection. But there remains a lot of room to improve. Some examples of the key input parameters that drive drilling cost are described below. "NOAK" is an industry term for "nth of a kind", meaning where the learning curve can progress after many Eavor-Loops.

Parameter	Current	Potential NOAK	Evidence
Rate of Penetration (ROP)	• 25 m/hr demonstrated in Geretsried (sedimentary)	• >40 m/hr in sedimentary • 30 m/hr in basement rock	• Oil and Gas wells in NA demonstrate >40 m/hr • Recent wells drilled in hard rock (ex: Colorado, Utah, New Mexico, etc.) show ROP of 30 m/hr is achievable
Bit run length (m)	• 1-2 runs per lateral (sedimentary) • <600 m run length assumed in basement	• Single run laterals (sedimentary) • 1500 m in basement	• Three single run laterals already demonstrated in Geretsried (sedimentary) on 3 of the last 4 laterals • Recent wells drilled in hard rock (ex: FORGE and Fervo in Utah) already show >800 m run length and it continues to improve
Trip speed	• 350 m/hr	• 800 m/hr	• >750m/hr trip speed average is demonstrated in US pad drilling
Sidetracking time	• 5 days demonstrated in Geretsried	• 2 days	• 2 days demonstrated in vertical well by Eavor in hard rock
Ranging/ intersection time	• 2 days demonstrated at Geretsried	• 1 day	• Clear line of sight to reducing ranging time with Eavor-Link™ AMR

Energy sales

The key to increasing electricity output is simple: drill deeper and hotter, which typically means drilling into basement rocks. Drilling deeper is fundamentally the biggest lever to ultimately reach <75 $/MWh “geothermal anywhere”, without having to locate in rare geothermal hotspots.

When the Eavor-Loop system is drilled in crystalline basement rocks, we call this “Eavor-Loop 2.0”. From a drilling technology perspective, it is the same as at Geretsried but drilled into harder granitic or basement rocks. Eavor and our development partners are actively progressing projects for both district heating and for power markets using the Eavor-Loop 2.0 configuration.

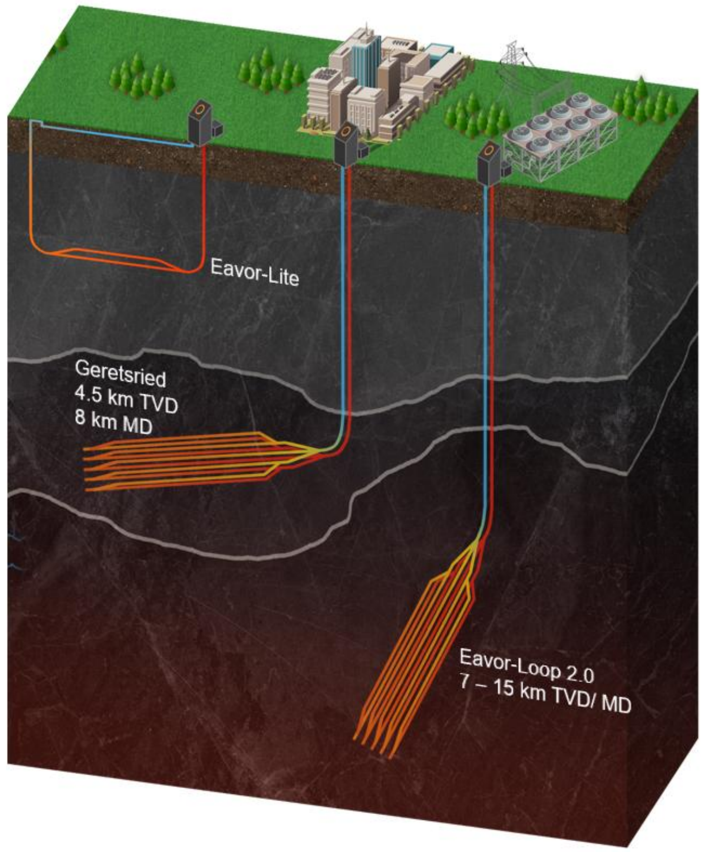

Depth is a continuum. Our first loop was already 8 km Measured Depth, and subsequent loops at Geretsried will be incrementally deeper/ hotter/ longer. The same incremental evolution in depth will apply to future Eavor-Loops. As we continue to push depth, output increases exponentially. For example, in a 30 C/km gradient, output of a loop at 15 km depth is 20 MW electric, compared to 2 MWe at Geretsried.

Reaching depths below about 9 km TVD, current drilling technology faces a limit. To drill substantially deeper, Eavor and our partners are building a drilling system that uses existing technology but is optimized to construct Eavor-Loops at these depths. We are building this system as we speak, and will have a working system drilling in the field in a few years.

This depth progression is made over a gradual evolution of the technology rather than an immediate step change. The world has witnessed this dynamic before with wind. Once the initial challenges are overcome and the system is proven at meaningful scale, size was increased.

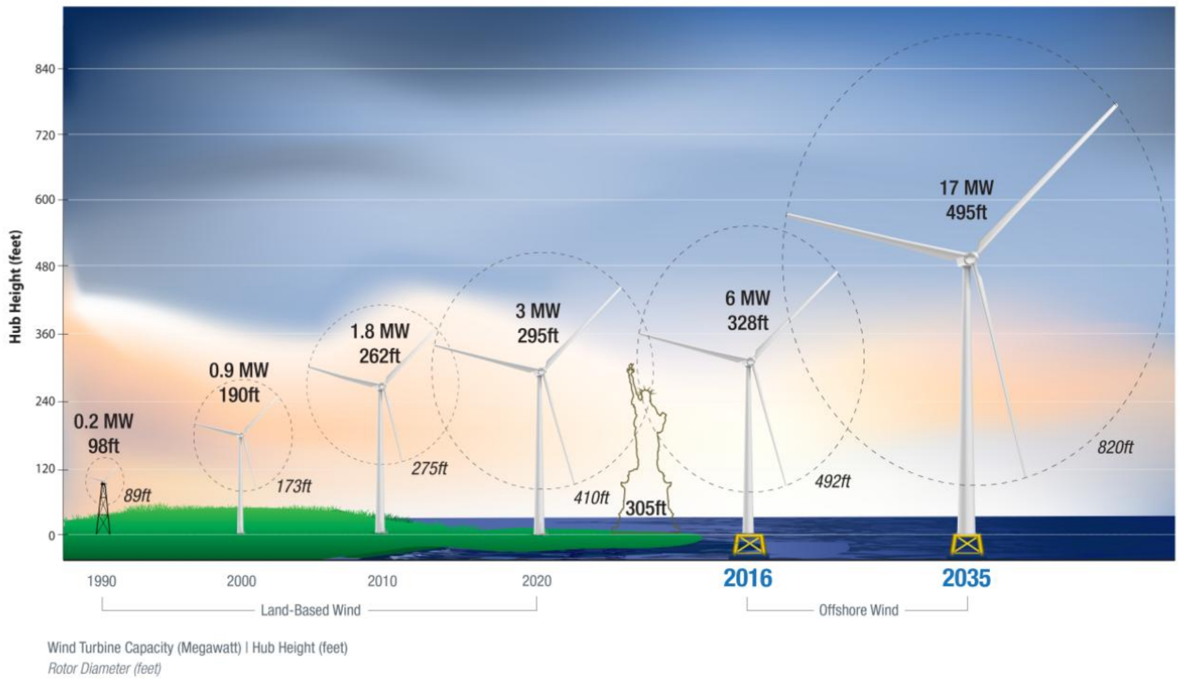

U.S. Department of Energy. (2024, August 21). *Illustration of increasing turbine heights and blades lengths over time* [Infographic]. Wind Turbines: the Bigger, the Better.

OPEX

Due to the long-term nature of the loop and the fact there is no reservoir and no moving pumps, many of the top budget items traditionally required in geothermal are eliminated. There is no sustaining capex, no pump installation or replacement, no pumping parasitic load, no scaling/reservoir plugging issues, etc. So Eavor-Loops have significantly less OPEX than many other technologies, at about 5% of LCOE, which is close to wind/solar in that regard.

The data below compares the contribution to LCOE from capex and opex of various firm technologies. The implications are that; a) capital efficiency is a poor metric to compare technologies for a given project, as the end user actually cares about energy price not capital efficiency; b) there is not much room to reduce LCOE by reducing Eavor-Loop opex; and c) there are effects for cost of capital and project structure.

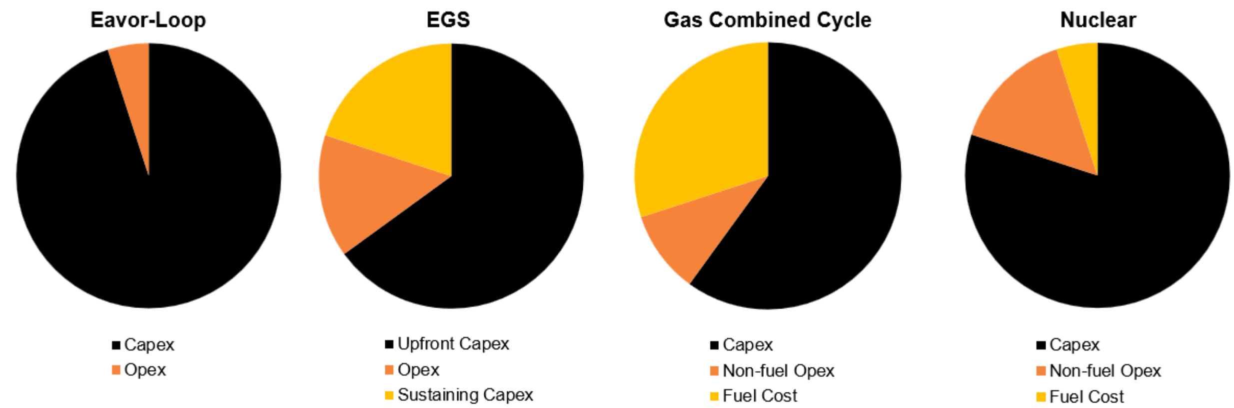

Contribution of opex and capex to LCOE for various firm technologies. Sources: Eavor - internal estimates based on Geretsried data, EGS - based on Fervo's S1 filings, Gas – Lazard LCOE+ 2025, Nuclear - OECD Nuclear Energy Agency.

Cost of Capital

The final impact on LCOE is cost of capital. This has a large impact on LCOE, especially for systems with high percentage of LCOE as CAPEX such as nuclear or Eavor-Loop. Initially, cost of capital is high for all emerging technologies. As projects are built and the execution becomes more standardized and predictable, cost of capital comes down.

For Eavor-Loops almost all of the risk is in the construction stage. Once the loop is built it is very predictable over long time frames (30 -100 years), with low maintenance. Resource risk is largely eliminated, especially after the first loop is built. Therefore, projects will typically be refinanced to a high debt ratio after Commercial Operation Date (COD). Project structure after COD is most closely aligned with mature energy technologies which are owned and operated by utilities or by infrastructure investors.

Can Eavor scale globally?

The last question we typically get asked is how can the company scale projects across multiple continents and in different jurisdictions globally? More on this to come, but let's explore briefly.

Eavor's strategy has always been to shift towards a technology licensing model over time, ever since the first pitch deck in 2018. Our business model now consists of several elements:

- License technology
- Provide Eavor-Loop-specific services:
    - Engineering and technical services (we've developed specialized know-how and software)
    - We directly provide certain specialized well and construction services after FID
    - We provide certain operating services
- Eavor will continue to have some "skin-in-the-game" in certain projects, the most obvious is Geretsried where we are the developer.

The advantages of this model are that it is capital-light, higher margin, and we bring the capabilities of world-class partners (both operating and financial) to scale globally. The primary drawback (for Eavor) is a loss of control. This model also requires us to retain the position at the center of an ecosystem, and thus includes barriers to entry such as:

- A web of patents at both the geothermal system level and at the component level
- A series of exclusive commercial agreements with suppliers, technology partners, and developers
- Expertise, know-how, and data
- We gain more data, experience, technology advancements, and partners over time. The product therefore becomes more valuable as more users adopt it.

Future business model at Geretsried

Eavor remains fully committed to the Geretsried project along with our partners. It is a strategic project that supports the deployment of our technology into the wider European market and drives cost compression that can be applied to projects globally. As the project developer, we are engaging with suppliers, partners, investors, and heat offtakers to improve the project and are excited about the path forward. We will share additional updates on the project as they become available.

Conclusion

At Geretsried we proved that Eavor-Loop technology works. We drilled and connected commercial-scale laterals at depths greater than 8 km MD, generated thermal production consistent with predictions, produced electricity as expected, and demonstrated the beginning of the learning curve that underpins the future economics of the technology.

The project also forced us to solve hard problems in the field. Some of those challenges were expected in a FOAK project. Others were not. But this is how industrial technologies mature: by encountering constraints at scale, solving them, and systematically improving performance through repetition. That is exactly what happened at Geretsried.

Returning to the three key questions noted earlier:

Does the technology work?

Yes. The Geretsried project has accomplished the intended objectives.

What is the cost of energy today and where is it going?

The economics are already competitive in European district heating and certain power markets today. More importantly, we now have a clear technical and operational path toward substantial future cost reduction through learning curve improvements and deeper, hotter loops.

Can Eavor scale globally?

Yes. The technology is specifically designed around scalable industrial processes and existing global supply chains. Combined with a licensing-based business model and strong strategic partners, this creates a realistic path to deployment across multiple continents and energy markets.

Eavor-Loop can satisfy what the market really wants: energy systems that are simultaneously secure, dispatchable, scalable, locally sourced, and low carbon.

Geretsried was the first proof of that thesis but is not the final destination. The next phase is about applying what we learned, improving performance through repetition, drilling deeper and hotter systems, and scaling the technology globally. This is only the beginning.