imec xpand

Bridging Vision and Early Revenue – MVP and Beachhead Strategy in Semiconductor Startups

April 2026

The journey from breakthrough technology to a scalable semiconductor company begins with two tightly linked strategic choices.

The first is defining the **Minimum Viable Product (MVP)** [1] in combination with the **beachhead market**, a specific market segment that a company dominates first to establish a secure, profitable foothold before expanding into larger, adjacent markets.

The second is marrying the MVP and paired beachhead market with a clear **strategic vision**.

These choices must be made early, deliberately, and with discipline. The startup executive team faces tough and often paradoxical expectations, as investors simultaneously demand *fast time to revenue* and a *compelling long-term vision* with potential for massive market impact.

It is precisely this tension between immediate deliverables and long-term ambition that makes the MVP/ beachhead strategy so important for a semiconductor startup.

Investor Expectations: Fast Revenue and Huge Markets

Investors are pulled in two strong but opposite directions. On one hand, they want companies that can **generate revenue quickly**. On the other hand, they expect founders to articulate an ambitious vision that leads to **large market opportunities**.

This contradiction means a founder must simultaneously:

- demonstrate a **credible path to Technology Readiness Level (TRL) 9** [2] **and revenue** within a few years, and
- paint a vision of a **big, inspiring, long-term opportunity**.

The best way to satisfy both requirements is through a well-chosen MVP *paired with* a beachhead market and application that enables early market access and product revenue, while clearly mapping how the MVP evolves toward the broader vision. The vision paints an exciting picture of what the company could achieve in the years after the MVP launch; this should represent multibillion-dollar markets, with the potential to be dominated by the startup's follow-on products.

In the next sections, we will explain the key elements of this "MVP + beachhead market and application" combination and how to address them.

The company should paint a credible picture of how to get from the MVP to the vision in steps, with each step representing a real product with a real market. We refer to this as the **"MVP-to-vision strategy."**

The picture below, although somewhat of an oversimplification, represents this graphically for a long-term vision of selling a fully specified car. Rows 1 and 2 represent distinct approaches, and it is crucial to choose your approach deliberately.

The first row represents incomplete technologies (steps 1-3), with only the final step representing a sellable MVP. With this approach, the startup will only create real product revenue when it reaches step 4. However, at some point along this trajectory, investors and customers may lose patience. Moreover, none of the intermediate steps will ever validate market traction, which is key feedback and necessary to convince investors that it is worth continuing to invest in the product development.

The second row represents a different approach, in which each technology point represents a sellable product, and each product is an upgrade of the previous product in credible, explainable steps. Through this approach, the company can start creating revenue with product 1, with corresponding customer validation.

The first approach is a valid option, if looked at differently. The startup could begin selling wheels (as a component) to car manufacturers as a first credible MVP step, and then extend the MVP (selling the complete chassis), and so on.

This approach offers the advantage of starting small and moving up the value chain; the downside is that the customer might change as the company moves up the chain. Furthermore, the company misses the opportunity for early feedback geared toward the final product.

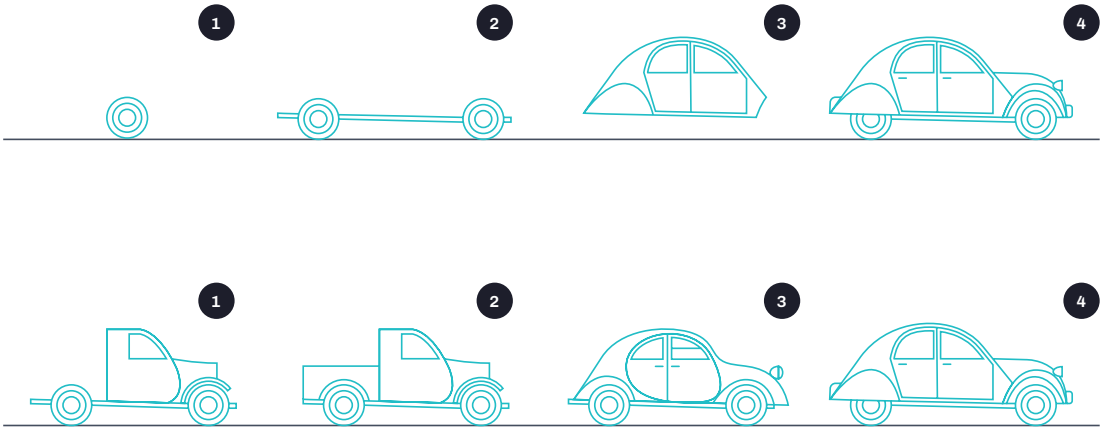

Figure 1 - MVP-to-vision strategy

It is very important to note that approximately 95% of a company's resources should be focused on the MVP/beachhead definition, validation, and development (in Figure 1, product 1 in the second row). The organization's entire goal is to develop and launch the product and reach product revenue as quickly as possible. The CEO and CTO, with support from the CFO and strategic marketing, are responsible for painting the vision and strategy to investors, strategic customers, and the board.

The MVP as a Real Product, not a Prototype

Founders often fail by confusing early technology demonstrations with products. An MVP in a semiconductor business is a **real, sellable, scalable product**, not a research prototype. Since investors have limited patience waiting for evidence of manufacturability, reliability, and market acceptance, semiconductor startups are expected to reach TRL9 with the MVP within four years.

This is why defining the MVP is not a brainstorming exercise to test assumptions but a **continuous process of assumption validation**, driven by a strong product manager/owner embedded from day one.

The MVP must have:

- product specifications that are validated with customers on the beachhead application;
- reliable and robust behavior that covers all the customer cases for which the product is intended;
- support documentation;
- a manufacturable supply path that can cover production for at least the coming two to three years.

MVP + Beachhead Market: A Single, Well-Defined Combination

Semiconductor startups operate under intense resource constraints. For this reason, the MVP and beachhead market must be defined as **one tightly integrated pair**. The startup team must focus all early sales, marketing, and engineering efforts on **one specific application in one segment**. Furthermore, the beachhead should have adjacent markets that the startup can enter after having conquered the beachhead.

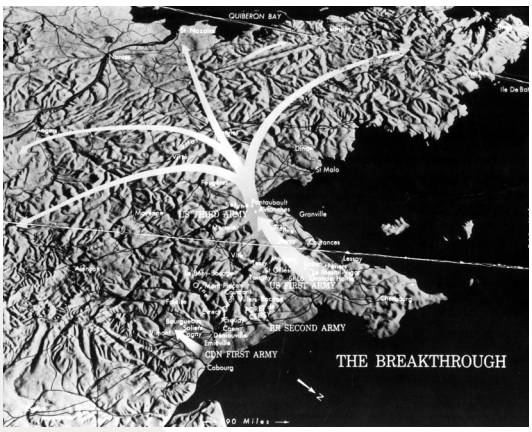

Figure 2 - Invasion routes following the Normandy landings

The term *beachhead market* comes from military strategy, vividly illustrated by the Allied landings in Normandy on D-Day. In that operation, forces first secured a narrow strip of coastline - a “beachhead” - that served as a defensible foothold in hostile territory. From this initial position, they reinforced troops, stabilized supply lines, and then advanced inland to liberate larger regions.

In business and innovation strategy, a beachhead market follows the same logic: Instead of trying to capture an entire market at once, a company focuses on a small, well-defined segment where it can establish dominance and build momentum. Once firmly positioned in that initial niche, it leverages credibility, resources, and customer traction to expand into adjacent markets, just as the Allied forces expanded from the beaches of Normandy into mainland Europe.

There are several reasons why this focus is essential. First, developing a semiconductor product is an extremely complicated endeavor. It is technologically challenging, has a complex supply chain, and requires strong partnerships and demanding customers.

Furthermore, semiconductor markets are unforgiving. A design-in is a major investment for the customer. They may already have spent years integrating an inferior solution, and

switching is risky. Making the “wrong choice” can be career-limiting for an executive.

Finally, competitive incumbents will actively discourage prospective customers: “*Why work with a startup?*” and “*We’ll launch something similar in 18 months.*”

By dominating a narrow segment in which requirements and behavior are homogeneous, a startup can have a reasonable chance to overcome these frictions.

Once the company proves the MVP in a focused beachhead market, its cost of sales declines, customer trust increases, word of mouth takes hold, and the company gains bargaining power — all factors that accelerate growth.

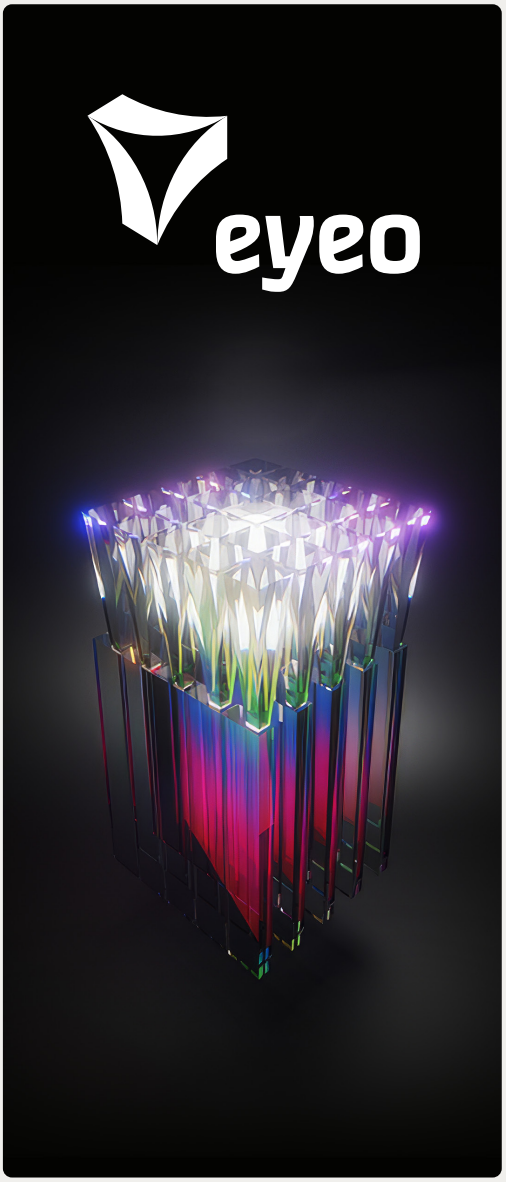

EXAMPLE

eyeo, a venture coming out of research center imec, is developing a new type of imager with vertical nanophotonic lightguides and color splitters that are disruptive for many imager applications.

The dream is to bring this imager to a level that it can replace the main imager(s) in the smartphone. It is a huge, existing market.

However, it is a very difficult market for many reasons:

- The functional requirements for a smartphone imager are enormous.
- The supply chain requirements for high-volume smartphones are very tough.
- Aligning with the smartphone ecosystem is complex from a relational and technical perspective.

Therefore, it is smarter to target a smaller and more benevolent application area. One in which the USPs of the innovation still shine but market access, supply chain, and technical and relational challenges are much easier.

So, the company takes a more gradual approach: first, going after the security camera market; next, aiming for the consumer camera market; and later, targeting the smartphone camera market.

The MVP–Vision Table: A Clear Strategic Separation

The five-point MVP and vision checklist below shows the difference between the MVP and the long-term vision. This table separates the product initially being built from the company's targets for the future.

The left side of the table details what the **MVP** must be: a tightly defined product that solves a concrete customer problem while still reducing risk to enable a fast time-to-market. The right side describes the vision: a far broader scope in line with societal trends, potentially spanning more and bigger markets. The vision allows for dreaming, where it is not clear how all future technological challenges will be solved. This side is less focused on customers but should appeal to investors.

Founders are often lured into trying to build the vision first. That leads to overcomplexity and missed timelines with devastating consequences for the startup.

The MVP should be selected not because it is the "ultimate product" but because it is the fastest and most certain path toward proving the company's value. Every (new) feature in the MVP represents a risk: it may malfunction, delay development time, or introduce unexpected complications. Thus, a company should very consciously decide whether it should take the risk of implementing any new feature (in the MVP) or avoiding it altogether by restricting the beachhead market and application. Avoiding it does not mean that the feature will never be implemented, it is just postponed to a later, more mature version of the MVP.

While it is crucial not to prioritize the vision over the MVP, failing to articulate the vision leaves investors unconvinced of long-term potential. Therefore, combining the MVP and vision consistently is the way to go.

	MVP	Vision
The promise	Solve a short-term customer problem	Paint a vision that resonates with social trends
Target group	Appeal to customers	Appeal to investors
USP	Convincing benefits at customer application level	Maximize potential and ROI
Focus	Fast time to market and risk reduction	Longer term
Language	Engineering mode & factual	Dreaming & emotions

What Makes the MVP Strong Enough to Dominate the Beachhead

We have identified **five defining properties** of an MVP that is capable of dominating the beachhead market. These properties ensure that the MVP is not merely a functional product but also a commercially viable one.

A true customer-centered 10x improvement

This is not a technological 10×, but a 10× in a functionality the customer cares about most. Most often, these are not specific analytical properties or capabilities of the product, but rather the true value the product offers at the application level. When looking for this customer-centric 10×, look out for customers expressing strong, urgent pain points, and brainstorm ways an MVP could resolve them decisively. There is an important nuance: Founders must avoid “false positives”, enthusiastic people in large organizations who like the technology but are not decision-makers in buying and integrating the product. Real validation requires commitment from buyers who actually carry adoption risk.

A defendable moat

This ensures that once the startup wins, competitors cannot easily copy or replace the product. The moat may come from IP, system architecture, manufacturing know-how, or device integration.

Low technological risk

The MVP must increase the predictability of development. Founders are warned that engineers tend to add bells and whistles that create risk and delay. If complexity creeps in, it is often better to narrow the beachhead application rather than expand the MVP functionality.

Scalable production within two to three years

Since scaling from low to mid volume takes at least 18 months, the MVP must allow the company to serve early demand without having to redesign the product or its supply chain.

Early revenue potential

The MVP should be able to generate about **€10M/year** in revenue in the beachhead market, within one to two years of launch. Of course, this is just a

ballpark figure. The key is that the MVP is a real, recurring, revenue-generating product instead of non-recurring engineering revenue.

A good beachhead is often in the range of **€100-300M** — not massive but substantial. If the market is larger, the startup cannot dominate it early; if it is smaller, the opportunity is insufficient for investors.

We note that financial VCs more often than not balk at such small beachhead markets. However, we find this concept essential. The startup must pick an "entry battle" it can surely win. To use the words of Sun Tzu [4]: "He will win who knows when to fight and when not to fight." Translated to startup language: You only ought to commit to a fight for a beachhead market if you can control the "when" (when you have a clear advantage) and "how" (when the terrain and tactics allow you to be victorious).

MVP Definition and Validation

Developing a semiconductor product takes years. As such, the product validation techniques of Eric Ries' "The Lean Startup" cannot be applied to most cases of semiconductor development. Fast, two-week iterations based on software prototypes that are continuously improved through customer feedback are not possible here. Thus, there is an enormous danger of leaving customer assumptions unvalidated. However, every unvalidated assumption is a massive risk multiplier.

All company creativity should go into **validating assumptions with prospective customers as early as possible.**

Companies should:

- Document every detail of the product definition and application use cases (that make up the beachhead market).
- Validate these detailed documents with prospective customers.
- Ensure that the product's customer use and benefits are well understood and validated with the customer.
- Establish a documented system-level view, including components delivered by the customers and/or partners. The component is often only a part of a full solution, so even if the startup's part meets all specs, it will not sell if the full solution does not also meet those specs.
- Use creativity to validate assumptions the best way possible when a full prototype is not available:
  - Engage with prospective customers using simulation models, emulator models, and eventually prototypes

to validate product behavior and design-in criteria;
  - Use mockups to test specific parts of the user experience;
  - Use proofs of concept to validate part of the benefits;
  - Organize customer interviews, including detailed spec sharing and feedback.

- Exercise caution regarding evaluation kits (EVKs). Although they can be useful tools, they do not usually validate product constraints and assumptions. Companies often buy an EVK out of curiosity and to follow technology trends, so the number of EVKs sold does not necessarily reflect customer interest and assumption validation.

- Build an early milestone into the road map that, if an important aspect of the product cannot be validated, allows that particular aspect/assumption to be validated in a better way. It is better to find out earlier - rather than at product launch - that a wrong assumption was made.

Validation shouldn't be a one-off exercise. Markets evolve while a chip is still in development, and sometimes, assumptions made at the start of development may lose validity after a year or two. Therefore, validation must be **continuous**.

The bottom line is that if a go-to-market fails because of a wrong assumption that could have been validated early, the company may fail where it could have succeeded by catching that incorrect assumption early. This is a potentially company-destroying mistake.

Why MVP Launches Fail

Based on our experience, these are the three major causes of MVP launch failure along a timeline that spans technology ▸ product ▸ commercialization ▸ scaling:

- **The MVP is not minimal**
  Too many features, too much customization, or chasing too many markets simultaneously.

- **The MVP is not a product**
  Teams confuse prototypes with products; reliability, documentation, and manufacturability are insufficient or even absent.

- **The MVP is not viable**
  The company has not validated customer needs, has misunderstood willingness to pay, or cannot support integration. Alternatively, the product is not unique enough to dominate the beachhead market, as there are enough other viable options for the customer.

All three failure modes push out the timeline, delaying revenue and destabilizing the company. Avoiding these failures requires tight scoping, rigorous customer engagement, and continuous evidence gathering.

CASE

Ultrasound Patches: MVP vs. Vision

pulsify medical

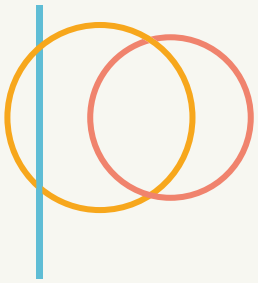

This example offers a concrete illustration of how a semiconductor startup should define its MVP, beachhead market, and long-term vision.

The company

Pulsify Medical, a venture coming out of nanoelectronics research center imec, incorporated in 2019, developed an ultrasound patch for monitoring human organs in three dimensions and deriving clinical parameters from that model automatically and in real time.

This capability enables a range of new possibilities since it allows for noninvasive and safe continuous monitoring of important clinical parameters. The patch technology could be produced in high volumes and is low enough in cost to become disposable, which is important, as the lifespan of the ultrasound transducers is limited.

The patch opened up multiple applications in many large markets: cardiac and vascular, liver, lung fluid, kidney, and bladder monitoring, as well as intensive care unit (ICU), hospital, and home monitoring settings. Even deployment for consumer applications in sports could be considered. With this technology, the company was able to dominate the new and wide-open ultrasound monitoring space.

Here is how the company went about defining the MVP and the vision:

	MVP	Vision
The promise	No electronics integrated Companion device	Electronics integrated Standalone device
Target group	Immobile patient	Mobile healthy person
USP	Integrated in ICU - ECG	Standalone
Focus	Cardiac output	Human organ monitoring

The MVP & beachhead market

In the MVP, the ultrasound patch was conceived as a **wired device intended for use in ICUs**. It did not contain integrated electronics; instead, it relied on a **companion device** to perform signal processing. This reduced the system complexity dramatically: The patch itself became a simpler transducer, avoiding premature challenges such as on-chip power management, wireless communication, and local computation. Because the device was wired, it did not require batteries, power-related safety certification, or ultra-low-power operation. All of this focused on risk reduction while still targeting a very attractive beachhead market.

The MVP was also designed for **immobile patients**, which further simplified the product. ICU patients are typically stationary, monitored continuously, and connected to multiple existing systems (ECG, ventilators, infusion pumps). This means the MVP could be physically tethered without compromising patient experience or clinical workflow. It also enabled the patch to use the ICU's **existing infrastructure**, making integration with hospital monitoring systems straightforward and reducing regulatory burden.

Most importantly, the MVP focused on a **single established clinical parameter: left ventricular volumes**. In ICU settings, this is extremely valuable for managing critical patients, and existing monitoring solutions can be invasive, cumbersome, or insufficiently continuous. In this particular case, the focus was on a specific application, namely, sepsis patients in the US. For such patients, invasive methods must be avoided, while minute-by-minute tracking of their cardiac volume is essential for tracking therapeutic effects.

The MVP, therefore, solves a real, high-urgency problem for a customer base that is motivated to adopt improved tools. This is exactly the type of acute customer pain that creates a viable beachhead: a controlled environment, a captive patient population, and clinicians desperate for better information.

The vision

In contrast, the long-term vision imagined a **fully wireless, battery-powered, standalone ultrasound patch** capable of supporting mobile, healthy users and patients outside the hospital. This future version would integrate all electronics into the patch, communicate via Wi-Fi or cellular connection, and perform expanded physiological monitoring.

The user context also shifted radically: Instead of immobile ICU patients, the vision product would **support people running, walking, or going about daily life**. This expansion could open more, larger markets, such as remote health monitoring, fitness, and chronic disease management, but it would also require solving vastly more complex engineering challenges: ultra-low-power design, secure wireless communication, motion compensation algorithms, compact packaging, robust consumer-grade reliability, and full regulatory approval for use outside clinical supervision.

Unfortunately, the fundraising climate in the MedTech industry was incredibly rough, and Pulsify Medical failed to raise sufficient funding to execute its plan. But we fully believe the strategy was set up exactly the right way.

About

Peter
Vanbekbergen
Partner & Co-Founder
imec.xpand

The author

Peter Vanbekbergen is a deeptech investor and technology entrepreneur with extensive experience in semiconductor R&D, venture creation, and early-stage company building.

He has been involved with several deeptech semiconductor startups as a research manager, product development engineer, VP of engineering, CEO, investor, and board member.

As a co-founder of imec.xpand, he focuses on transforming breakthrough research into scalable startups and guiding founders through the complex lifecycle of deeptech commercialization.

imec.xpand

imec.xpand is one of the world's largest independent venture capital funds dedicated to early-stage semiconductor innovation. It targets ambitious startups where the knowledge, expertise and infrastructure of imec, the world-renowned semiconductor and nano-technology R&D center, can play a determining role in their growth.

imec.xpand has an outspoken international mindset towards building disruptive global companies and strongly believes that sufficient funding from the start is key to future success.

References

1 The concept of the MVP was first introduced by Frank Robinson in 2001 and later popularized by Steve Blank and Eric Ries. While Ries' primary goal of the MVP was fast, experimentally validated learning and hypothesis testing, we retain a stricter definition of the MVP as being the first selling product, catering to the least demanding customers, such as early adopters.

2 Mankins, John C. "Technology Readiness Levels: A White Paper." NASA, Office of Space Access and Technology, 1995.

3 Schoonhoven, C. B., Eisenhardt, K. M., & Lyman, K. "Speeding Products to Market: Waiting Time to First Product Introduction in New Firms." Administrative Science Quarterly, 177-207, 1990.

4 "The Art of War," Sun Tzu, original dating from approx. 400 BCE.

imec xpand