


The Capital Behind the Clock: Applied Materials' EPIC Center and the Economics of Semiconductor R&D

Why an equipment maker is funding a shared research center for competing chipmakers, and the ROI it expects.



OCTOBER 2026

Author:
Stephen Sopko
Practice Lead – Semiconductors
& Deep Tech



1. Executive Summary

The companies building AI hardware now bring a new generation to market roughly every year. The materials those chips are made from move on a far slower schedule, often taking many years from first research to production. That gap shapes the pace of the entire industry, and it is the subject of this paper.


Our first observation is that the slower schedule is not set by science alone. The work of turning a new material into a manufacturable process passes through a chain of separate organizations: university laboratories, equipment makers, chipmakers' integration teams and, finally, the fab. Each operates on its own budget, incentives and definition of success. Time is lost where work crosses from one to the next, through rework, requirements that change meaning in transit, and research aimed at targets that later prove impossible to manufacture. The last of these is the most costly and the least measured. Across four of the major process transitions of the past three decades, the hardest problems sat in the integration work between organizations, not in the discovery of the materials themselves.

Our second observation is that Applied Materials' EPIC (Equipment and Process Innovation and Commercialization) Center, unveiled on October 12th 2026, is designed to address that loss directly. Applied plans to invest \$5 billion, with

spending scaling as customer programs begin. The center brings chipmakers' engineers, Applied's materials scientists and university researchers into joint programs multiple generations ahead of production. It is designed to work alongside the labs Applied keeps near its customers' fabs, not replace them. The partner list suggests the industry is taking the model seriously. TSMC, Samsung, Intel, SK hynix and Micron are founding partners, and they are joined by Broadcom, other equipment makers and leading research universities. The most revealing commitment may not be from a chipmaker at all, we examine why in Section 5. Programs are already defined and underway ahead of full tool installation.

Our third observation is the one we consider most important for investors, EPIC is less a building project than a change in where Applied competes. We will develop this in detail in Section 7.

The model carries real execution questions, including confidentiality among competing partners, the concentration of work in one region and the time it takes to measure results. We examine each in Section 6. Applied has navigated comparable organizational challenges before; we find none of the current risks likely to be insurmountable, though they will require careful management and oversight. The sections that follow set out the history, the organizational problem, the partner base, the capital case and the signposts to watch.



2. New Materials are Important; High-Volume Yield is Critical

Most accounts of the AI trade start at the accelerator, or one step earlier at the foundry. That is usually where the capital is visible and where the product announcements land. We think the story starts further back, at a question most investors never have to ask: what does the chip have to be made of, and does anyone yet know how to make it at volume?

The answer is rarely yes on the day an architecture is conceived. Chip designers can draw structures years before anyone can build them. What closes the distance is a chain of work that runs from university basic research through applied research, into equipment and process development, then integration, and finally the high-volume fab. It is tempting to picture that chain as a sequence of activities. We find it more useful to picture it as a sequence of organizations. The university lab, the equipment maker's research group, the chipmaker's integration team and the fab each sit inside a different institution, with its own budget cycle, its own incentives and its own definition of 'done.' Work moves between them by transfer, and every transfer crosses a boundary. Those boundaries are the subject of this paper.

The industry's own history suggests why. Across four of the transitions that defined the last three decades of scaling, the pattern repeats with regularity that teaches us something. **The architecture was unimaginable. It was not manufacturable until a materials or process capability existed.** And the hard part was rarely the material itself.

Copper wiring is the first case. By the late 1990s, at the 0.25 and 0.18 micron generations, the limit on clock speed had shifted from the transistors to the wires between them, and aluminum's resistance had become the bottleneck. Copper conducted

better and held up better under current. The difficulty was that copper could not be patterned by plasma etching the way aluminum was. The industry therefore rebuilt the back end of the process around a different approach, depositing metal into etched trenches and polishing it flat. Getting that flow to yield as reliably as the aluminum process it replaced required qualification across the full integrated flow rather than step by step.

High-k metal gate followed a decade later, across the 45nm to 28nm generations. The silicon dioxide gate insulator had been thinned to a few atomic layers, and electrons were tunneling through it, burning power in chips that were supposed to be idle. Replacing it with a hafnium-based insulator and a metal gate cut gate leakage by more than a hundredfold. Yet identifying hafnium oxide was the easier problem. The harder one was building an interface less than a nanometer thick, free of defects, with separately tuned gates for the two transistor types. That appears to have required several deposition and treatment steps to run in sequence without the wafer ever touching air, which in practice meant integrating them on a single platform.

Cobalt contacts, adopted across the 10nm and 7nm generations, tell a similar story at a smaller scale. Tungsten had filled the tiny vertical connections to transistors for years, but the liners and seed layers it needed consumed so much of each opening that, a few generations out, the plug would have been nearly all overhead. Cobalt offered a way forward. Delivering it meant a multi-step flow held under vacuum, plus a way to find voids without destroying the wafer, because a single void can kill a die.

The current transition, from FinFET to gate-all-around nanosheet transistors, moves the pattern into the transistor itself. Gate-all-around transistors first entered production at 3nm in 2022 and have since become the leading-edge architecture at 2nm and below. A fin lets the gate control three sides of the channel. As transistors shrank, that was no longer enough, and fins made tall and thin enough to compensate had begun to bend under stress. So the channel was laid on its side and stacked, with the gate wrapping every surface. The consequential change is where the channel's thinnest dimension is set. In a FinFET, that dimension, the width of the fin, was defined by patterning. In a nanosheet, it is the thickness of each sheet, set by growing alternating layers of silicon and silicon-germanium. The silicon-germanium is then selectively removed to free the sheets, and the gate stack must be deposited evenly into gaps of roughly ten nanometers. Sheet width is still patterned. The dimension that governs how well the gate controls the channel has moved from patterning to materials.

Read together, the four cases support a narrow claim, and we want to keep it narrow. In each, the new material was arguably the easier part. The hard part was integration, yield and measurement in high-volume manufacturing. That work sits across organizational lines, because neither the party that invented the material nor the party that runs the fab can finish it alone. It needs both, plus the equipment maker in between, iterating against each other's constraints. Anyone who has

managed a product transition across three companies knows the spec sheet is the easy document. The handoff memo and meetings are the hard lift.

Applied was among the companies working in that gap at each of these transitions, and its integrated platforms and inspection tools were designed for problems of exactly this kind. We do not suggest it closed the gap alone. These were industry efforts involving chipmakers, rival equipment makers and materials suppliers. The more interesting observation is how the applied-research work was organized. Each time, it appears to have been assembled around a specific transition, program by program and customer by customer, and then dispersed once the process moved into production. While the capability was real, the process structure was temporary, based on the demands and parties involved.

That is the conclusion this section earns. Applied has operated at the applied-research layer before, episodically, when a transition required it. EPIC is designed to make that position standing rather than situational, with tools, customer engineers and university researchers working in one place three to four generations ahead of production instead of being assembled anew for each transition. Whether a permanent structure outperforms a series of temporary ones is the question the rest of this paper examines. History says the problem recurs. It says less about whether the answer can be built in advance.





3. The Telephone Game: Eliminating Handoff Friction

If materials research runs through a sequence of organizations, the next question is what happens at the joins. Our view is that the joins cost more in delay or re-direction than the science does, and that most of that cost rarely shows up where anyone is measuring.

The model the industry inherited is a relay. Basic research hands a promising material to applied research, which hands a candidate process to equipment development, which hands a qualified tool and recipe to the chipmaker's integration team, which hands a working flow to the fab. Each runner is fast inside its own lane. Speed is lost at the baton. Takeuchi and Nonaka described the alternative in 1986, studying companies such as Honda and Canon: a rugby line, where overlapping teams move the ball down the field together and phases run concurrently rather than in strict sequence. Four decades on, that insight is standard doctrine in autos, consumer electronics and software. Semiconductor materials R&D largely skipped it, and the reason is mostly economic.

In a car program, the rugby line is affordable because prototypes are cheap relative to the product. In chipmaking the prototype lab is the tool set, and a leading-edge tool set is priced like a small fab. Putting materials scientists, equipment engineers, chipmakers and university researchers around the same full process flow, generations ahead of production, has historically meant building fab-class cleanroom capacity for work that generates no wafers to sell. Few organizations could justify that for their own programs. Fewer still could justify it for someone else's. EPIC is designed to remove that barrier, with a single capital asset serving many bilateral programs. Applied's framing is that the aim is to make each runner faster and also to

change the shape of the race. We think the second half is where the value sits.

The costs of the relay are not equally visible across multiple dimensions.

The first is **latency**. Every time work crosses an organizational line, some of it has to be redone. A process developed on one set of chambers must be requalified on another. Recipes are transferred, tuned, and sometimes rebuilt. Rework accumulates quietly, weeks at a time, across every handoff in the chain. Latency is the friction the industry already measures, which is exactly why it gets the most attention.

The second is **meaning loss**. A technical requirement written in one organization is read in another, often in another country and another language, by engineers with different tools and different definitions of success. A term such as acceptable defectivity or thermal budget can shift meaning in transit without anyone noticing, because each party reads the document correctly by its own standards. The error only surfaces when the process fails to integrate, months later and several organizations downstream. Nothing in a schedule captures this. It looks like bad luck.

The third, and in our judgment the most expensive, is **unintentional misdirection**: research effort spent on targets that someone further down the chain already knows cannot be commercialized. A material may perform beautifully in a published result and still be unable to survive the temperatures of later processing steps, or be impossible to deposit uniformly at production throughput. The equipment maker often knows this early. The university lab frequently does not, because the

information sits two organizations away. Years of careful work can then flow toward a dead end, and conventional cycle-time metrics will record that work as efficient, because each stage was completed on schedule. Misdirection measures whether the project should have existed at all, a question speed metrics are not built to ask.

This is why we treat misdirection as the center of the argument. **Latency and meaning loss slow good ideas down. Misdirection spends the budget on the wrong ideas.** The remedy for misdirection runs through the requirement. When the people who will build and run the process are present at the start, the customer requirement can shape tool and process design early, before choices harden into a starting point that later proves expensive to unwind. Toyota's product development research describes the value of front-loading decisions in this way, and we borrow only that narrow principle here. The payoff from getting early decisions right shows up less as calendar speed than as a higher share of programs that reach production, a point we return to in Section 6.

The strongest objection deserves a hearing. The relay exists for reasons. Stage gates protect fabs from immature processes, specialization concentrates expertise, and clean handoffs create clear accountability and clear intellectual property boundaries. Takeuchi and Nonaka themselves noted that overlapping development demands more coordination, and suggested the approach fits some projects better than others, with fundamental breakthroughs among the harder cases. A rugby line can drop the ball, and when it does, ownership of the fumble is less obvious than in a relay.

We think the objection narrows the claim without defeating it. EPIC's overlap is concentrated between applied research and development, where the material is known but its manufacturability is not, rather than at the level of fundamental discovery, which stays with university partners upstream. Roles inside each program are defined. Applied develops the product against the customer's requirement, and the customer's engineers supply the requirement and validate against it in rapid learning cycles. That structure keeps accountability legible while removing the transfer itself. The harder question, how competitors share one building without sharing their roadmaps, belongs to Section 6.

The relay made sense when each runner could afford only its own lane. The question now is whether a single lane, paid for once and used by many, changes what the race is worth running for.





4. Distributed Labs, and the Layer Above Them

A fair question at this point is whether EPIC duplicates something the industry already does. Equipment makers (including Applied) have long run applications and development labs close to their largest customers, and several have expanded those networks recently. If the problem is slow development, more labs near more fabs looks like the direct answer. We think it answers a different problem, and that distinction is where much of EPIC's rationale sits.

Start with what a customer-adjacent lab is for. It sits near a fab and works on the customer's current generation and the one after it. Its job is incremental: tune a recipe, qualify a chamber, chase down a defect that appeared on the production line last month. Its value comes from proximity. Engineers can move wafers between the lab and the fab in hours, and the problems are defined by a process that already exists. Applied operates labs of this kind and, as we understand it, intends to keep investing in them. Nothing about EPIC appears designed to pull that work away from the fabs it serves.

EPIC is aimed one layer up the continuum described in Section 2. Its programs target technologies Multiple generations ahead of production, across the full process flow rather than a single step, with the customer's own engineers participating. One way to picture the arrangement is as a hub with spokes. The hub is where future-state work is shaped, and the regional labs carry maturing work to specific fabs and specific production lines. The horizon at the hub is longer and the questions are

less defined. The difference that matters most, in our view, is tolerance for failure.

That difference is easy to understate. **A lab next to a running fab cannot afford much speculative failure.** Its tools, its engineering hours, and its standing with the customer are spent on problems that matter this quarter, and an idea that does not pan out is a cost with nothing to offset it. At the applied-research layer the economics reverse. Several generations out, the cheapest thing an organization can do is abandon an idea early, before a tool program or a process flow has been built around it. A candidate material that cannot survive later thermal steps, or a structure that cannot be measured at production speed, is far better discovered in the first months than in the fourth year. EPIC appears designed to make that kind of early filtering routine, testing ideas against the customer's requirement while they are still inexpensive to drop. This is the practical remedy for the misdirection cost described in Section 3, and it is work a fab-adjacent lab is poorly placed to do, however capable its engineers.

This is also why we do not read a larger network of distributed labs as an equivalent response. Adding labs adds capacity. More engineers run more experiments in more places, and development throughput rises, which has real value. But the costs we identified in Section 3 accumulate between organizations rather than inside any one of them. A lab near a

fab does little to shorten the path from a university result to a qualified process, to stop a requirement from shifting meaning as it passes from chipmaker to equipment maker, or to keep research effort from flowing toward a material the fab will never accept. A wider network may, if anything, add more joins to manage.

None of this makes the two models rivals. Our read is that they are complements, sequenced along the continuum, with the regional labs handling proximity and the hub handling direction. The open question is whether the hub can hold its failure tolerance under commercial pressure. A center built to kill ideas early only earns its keep if programs are in fact allowed to end early, and that discipline is easier to describe in a launch document than to sustain across years of partner expectations. We return to how it might be measured in Sections 6 and 7, along with the question of concentrating this much future work in one region.

How Others Are Approaching the Same Challenge

Applied is not alone in identifying slow development as the constraint worth attacking. Its largest peers have each made substantial commitments of their own, and each addresses a different part of the friction described in Section 3.

Lam Research announced in August 2026 that it plans to invest more than \$3 billion over five years to expand its global lab network, with a goal of raising experiment capacity by more than half. The network links labs specialized for foundational research, process development and qualification near customer sites, and runs them as a single around-the-clock system that shares tools and data across time zones. Lam reports that the approach has shortened process development by up to two and a half times in some customer engagements. It is a serious answer to latency: more experiments, run faster, with learning carried from one site to the next.

Tokyo Electron is taking a different route. It plans to spend more than 1.5 trillion yen, roughly \$10 billion, on research


and development over the five years through fiscal 2029.


In October 2025 it completed a large development facility in Kumamoto, close to TSMC's operations there, with next-generation lithography processing among its priorities. TEL has also said it works with major customers on developments spanning multiple generations. Its bet centers on proximity, placing development near where a leading customer's volume manufacturing happens, which speaks most directly to meaning loss at the join between equipment maker and fab.

The joint lab ASML and imec opened in Veldhoven in 2024 is an adjacent model rather than a competing one, since ASML does not supply the deposition and etch steps Applied does. It gives logic and memory makers, along with materials and equipment suppliers, access to a prototype High NA EUV scanner and its surrounding tools before the systems reach their own fabs. It demonstrates the value of early, shared access to a single high-value capability.

Read together, these are three reasonable answers to the same problem, each strong where it aims. Lam scales depth and speed across a connected network. TEL places development beside the customer. ASML and imec open early access to one essential tool. EPIC's bet is different in kind. It assumes the hardest remaining problems span many adjacent process steps at once, and that resolving them requires the customer's engineers, the equipment for the full flow and the research partners working in one place at the same time. Its peers are betting, at least implicitly, that distributed depth linked by shared data and close customer contact is sufficient. Neither bet has been proven, and it is possible each suits a different class of problem. Section 7 considers what follows if the distributed answer proves good enough.

To summarize: distributed labs near the fab keep today's node moving. EPIC is an attempt to decide, earlier and with more of the relevant people present, what tomorrow's node will be made of.





5. What the Launch Roster Reveals

A strategy announced in 2023 is a proposition. By the time the doors open, the more useful evidence is who has agreed to work inside it. Economists call this revealed preference: what companies do with their capital and their engineers tells you more than what they say about it. On that measure, EPIC reaches its unveiling in a stronger position than most new research facilities, and the composition of its partner list tells us more than simply measuring its length.

Start with the founding partners. Samsung committed in February 2026, Micron and SK hynix in March, TSMC and Intel in October. Between them they represent the leading edge of both logic and memory manufacturing, and we think the pairing matters as much as the names. AI accelerators are increasingly limited by how quickly data moves between processor and memory rather than by raw computing power, which is why high-bandwidth memory is now stacked and packaged directly beside the processor. Progress on that problem depends on logic, memory and the packaging that joins them advancing together. Micron's engagement spans DRAM, high-bandwidth memory and NAND, connected to its innovation hub in Boise. SK hynix's covers memory materials, process integration and 3D packaging for next-generation DRAM and high-bandwidth memory. A facility where logic and memory programs draw on the same materials and process base is, at least in principle, positioned to catch interactions that separate programs would find later.

These four companies also compete with one another, in foundry services and in memory. That each chose to commit

engineers suggests each judged the access worth the exposure. Confidentiality is a real risk, and we treat it at full strength in Section 6. The point here is narrower: the most sensitive buyers in the industry appear to have run that calculation and signed.

It helps to be precise about what each partner joined. EPIC is a platform, with the Silicon Valley center at its core and related work elsewhere, including advanced packaging commercialization in Singapore. Broadcom's engagement, announced in May, is with the platform and centers on advanced packaging for AI systems. SCREEN Semiconductor Solutions joined the Silicon Valley center itself to develop co-optimized wafer cleaning. Advantest, a test equipment specialist, joined as well. We consider each engagement in the scope announced rather than folding every name into one building.

We visited Applied's Tampines facility during the expansion and discussed packaging R&D at the APDC. Two things stayed with us. The first was a robot on the factory floor being trained by a human engineer to assemble part of a tool. The robot was not running a pre-loaded routine, but learning, repeating a sequence while the engineer corrected its movements in real time. The complexity of the machines EPIC is designed to develop is easy to understate in a paper like this; watching one being built makes the point in a way that specifications do not. The second was the factory itself. The scale and precision required just to manufacture the equipment (before it ever reaches a customer's cleanroom) is staggering, and it puts a floor under the barriers to entry that the rest of this paper discusses in abstract terms. These are not observations that appear in a press release.

They are the reason we think the co-location model matters: the feedback loops that EPIC is designed to shorten are not administrative inefficiencies. They are engineering problems embedded in machines of extraordinary complexity.


The equipment partners deserve more attention than they will likely get in initial reporting. A manufacturing flow includes steps Applied does not supply, and cleaning and test sit directly alongside the deposition and etch steps it does. Tuning those steps to synchronize earlier is the same integration problem Section 2 identified across four prior transitions. Inviting other suppliers in also tests whether EPIC can operate as a convener for the whole flow rather than a showcase for one vendor's tools. The early signs point toward the former, though two partners do not yet establish a pattern.

The universities address the front of the continuum. Arizona State, Rensselaer and Stanford joined as inaugural research partners in May, with UC Berkeley following in August. Their role bears most directly on the misdirection cost described in Section 3. A research group with early sight of what the manufacturing end can accept may choose different problems, or abandon unpromising ones sooner. That benefit is plausible, hard to measure, and slow to appear.

Broadcom is the most telling early name. It does not run fabs. As a designer of custom AI silicon and networking chips, it commits to designs years before they are manufactured, and it needs to know what will be buildable when that time arrives. Its presence suggests EPIC's value may extend beyond chipmakers to companies that set requirements from further down the chain. If more design houses and systems companies follow, we would read that as a meaningful signal, and we return to it in Section 8.

Applied also indicates that programs are already defined and running ahead of full tool installation. That matters because it means the partnerships have moved past signing ceremonies into scoped work.

What the roster does not show is equally important. Partnership announcements disclose intent, not depth. Program counts, engineering headcount and funding terms are not public, and commitment almost certainly varies from one partner to the next. The list is also still growing, so any count we offer would be out of date within months. We read the roster as strong evidence of where the industry's most demanding buyers believe applied research should now happen. It is not yet evidence that doing it there pays off. That question belongs to the capital case, which comes in Section 7.





6. What Has to Go Right

The case we have made rests on a model that has not yet run at full scale. We see several risks worth taking seriously. We take each in turn, along with what addresses it and what it leaves open.

Confidentiality. TSMC, Intel and Samsung compete for the same foundry customers, and Samsung, SK hynix and Micron compete in memory. All are committing engineers and roadmaps to a facility run by a supplier to each of them. The visibility we credit to Applied in Section 7 is precisely what these partners would want protected. If their confidence slipped, they might bring less of their real roadmaps into the building, and the programs would weaken without anyone formally leaving. EPIC's structure is designed with this risk in view. Programs are bilateral, one customer working with Applied on its own requirement, not a shared pool where competitors pursue a common project. Applied also has decades of practice holding information boundaries between competing customers, since it already supplies all of them and works inside their processes. The partners' own commitments are evidence that the most demanding buyers in the industry judged those safeguards adequate. What remains is that concentration raises the stakes of any lapse. Over time, whether partners expand their engagements will be the clearest measure of their confidence.

Geographic Concentration. Placing the center of future-state development in one region adds exposure to seismic risk, regional costs, talent competition and policy change. It also puts that work at a distance from the fabs in Taiwan, Korea

and elsewhere where the resulting processes will run. The platform model addresses much of this. Packaging work runs in Singapore, university partners span several states, and the customer-adjacent labs described in Section 4 remain near the fabs. The Silicon Valley choice also has its own logic, because the region concentrates chip designers, systems companies, suppliers and research universities to a degree few places match. We read the concentration as a deliberate choice about where future-state work happens, not a step away from customers. The physical exposure that comes with any flagship site is a risk managed rather than removed.

Measurement Lag. EPIC is designed to pay off on the long clock, measured in technology generations, and conventional R&D metrics show little in between. Investors will be funding the model for years before a full cycle can confirm the result. Leading indicators help close that gap, even if none is conclusive on its own. Partner renewals and expansions, the arrival of more fabless and systems designers, integration problems surfacing earlier in programs, and programs ended early by choice would each suggest the model is working before any full cycle closes. How much of this becomes visible to outside observers will depend on what Applied chooses to report, and that disclosure will shape how the model is judged.

Execution. Co-location does not guarantee collaboration, and shared facilities can drift into separate teams working under one roof. As we noted earlier, the model also depends on ending programs early when the evidence calls for it, a discipline that is hard to hold in any partnership. The defined roles described in Section 3 are the main safeguard. Applied

develops the product against the customer's requirement, and the customer's engineers validate it in rapid learning cycles. Applied also indicates that operating practices are already being tested on live programs ahead of full tool installation. That is an encouraging early sign. The fuller test will come when the building reaches full occupancy.

Competitive Response. EPIC's case rests on the premise that co-location solves problems a distributed network cannot. If Lam's expanded network, Tokyo Electron's development presence near customer fabs or approaches not yet announced resolve enough of the handoff problem at lower cost and with less organizational complexity, the advantage of EPIC's hub model over them would narrow. A distributed model also

sidesteps some of the risks above, including geographic concentration and the demands of hosting competing partners under one roof. Part of the answer lies in scope. The approaches described in Section 4 appear well suited to latency and proximity, while the multi-step integration work EPIC targets is a different problem, one that seems to benefit from customer engineers and the tools for adjacent steps being present together. Applied also operates its own customer-adjacent labs, so it is not choosing one model to the exclusion of the other. What remains is that the relative value of the two approaches is an empirical question the industry has not yet answered.


From a relay to a line moving together

How the EPIC Center changes the path from research idea to production chip

■ Research ■ Equipment and process development ■ Integration, pilot and volume ramp

• BEFORE EPIC

A relay: each organization runs its leg, then hands off



• WITH EPIC

A line moving together: the same work, overlapped, in one place



Source: Applied Materials; HyperFRAME Research analysis. Phase lengths are illustrative midpoints of Applied's ranges. 6–8 years is Applied's stated target for EPIC and related initiatives, not a forecast.


History offers a useful check on all of these risks. Sematech, formed in 1987 as a consortium of U.S. chipmakers with matched federal funding, is the industry's best-known experiment in collaborative R&D. Assessments by Hart for the Bipartisan Policy Center and by Wessner and Howell for CSIS suggest its strongest results came from roadmapping, improving supplier quality and pooling learning on equipment. The structural contrasts with EPIC are significant. Sematech was a public consortium of competing device makers. EPIC is vendor-hosted, privately funded, bilateral and aimed at commercialization, with roles defined inside each program. The lesson we draw is that collaboration among competitors has worked when roles are clear and what is shared is infrastructure and learning rather than the product itself. EPIC appears designed along those lines.

EPIC also differs from today's collaborative research centers, such as imec in Belgium, providing neutral, shared access for pre-competitive work. EPIC offers depth with one customer's

roadmap and engineers at a time. We see these as different models serving different purposes rather than rivals, and Applied works with both.

That returns us to the question left open at the end of Section 2: can a standing structure outperform the temporary ones assembled for past transitions? Our tentative answer is yes. A temporary structure disperses what it learned once a transition reaches production, while a standing one can carry that learning into the next. Whether EPIC realizes that advantage depends on how it manages the risks above. None of them appears disqualifying, and all will bear watching as the rollout continues.





7. The Capital Question

Every investor reading about EPIC will eventually ask the same thing. Why would an equipment company spend billions on a facility used by customers who compete with one another, to develop technology those customers will own and run in their own fabs? There are two readings. In the first, Applied is underwriting its customers' research, a subsidy that is tempting to see as cost today in the hope of goodwill tomorrow. In the second, Applied is acquiring strategic assets it cannot buy any other way and promoting itself to the logical partner in future nodes based on co-research. We think the second reading is closer to right, though not entirely, and the balance between the two is what determines whether the investment pays.

Start with the size and shape of the commitment. Applied has said it plans to invest \$5 billion, up from the figure of up to \$4 billion it cited when EPIC was announced in 2023, and that spending is expected to scale toward that amount as customer projects begin. The second part matters more than the headline. Capital that follows program demand is a different risk profile from a lump sum committed on faith. It suggests the facility is better understood as the physical routing point for Applied's forward R&D than as a standalone real-estate line, and that the pace of spending is designed to track evidence of partner commitment.

The figure is large, though not unique. Lam plans more than \$3 billion over five years for its lab network, and Tokyo Electron more than 1.5 trillion yen, roughly \$10 billion, for its entire research budget over a similar period. The three numbers measure different things - a single facility, a lab network and a

company's full research spending - so they do not rank neatly against one another. The more useful comparison is between the models the money pays for, which is where this paper has focused.

What does that capital buy? We see three assets.

Greater roadmap visibility. Equipment makers have traditionally learned the shape of a new node when their customers were already close to defining it, then competed to supply tools into that definition. Working multiple generations ahead with several leading logic and memory makers at once would give Applied an earlier and wider view of where the industry's requirements are heading. The value lies less in any single customer's plans than in seeing where they converge. That view can direct Applied's own research budget, and in effect reduces the misdirection cost from Section 3 inside Applied itself: fewer internal bets placed on directions the customer base is not going to take.

Early design-in. In our experience, the supplier whose tools are used to develop a process often has an advantage when that process moves into volume, because requalifying a different tool late carries cost and risk. Installed tools also anchor service relationships that can run for the life of the equipment. Being present before architectures and process flows freeze would, in principle, position Applied for a richer tool mix and longer service revenue. We would not overstate this. Leading chipmakers multi-source deliberately, and development presence is an advantage in a competition, not a guarantee of the outcome.

Improved R&D productivity. A single shared facility serving many bilateral programs avoids replicating leading-edge tool sets at each customer site for this class of work. Earlier access to customer materials and requirements should also pull the discovery of integration problems forward, when fixing them is cheaper and requires fewer iterations. This is the asset most directly tied to the cycle-time question. Applied's own framing is that the path from early research to a chip in production has taken 10 to 15 years, and that it is targeting 6 to 8 with EPIC and other initiatives. We treat that as a target, not a forecast. Even partial progress against a baseline that long would represent years recovered per technology generation, which is why we think it merits attention despite the uncertainty.

Calendar speed, however, may not be the most important form of return, and the pharmaceutical industry offers a useful caution. Scannell and colleagues documented in 2012 that drug R&D output per dollar had fallen by roughly half every nine years since 1950, despite enormous gains in underlying science. They called it Eroom's Law, Moore's Law in reverse. Semiconductor research faces a version of the same problem: each new material must displace an incumbent process that decades of refinement have made extraordinarily good, which makes the bar for replacement steadily higher.


AstraZeneca's response is instructive. Morgan and colleagues reported in 2018 that after the company restructured how it made early project decisions, the share of candidate drugs that went on to complete Phase III trials rose from 4 percent in 2005 to 2010 to 19 percent in 2012 to 2016. The gains came mainly from **better choices at the start rather than faster**

execution later. By analogy, EPIC's return may show up less as shorter programs than as a higher share of programs that reach production. If that analogy holds, some improvements might become visible through intermediate signals, such as fewer late-stage reworks or problems surfacing earlier in the learning sequence, before any full development cycle has closed. That is our speculation about how progress could be observed, not a commitment Applied has made.

The analogy has limits, and we should name them. Drug development is gated by regulators and clinical outcomes; chip development is gated by yield and cost at volume, which makes it faster to learn from failure. AstraZeneca also improved its success rate partly by cutting projects aggressively, a discipline that is easier inside one company than in a shared facility where each program has a partner with its own expectations. As we noted in Section 4, whether EPIC can end programs early may turn out to be one of the better tests of whether the model works.

So is Applied underwriting customer R&D or acquiring advantage? Some of both. Partners gain real access at Applied's expense, and part of the \$5 billion will function as a cost of admission to the conversations that set future nodes. What Applied appears to receive in return is position: a place at the layer where requirements are formed, rather than only the layer where they are bid. That is a more durable asset than any single tool win, if it can be defended. Whether it can depends on Applied's ability to manage the risks and communicate the rewards.





8. What to Watch

The unveiling marks a beginning rather than a result. For readers tracking EPIC from here, we see five signposts worth following over the next several years. None will settle the question alone, but together they should show whether the model is working well before a full technology cycle has run its course.

The first is who joins next. The founding partners establish that the leading logic and memory makers see value in the model. The more telling additions would be further fabless designers and systems companies, the firms that commit to designs years ahead and need an early read on what will be buildable. Broadcom's engagement suggests that demand exists. A broader set of design-side partners would suggest it is a pattern.

The second is how existing engagements evolve. Renewals, expansions into new technology areas and partners bringing more of their own engineers would each indicate that the model is delivering value to the companies closest to it.

The third is what Applied chooses to disclose. Leading indicators such as issue discovery moving earlier, program milestones and partner growth are the evidence outside observers will need. The more of it Applied reports, the sooner the model can be judged on results rather than intent.

The fourth is the pace of capital. Applied has tied spending to the start of customer projects. Tracking how that spending scales against program activity would give a practical read on demand.

The fifth arrives later and is harder to see: technologies developed in EPIC programs appearing in customers' production process announcements. Attribution will rarely be clean, because customers will present these as their own achievements, and they will be right to. Even so, the timing of those arrivals against historical norms will be the fullest test of the model.

Behind all five sits the contrast this paper began with. AI products now turn over roughly every year. The materials those products depend on still take the better part of a decade or more to move from laboratory to fab. For most of the industry's history, equipment makers have sold into that longer cycle once its direction was already set. EPIC is an attempt by Applied to work inside the cycle itself, reorganizing the research that feeds it so that it moves less like a relay and more like a line advancing together. If that works, the pace of the slower clock starts to set the pace of the faster one, and the company helping to compress it holds a stronger position than one supplying tools after the fact.



This paper was commissioned by Applied Materials.
HyperFRAME Research conducted the analysis independently,
and the conclusions are our own.

ABOUT HYPERFRAME RESEARCH:

HyperFRAME Research delivers in-depth research and insights across the global technology landscape, spanning everything from hyperscale public cloud to the mainframe and everything in between. We offer strategic advisory services, custom research reports, tailored consulting engagements, digital events, go to market planning, message testing, and lead generation programs.

Our industry analysts specialize in rigorous qualitative and quantitative assessments of technology solutions, business challenges, market forces, and end user demands across industry sectors. HyperFRAME Research collaborates closely with your Analyst Relations, Product, and Marketing teams to build and amplify your thought leadership, positioning your expertise to enhance brand and product recognition. Through content that engages readers, viewers, and listeners alike, we ensure your voice resonates across channels.

CONTACT HYPERFRAME RESEARCH:

Steven Dickens
CEO & Principal Analyst | HyperFRAME Research
Email Address:
steven.dickens@hyperframeresearch.com
Telephone Number:
+1 845 505 1678
X: [@StevenDickens3](#)
LinkedIn: Steven Dickens
BlueSky: Steven Dickens

CONTRIBUTORS

Stephen Sopko

Practice Lead – Semiconductors & Deep Tech

INQUIRIES

Contact us if you would like to discuss this report and HyperFRAME Research will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "HyperFRAME Research." Non-press and non-analysts must receive prior written permission by HyperFRAME Research for any citations.

LICENSING

This document, including any supporting materials, is owned by HyperFRAME Research. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of HyperFRAME Research.

DISCLOSURES

HyperFRAME Research provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

