


From Lab to Fab

Using AI to Scale New Materials into Production

Atomscale's thesis and proven approach, with field results and an adoption roadmap.



From atom to scale.

–	Introduction	03
---	---------------------	----

–	Our Thesis: Charting the Path from Atom to Scale	04
	Why the materials of the future are stuck in the lab, and how AI built on the data tools already generate can bring them to production scale.	

[01]	Measurement: Next Week’s Data, Today	11
	How physics-informed models turn the signals a tool already emits into a live estimate of material state, shown on a composition metric calibrated from two reference runs.	

[02]	Understanding: The Suspicion You Can’t Test	16
	How runs already on disk reveal how a process actually behaves, shown on a year-old control-loop suspicion settled from the archive in days.	

[03]	Detection: The Excursion Started Weeks Ago	20
	How changepoints with attribution separate intended change from drift, shown on an MBE surface change caught 40 seconds before manual observation.	

[04]	Action: Steer the Run While It’s Still Growing	24
	How live state, learned dynamics, and detection combine to steer runs as they grow, shown on growth outcomes predicted from the first minutes of a run.	

[05]	From Pilot to Production	29
	How teams adopt materials enablement in three phases, from a first result on historical data to a process that improves on every run.	

–	About Atomscale	35
---	------------------------	----


Materials innovation has a scale-up problem.

The materials that will define the next decade of computing, communications, and quantum technology mostly exist already. They have been made once, somewhere, in a lab. Getting them into production reliably is still slow, and it still runs largely on trial and error.

This collection brings together the thesis we published and the four-part series that followed it. The thesis makes the argument: materials innovation has a scale-up problem, and the information needed to solve it is mostly going to waste. The series breaks the answer into four capabilities, each built on the one before: measurement, understanding, detection, and action.

This edition adds a page after each part of the series that shows the capability at work on a real tool, with the results it produced and what they are worth to a development program. It closes with a chapter on how teams adopt materials enablement, from a first result on data they already have to a process that improves with every run, and on where that adoption stands across the industry today.

If your team's schedule is already full, start with the closing chapter, [From Pilot to Production](#): the first step runs on data you already have, needs no new wafers, and is provided at no cost.



OUR THESIS

Charting the Path from Atom to Scale

The materials of the future already exist in the lab. Frontier intelligence will bring them to the world.

The Promise and the Gap

In December 1959, Richard Feynman stood before the American Physical Society at Caltech and told a room of physicists there was “plenty of room at the bottom.” He was inviting them to join a new mission of exploration: the deliberate control of matter at the atomic scale. Nearly seventy years of world-changing progress followed downstream of that invitation: modern electronics, Moore’s Law, and our deepest grasp of physics itself. We learned to engineer the world atom-by-atom, and nearly everything we now call technology rests on that work.



Richard Feynman · American Physical Society Meeting
“Plenty of Room at the Bottom” · Caltech – December 29, 1959


Thanks to that work, the materials that will drive the next wave of technology — for AI, for quantum, for energy and electrification — are, for the most part, not waiting to be discovered. They are already known. They are valuable. They can be made in the lab — but they are stuck there. We cannot make them at production scale.

Put simply, materials innovation has a scale-up problem, not a discovery problem.

Every major technological shift begins in the physical world of materials. The intangible breakthroughs we celebrate — the model, the qubit, the grid — are all downstream from someone learning to manufacture a substance reliably, at yield, inside a real device. When that manufacturing stalls, the future stalls with it.

“The breakthrough was not the material. It was learning how to process the material at scale.”

Intel provides a clear illustration. By the 2000s, the silicon dioxide that had insulated the transistor gate for four decades had been thinned to a few atoms across, and it was leaking. Intel knew it needed a high-k dielectric years before it could ship one. The hafnium-based material it landed on was not a eureka discovery and required synchronized changes to other materials in the stack to be effective. More than a decade of work went into making that material manufacturable: integrating it into a real transistor stack, at yield, without breaking everything around it, by depositing one atomic layer at a time.



Intel “Penryn”-family Processors
45nm Hafnium-based High-k Metal Gate Wafer · January 27, 2007

When Intel finally shipped it at the 45-nanometer node in 2007, Gordon Moore called it the biggest change in transistor technology since the late 1960s. The breakthrough was not the material. It was learning how to process the material at scale.

The materials scale-up bottleneck is a bottleneck on our future.

02 / THE HURDLES

Why This Problem Remains


The hurdles to solving this problem are twofold: physical and informational.

The physical difficulty is that materials do not exist in a vacuum. They live in context, nested inside heterogeneous device structures, one material grown onto another, each environment changing both what optimum looks like and the path to reach it. Nearly every degree of freedom in the design of material and

process is continuous and interconnected. Simulation and digital twins can provide guidance but not reach the last-mile fidelity of a real material. So synthesis remains trial and error, guided by hard-won expertise and informed intuition. Tuning a synthesis process is like operating tweezers while wearing oven mitts.

“That’s the heart of it: the throughput of data generation has exploded, but our ability to use it has not.”

The informational difficulty compounds the challenge. Characterization is fragmented across complementary, narrow probes — XRD, XPS, RHEED, TEM, AFM, and many more — each with its own hardware, its own software, and its own sub-specialty to separate the signal from the noise. Practitioners stitch the picture together by hand: serial, operator-biased, lossy. Metadata is left behind. Null and negative results, which carry real information, are routinely thrown away. Skilled engineers may feel that their analysis is already sufficient and the data captured is excessive, but nothing in the toolchain ever surfaces what’s being discarded.



Molecular Beam Epitaxy (MBE) Tool · National Renewable Energy Laboratory (NREL)

That’s the heart of it: the throughput of data generation has exploded, but our ability to use it has not.

Why Now

So why is this problem solvable now, when it was not before?

Two forces have converged. Modern tools have become sensor-rich and high-throughput, producing real-time data in volumes that were unimaginable a decade ago. AI is now capable of using such data, having passed significant thresholds in compute, memory, bandwidth, and transfer learning that works on limited datasets rather than demanding oceans of examples.

“We need systems that use the immense data that already exists to intelligently guide growths — instead of only explaining failures.”

The honest remaining constraint is the same thing that makes the problem both hard and defensible: there is no “internet of materials” to scrape. The data that does exist is rich, but it is proprietary and siloed. A generic, web-trained AI model cannot solve this, and the opportunity isn’t in waiting for more data — it’s in building a system that finally puts to use the immense amount of data already being generated.


What Atomscale Does

The bottleneck has never been a shortage of promising candidate materials. It is the decades of trial and error it takes to manufacture even one of them reliably. Atomscale turns that grind into a guided science. It reads the data your tools already produce and uses physics to turn raw signal into insight in real time — so you can steer a run while it is still growing, rather than explaining it after it has failed. Every run becomes legible. Every insight becomes cumulative. Knowledge compounds instead of dying in a silo.

Atomscale’s architecture uses a hierarchy of models. Physics-informed adapters sharpen signal-to-noise at the base and pass only meaningful information upward to time-series and reasoning layers. That hierarchy acts as a sandbox, constraining general-purpose AI to domain-relevant data, so the system behaves like a knowledgeable colleague instead of a generic LLM. Domain expertise is built into its structure, combining the leverage of automated reasoning with deterministic outputs grounded in physics. This is not a promise; it is already measurable. In encoding information about an unseen run, Atomscale performs 43x better than baseline unsupervised machine learning.


Today's tools

SILOED, SINGLE-USE MODELS



Atomscale

A HIERARCHY OF FIT-FOR-PURPOSE MODELS



Today's tools are siloed, single-use models: raw process and metrology data feeds isolated statistical, physics, and generic models and a chat assistant, most signals are lost, and insights dead-end with no carry-over across runs. Atomscale is a hierarchy of fit-for-purpose models that extract and record 43x more useful information and steer every run as it grows with real-time insights.

Atomscale changes the daily texture of work. Before: manual analysis, fragmented tools, and insight trapped in one person's head. After: automated feature fingerprints, queryable data, runs steered in real time, recipes once too complex to maintain now within reach. This lifts the burden of operating today's processes while improving tomorrow's because knowledge no longer dies at the end of an investigation. It accumulates across the organization, on a level playing field. Engineers stop being the workhorses of their processes and instead become the managers. The whole organization can begin to aim for 10x gains instead of incremental ones.

“We clear away the data drudgery so people do more of the thinking, not less. We augment the expertise that already exists.”

This is not a replacement for engineers; on the contrary, there has never been a more exciting time to wield the levers of physics and develop exceptional real-world technologies. Design, conceptualization, creativity, the setting of goals and requirements — these stay human. We clear away the data drudgery so people do more of the thinking, not less. We augment the expertise that already exists.



Feynman collaborating with prominent physicists
Shelter Island Conference on Quantum Mechanics · June 1947

05 / THE BELIEF

What We Believe


Our claim: the right AI today can extract far more meaningful information in aggregate from the data already collected than human analysis. Parts of the field may disagree. The teams that embrace it are the ones that will lead the coming wave of materials innovation.

We're headed to a place where materials measurement data becomes a true asset: a physics-rich, interoperable space that reasoning agents can navigate and monitor continuously. Scale-up timelines will compress from a decade toward something far shorter. The materials trapped in the lab today will begin to reach the world on a human timescale.

We drilled to the bottom decades ago. Feynman was right: there was plenty of room, and we have spent nearly seventy years proving it. The first phase of exploration is complete, and the next phase is different. We must learn to navigate what we found and bring it back up, from atom to scale.

— Atomscale

BOSTON, MASSACHUSETTS



[01] MEASUREMENT

Next Week's Data, Today

Turn the signals your tools already collect into a live estimate of material state, so learning no longer waits on the characterization queue.

MATERIALS ENABLEMENT
PART 1 OF 4

Somewhere on one of your tools, a film is growing right now whose composition nobody will know until next week. The run will end, the wafer will join the characterization queue, and the first hard fact about the material will arrive days after every decision about it has already been made.

The industry has spent several years pursuing the promise of AI for materials discovery, and this effort is warranted. Screening and structure prediction now surface promising candidates faster than any lab can grow them. But discovery answers the question “what material?” The question that new technology rests on is “how do we make it, every run?”, and every advance on the first question widens the gap to the second.

Materials enablement is our name for the work of closing that gap, the same gap our thesis describes: applying AI to the journey from “works once” to “works every time,” built on the process data a lab or fab already generates.

No such system exists off the shelf, because the two disciplines that might have built it were aimed elsewhere. Statistical process control assumes a mature, settled, high-volume process. Discovery AI points upstream, at the candidate. Between them sits the actual decades of scale-up and uncertainty, still run on trial and error and the memory of whoever has been with the tool longest.

Four needs define the missing system: measurement, understanding, detection, and action. Each one unlocks the next, so this series takes them in order. Measurement comes first because everything else depends on it.

A team cannot learn from a run it never saw, catch drift it never measured, or steer a process it cannot observe. Everything downstream in this series rests on one capability: a trustworthy estimate of the state of the material on every run, arriving fast enough to act on. How fast depends on the process and the signals it emits, and the direction of travel is toward an estimate available while the run is still live.

TODAY'S CEILING

Late, Partial Signal

The estimate most teams actually have is retrospective. Ex-situ characterization is slow, often destructive, and samples a sliver of the material: a few checkout wafers per campaign, a few sites per wafer, with results hours, days, weeks, or months behind the run. Meanwhile the properties that decide success (composition, thickness, stress, defect density) have no instrument that reads them mid-growth.

The chamber does stream varied data while the film grows. Diffraction, spectroscopy, reflectometry, pyrometry, dozens of tool-trace channels. But these are proxies, raw signals correlated with the properties you want in ways that nobody has fully written down.

Teams bridge the gap themselves where they can. A staff engineer fits a calibration tying one proxy to one property, and it holds near the center of the process window where it was built, then degrades toward the edges, which is exactly where new materials development happens. The obvious upgrade, a proper machine-learning model, runs into arithmetic: classical ML wants hundreds or thousands of labeled runs, and a team scaling a new material has a few dozen.

“If structural data takes two weeks to come back, then two weeks is the floor under your learning rate.”

None of these constraints comes from a shortage of data. A materials process tool produces more data each month than its team can fully analyze; what the data lacks is interpretation. And until it gets interpreted, the characterization queue sets the speed limit of the entire program. Experiments get batched to fit the queue. Decisions wait on it. If structural data takes two weeks to come back, then two weeks is the floor under your learning rate, whatever pace the tool itself could sustain.

THE SOLUTION

Building a Live Estimate of Material State

The way out is to change the job ex-situ metrology does. Atomscale's physics-informed models can fuse the sensor streams a tool already emits into a live estimate of material state. The concrete forms are virtual and computational metrology and learned composition metrics: models that infer the properties you care about from the signals you already collect, continuously, while the run is in progress.

Physics makes the approach workable from the earliest stages of development when volumes are low. Because the physical constraints are built into Atomscale's model, it needs only reference runs to calibrate, and a handful will do. That one fact moves the technique out of the expensive high-volume labeling regime where classical ML lives and into the regime of real-world early process development: low volume, high mix, processes still being learned.

Ex-situ characterization keeps its place in this picture, with a different job. It becomes the anchor, the ground truth that validates and periodically re-grounds a model running underneath every run.

Faster Cycle Times, Faster Scale-Up

Measurement becomes continuous: every run, every wafer, every second, with material state available as a live signal during growth. That signal is the core input to every later pillar in this series.

The first return shows up in cycle time. Development stops waiting on the queue, because the queue no longer holds the only copy of the truth. The second return is larger: measurements that never existed become routine. One compound semiconductor team we work with needed within-run wafer composition data they could not get at any turnaround time. A learned metric, calibrated on a few reference runs, now delivers it live, with a tenfold accuracy gain over their previous best estimate (RMSE near 1.0 down to about 0.1) against the same ex-situ reference.

That second category is the one worth a leader's attention, because the two returns land on different lines of the business. Continuous versions of yesterday's measurements make a program faster: fewer months to qualification, fewer wafers spent learning, a development budget that buys more experiments. That is a cost line. Measurements that never existed change the unit process itself, on every run, and that reaches the revenue line: higher yield when the process moves to volume, a material held to spec closely enough to sell, a product window a competitor working at queue speed will not make. The first makes the program you already have cheaper to run. The second makes viable a program that was not.

Through the rest of our series on materials enablement, we'll see how teams can practically use Atomscale to recover how processes actually behave using past runs, detect crucial moments when runs change, and eventually steer runs as they grow.

A Composition Metric the Team Can Control Against

One team we work with is building a thin-film stack for photonics and quantum devices, where film composition sets how the finished devices perform. Holding composition inside its window means measuring it during the run, and the team had spent more than a year building a metric from in-situ diffraction for that job. The metric tracked composition in parts of the process window and lost it in others, which left it too unreliable to control against.

Atomscale built a replacement from the same in-situ data. Calibrated on two reference runs, the new metric tracks composition across the whole process window, checked against XPS over 24 process conditions. It is accurate enough to control against, which gives the team a real chance of catching and correcting composition excursions during the run.

10x

Lower error than the team's best existing metric: RMSE from 1.0 to 0.1 against XPS ground truth

2

Reference runs to calibrate, where classical machine learning typically needs hundreds or thousands

24

Process conditions tested, spanning the full window

1 year+


Engineering time already spent on the metric it replaced

WHAT IT'S WORTH

The return lands on both lines of the business that Part 1 described. On the cost line, the learning cycle stops waiting on characterization: a program with a two-week turnaround that runs twelve learning cycles a year spends 24 weeks of that year waiting, and a live estimate can return much of that time to development. On the revenue line, a property the team can measure on every run is a property it can hold to spec, and yield at volume depends on holding it there.

WHERE IT FITS IN THE ROLLOUT

Teams usually start here or with understanding. In a pilot, historical runs paired with your existing metrology show whether an in-situ proxy holds on your process, and the live estimate follows once the tool is connected. See [phases 1 and 2 in From Pilot to Production](#).



[02] UNDERSTANDING

The Suspicion You Can't Test

Recover how your process actually behaves from runs you already have, and turn hard-won experience into knowledge that lasts.

MATERIALS ENABLEMENT
PART 2 OF 4

Every process on your line runs on numbers somebody once wrote down: response lags from the vendor spec, thermal constants measured during install, drift behavior remembered from the last excursion. Some of those numbers are wrong. Somewhere in your organization, an engineer already suspects which ones, and cannot prove it.

This series is about materials enablement, the application of AI to the journey from “works once” to “works every time.” Part 1 argued that the characterization queue sets the speed limit of a development program until material state becomes a live, continuous signal. This part is about the pillar that signal feeds first: understanding. Most teams do not actually understand how their process behaves. They know how it is supposed to behave, which is different, and the difference costs months.

The cost is silent, which is what makes it expensive. Experiments get designed against the assumed dynamics, so when the assumptions are off, the experiments answer the wrong question. Nothing fails loudly. The program simply converges slower than it should, and nobody can say why.

TODAY'S CEILING

Dynamics Taken on Faith

Start with the numbers themselves. At one lab we work with, a critical control loop turned out to respond more than fifty percent slower than the value every recipe had been designed against. The consequence was an overdriven process that had spent years oscillating around its setpoints instead of settling onto them. A senior engineer suspected it for the better part of a year. Suspicion was as far as anyone could get, because checking would have meant a dedicated experimental campaign that no development schedule could absorb.

What understanding does exist rarely lives anywhere durable. It accumulates in the heads of senior engineers, run by run, over years, and it walks out the door with them. Every departure is a small library fire.

“The people with the questions cannot reach the data, and the people who can reach the data do not know the questions.”

The record that could rebuild that understanding technically exists. Tool logs, in-situ sensors, metrology, the MES: every stream is captured somewhere. But each one is analyzed alone, and the relationships that explain process behavior run across them, invisible to any single-stream view. The archives themselves sit in formats only a software team can open, and the software team has a backlog. So the people with the questions cannot reach the data, and the people who can reach the data do not know the questions.

That leaves the designed experiment as the one sanctioned path to causality: rigorous, respected, and priced in wafers. A proper DOE isolates one factor at a time and bills weeks of tool time for the answer. The price does more than slow answers down; it rations the questions.

Recovering True Dynamics from Runs You Already Have

Atomscale's models recover how a process actually behaves from the runs a team has already completed. Response lags, time constants, drift modes: extracted from historical data, with no new experiments on the calendar. That recovery depends on keeping the record whole. Atomscale carries the full signal end to end at native resolution, so nothing dead-ends in an intermediate format and no compression step averages away the fine structure that dynamics live in. Cross-data-stream pattern detection does the same for relationships, surfacing correlations across tool logs, sensors, and metrology that no univariate analysis can see.

Just as important is where the capability sits. Atomscale puts these models in the hands of the material scientists themselves, so a question that once required a software project, or a wafer budget, becomes a query. And the live estimate of material state from Part 1 feeds the same models, giving them far more to explain than raw tool signals alone.

THE UNLOCK

A Process You Can Interrogate

The process becomes a model you can question. A hypothesis that once justified a DOE, or died for lack of one, gets tested in days against runs already paid for. That is how our customer's year-old suspicion about their control loop was settled: from the archive, in days, without a single new wafer. The lag was real, the recipes were rebuilt around the measured value, and the oscillation stopped.

Understanding also starts to compound. Knowledge captured in a model survives reorganizations and retirements; the library stops burning. For a leadership team this is the deeper return: process understanding turns into an asset that appreciates with every run instead of evaporating with attrition.

The next possibility is transferability. Dynamics learned on one tool can seed the model for the next tool, the next recipe, the next site, so understanding stops restarting from zero at every transfer. Programs begin to inherit from each other, and each one starts its learning curve partway up.

A Year-Old Suspicion, Settled from the Archive in Days

The control loop in Part 2 belongs to a photonics team we work with that processes thin-film oxides inside a very narrow window, under a highly multivariate control system. Within a few months the team suspected a lag between a setpoint change and the wafer's true response. Pinning it down would have meant pulling engineers off competing priorities, and the evidence sat in large, noisy tool logs and in-situ metrology the team had no clean way to read.

Atomscale isolated the effect of individual controls from data already on disk, aggregating the material's response to every setpoint change across many runs. The measured lag was 75% longer than the control system assumed, and the process had been overdriven as a result. The team retuned against the measured value, and the analysis became a workflow its own engineers now run whenever the question comes up again.

The same approach carries over to other deposition methods. On ALD runs of hafnia dielectrics, Atomscale turned raw in-situ ellipsometry into fingerprints at the resolution of individual deposition cycles, separating three substrate temperatures and showing the cycle at which Zr-doped and undoped films began to diverge. Ellipsometry is recorded on nearly every ALD run and rarely analyzed beyond film thickness, so the data was already on hand.

75%

Longer lag than the control system assumed

Days

From the question to a measured answer

0

New wafers or designed experiments needed

~1 year


Time the team had spent unable to prove what it suspected

WHAT IT'S WORTH

The first return is engineering time: the answer came from runs already paid for, where a dedicated experimental campaign would have cost several weeks of tool time and a wafer budget. Every experiment designed afterward returns more, because it is designed against the dynamics the process actually has. And because the result lives in a model and a reusable workflow, it stays with the organization when the engineers who found it move on.

WHERE IT FITS IN THE ROLLOUT

Understanding is often the first result a pilot delivers, because recovering dynamics needs only historical data. A team can have an answer to a long-standing question within days of sharing its archive. See [phase 1 in From Pilot to Production](#).



[03] DETECTION

The Excursion Started Weeks Ago

Distinguish intended change from unexpected drift with attributable signals that catch excursions while the run is ongoing.

MATERIALS ENABLEMENT
PART 3 OF 4

On every chart in the control room, last week's runs look fine. One of them was not. The evidence exists right now, spread across forty sensor channels, each individually in spec. You will meet this excursion in three weeks, at end-of-line characterization, after it has repeated itself through every run in between.

This series is about materials enablement, the application of AI to the journey from “works once” to “works every time.” Part 1 built the foundation, a live estimate of material state. Part 2 turned history into understanding, a model of how the process actually behaves. Detection is where those investments start paying operational rent: knowing the moment anything changes, in a process where change is constant.

That last clause is the paradox of scale-up. The process is supposed to be changing. Setpoints move, recipes get revised, and deliberate change from run to run is the entire point of development. That is precisely what makes unintended change so hard to see. When change is the baseline, drift wears camouflage.

TODAY'S CEILING

Excursions Found Weeks Late

Today, most excursions surface at end-of-line characterization, weeks and many wafers after they began. The direct cost scales with time-to-detect: every additional day is more scrapped material. The indirect cost is worse. Every experiment run during the undetected window was interpreted against a process that was no longer the process, so the excursion consumes wafers and corrupts conclusions in the same stroke.

The standard toolkit was built for a different world. Statistical process control assumes a settled process: stable mean, known variance, control limits earned over hundreds of runs. A development process violates every one of those assumptions by design. Teams that apply SPC anyway get a flood of false alarms, respond by widening the limits, and end up with charts that can no longer see anything at all.

Even a settled process hides two kinds of change from single-signal charts. One is the joint shift: forty channels each within limits while the relationship between them walks away, a change no univariate view can represent. The other is slow drift, from consumable aging, chamber conditioning, even seasonal effects in the facility, each too gradual to cross any threshold on any one chart.

“An ignored alarm system is worse than none, because it converts vigilance into ritual.”

The modern fix, generic anomaly detection, fails socially rather than statistically. A physics-blind model scores strangeness without meaning. Engineers learn within a month which alerts to ignore, and an ignored alarm system is worse than none, because it converts vigilance into ritual.

Changepoints with Attribution

Atomscale treats detection as a modeling problem before an alerting problem. Changepoint detection runs over the full multivariate stream, referenced against a physics-grounded model of what this particular run should look like, the same model hierarchy built up in Parts 1 and 2. Because the reference knows the recipe, intended change and unintended drift separate cleanly: moving a setpoint is expected, while the chamber responding differently to the same setpoint is an event.

That grounding is also the answer to alarm fatigue. An alert from Atomscale arrives with attribution: what changed, where in the run, and against which expectation. An engineer can evaluate that in minutes, and the alert channel keeps its credibility, which is the property alarm systems usually lose first.

Excursions Caught While the Run Is Live

Time-to-detect collapses from weeks to the middle of the run. The excursion in the opening of this piece gets flagged during the run where it began, with a pointer to what moved. Post-mortem mysteries become attributable events, and “was that run normal?” becomes a question with an automatic, trustworthy answer, for every run, without anyone staring at charts.

At one materials program we work with, changepoint detection over the existing sensor archive surfaced process shifts the team’s monitoring had never registered. The runs were already grown and the data already stored; what had been missing was a model of normal against which change becomes visible.

The next possibility is prediction. A model that knows what every healthy run looks like can begin flagging the run that is going to fail before it does, and maintenance can trigger on material state instead of wafer counts. Detection matures from catching change into anticipating it.

A Change Caught Forty Seconds Before Manual Observation

During an MBE growth run, Atomscale and an engineer watched the same live RHEED feed for the onset of a surface reconstruction. That transition signals a change in growth mode, stoichiometry, or defect formation. Atomscale flagged it about 41 seconds into the growth phase, and the engineer saw it at about 81. Forty seconds is enough time to adjust flux, substrate temperature, or shutter sequence before a defect propagates into the layers above.

On MOCVD, the same approach runs on data every production run already records. A model trained on recipe parameters and sensor logs flags the runs that will produce non-uniform wafers with greater than 96% accuracy. It is tuned to favor certainty, so each alert is trustworthy enough to act on without a manual check, and a drifting process gets caught before a full cassette goes through it.

“I have been hoping somebody would be able to do this for 20 years.”

STAFF SCIENTIST AT A TOOL MAKER, ON AUTOMATED DETECTION OF ATOMIC-SCALE DEFECTS DURING MBE GROWTH

40 s

Earlier than manual observation of the same RHEED feed

0

False positives from image-processing artifacts in that comparison

>96%


Accuracy flagging non-uniform MOCVD wafers from in-process data

WHAT IT'S WORTH

The return lands on yield, and it scales with time-to-detect. An excursion found at end-of-line characterization three weeks after it began has passed through every run in between, about 15 at five runs a week, each of them scrap or an experiment read against the wrong process. If caught during the run where it started, it costs one run and leaves the experiments after it intact.

WHERE IT FITS IN THE ROLLOUT

Detection starts on the archive: a pilot can run changepoint detection across past runs and show shifts your monitoring never registered. Live, attributed alerts arrive once the tool is connected. See [phases 1 and 2 in From Pilot to Production](#).



[04] ACTION

Steer the Run While It's Still Growing

Use live material state, learned dynamics, and detected changes to intelligently guide runs today and carry process knowledge into every run that follows.

MATERIALS ENABLEMENT
PART 4 OF 4

The recipe is loaded, the source opens, and for the next several hours the most advanced material your company makes is on its own. Nothing anyone learns tonight can help tonight's run. That is open-loop development, and it is still the default across the industry.

This is the final part of our series on materials enablement, the application of AI to the journey from “works once” to “works every time.” The first three parts built instrumentation: measurement gave a live estimate of material state, understanding gave a model of true process behavior, detection gave the moment anything changes. Action is what the instrumentation was for: intervening while intervention still matters, and never learning the same lesson twice.

TODAY’S CEILING

Open Loops and Dead-End Insight

The status quo is open loop. Recipes are set before the run and judged after it, and deviations are discovered once the material exists and the money is spent. Feedback of a kind does exist in this industry, but it sits where it is least needed; run-to-run control adjusts the next run from the last one, and it serves mature, high-volume, settled processes. The processes that most need the help – low volume, high mix, still being learned – have none.

Meanwhile the learning that does happen fails to accumulate. An engineer investigates, finds the cause, presents the slide, and there the analysis stops. Nothing carries the finding into the recipe or the model, so the next run proceeds exactly as it would have. Multiply that dead end across every investigation, every team, every year. Much of what a materials organization learns, it learns into a filing cabinet. And when a process finally works, the learning refuses to travel. Ramp to a second tool or transfer to a second site, and the trial and error largely restarts, thousands of wafers’ worth, at every step of scale.

“Skip those, and a closed loop is just a faster way to make bad material.”

There is also a reason nobody has simply closed the loop on a development tool. Every tool has a controller already, and it executes the recipe faithfully without knowing anything about the material coming off it. Closing the gap between those two means intervening on the process itself, and intervention without understanding is dangerous. You can only responsibly steer a process with a model of it you trust, which is why action is the one pillar that cannot be bought as a point solution. It inherits everything upstream: the live state estimate, the learned dynamics, the change detection. Skip those, and a closed loop is just a faster way to make bad material.

Closing the Loop on Development Itself

Atomscale closes the loop at two timescales. Within the run, the live state estimate guides intervention: when the material's trajectory deviates from intent, the run can be corrected while correction is still possible. In our process intervention work with customers, this is where it begins concretely, with deviations addressed during the growth itself and the post-mortem reserved for real surprises.

Across runs, every completed growth updates the model that guides the next one, so the process improves run over run by construction, whether or not anyone writes a slide. Carried far enough, development changes character: from trial and error to guided search, with each experiment chosen because the model expects it to teach the most.

Runs-to-Spec Collapses

Steering a run while it grows means a deviation becomes a correction, and a scrapped wafer becomes a saved one. Runs-to-spec, the number every program manager actually watches, collapses, because runs stop being spent rediscovering known failure modes. And because the learning lives in models, it finally travels: dynamics learned on one tool seed the next tool, the next site, the next ramp.

The end state is the one our thesis promised: scale-up as a guided science. The decade between a material that works once and a material that ships compresses, because the system gets smarter with every run.

Four Pillars, One System

Taken together, this series describes one system. The same hierarchy of models serves every pillar: the state estimate that measurement builds is what understanding explains, what detection references, and what action steers. Teams usually enter at measurement or understanding and mature toward action, so

the arc of these four stages is also, in practice, an adoption path. A point tool can imitate any single pillar. None can compound across all four, because compounding requires shared models and run-to-run memory, and that compounding is what defines materials enablement as a category.

Why now is a short argument. The data already exists, and the extraction finally works: our models encode 43x more useful information from an unseen run than baseline unsupervised machine learning, from sensors that are already installed. The architectures finally fit the problem, physics-informed and time-series-native, learning from dozens of runs instead of thousands. And the industrial moment demands it: reshoring, advanced packaging, and the quantum and photonics ramps are all, at bottom, materials scale-up problems.

The materials of the next decade are, for the most part, already made, once, somewhere, in a lab. What stands between them and the world is the journey this series described. We think closing that gap is the defining industrial problem of the decade, and we are building the system for it.

If any of the four ceilings in this series felt familiar, we should talk.

Knowing Early Which Runs Will Fail

Steering a run starts with knowing, early, which runs are headed for failure. On one program, Atomscale predicted whether a growth would succeed or fail in 90% of cases, trained on an initial set of roughly ten labeled samples. The prediction came from roughly the first ten minutes of growth. A run flagged that early can be corrected, or stopped before it spends hours of tool time and a substrate.

90%

Growths correctly predicted to succeed or fail with no prior baseline

~10

Labeled samples needed to train that prediction

22% to 74%

Estimated completed runs-to-spec before and after

3x


Estimated reduction in tool time per successful run

WHAT IT'S WORTH

The return shows up in runs-to-spec, the number program managers already watch. Each run saved from scrap returns its tool time, its substrate, and hours to days of schedule, and so does each run that no longer rediscovers a known failure mode. The larger return is harder to price: materials too difficult to reach by trial and error become reasonable programs to attempt, which is the revenue line Part 1 described.

WHERE IT FITS IN THE ROLLOUT

Action comes last by design, because it inherits the state estimate, the learned dynamics, and the change detection. Teams turn it on once they trust the first three, usually in phase 3. Predicting doomed runs can start earlier, as an alert an engineer acts on before anything is automated.



[05] ADOPTION

From Pilot to Production

Start with runs you've already paid for, and grow toward a process that improves with every run.

THE ADOPTION PATH

Part 4 ended on an observation from our deployments: teams usually enter at measurement or understanding and mature toward action, so the arc of the series is also an adoption path. This chapter follows that path from a first result to a process that improves on every run, with what each stage asks of your team and what it returns.

Start with a Problem Already on Your Schedule

Most engineering leaders we speak with agree with the argument in this book and then tell us their schedule is full. That is a fair objection, and it is why a good first project is one already on the schedule, a problem your team is working on this quarter:

- A qualification that keeps slipping because characterization comes back late
- A property you measure after the run that you need during it
- An excursion that cost a campaign and was never fully explained
- A control loop, chamber, or recipe that behaves differently from its spec
- A tool transfer or second site coming up, where the team expects to relearn the process

Each maps to one of the four capabilities, and each can be tested against runs you have already made.

THREE PHASES


The First Results Come from Data You Already Have

Adoption runs in three phases, and the four capabilities arrive in roughly the order this book presents them. Three of the four can show results in the first phase, on historical data, before anything is connected to a live tool.

PHASE	YOUR TEAM PROVIDES	YOU GET
Proof of concept on historical data	An export from one tool (in-situ signals and tool logs, with matching ex-situ measurements), one question, and a few hours a week from a process engineer	Recovered dynamics and cross-stream relationships, changepoints across the archive, and a first in-situ proxy for an ex-situ measurement, each tied to your KPIs. No new wafers, and no cost.
Live on one tool	A data connection (real-time streaming, file watchers, or the API) and a handful of reference runs to calibrate	Material state estimated on every run, attributed alerts during the run, and analysis your engineers run themselves
Always on, then across tools	Agreement on where intervention is appropriate, and controls integration through recipe files or SECS/GEM	Every-run monitoring, operator assistance, intervention where appropriate, and models that seed the next tool, recipe, or site

Three of four capabilities show results before a tool is connected

The adoption path in three phases, in order, not to scale




Understanding, measurement, and detection all show results in phase 1 on historical data. Action waits for phase 3, once the other three are trusted.

At the end of phase 1, the decision to go further rests on results from your own process. Atomscale connects alongside the instruments you already have and replaces none of them. It runs in the cloud, through an API, or on premises for facilities whose data can't leave the building, and process data stays inside the deployment boundary you agree. Pricing is usage-based, tied to the data processed, with no large upfront commitment.

Value Compounds as the Capabilities Stack

Each capability reuses the models built for the one before it, so the returns add up. The state estimate from measurement is what understanding explains, the learned dynamics are what detection references, and all three are what action steers with. Adding a capability puts models the team already has to new use, and every run on a connected tool improves them.



Illustrative. Each capability starts once the one beneath it is established and widens with every run; trial and error stays flat.

CAPABILITY	BENEFIT	WHAT CHANGES	EXAMPLES IN PRACTICE
Measurement	Shorter campaigns; yield at spec	Weeks spent waiting on characterization; properties held in window	10x lower composition error, from 2 reference runs
Understanding	Engineering time	Designed experiments avoided; experiments aimed at the real process; knowledge retained	A 75% lag error found from the archive in days, with no new wafers
Detection	Yield	Runs lost per excursion; experiments corrupted by undetected drift	A change caught 40 s ahead of manual observation; >96% accuracy on non-uniform wafers
Action	Runs-to-spec; new programs	Runs to qualification; materials worth attempting	90% of growth outcomes predicted from ~10 samples

A WORKED EXAMPLE

Take a program running twenty five growths a week, with a two-week characterization turnaround and two excursions a year that each run three weeks before anyone sees them. Mid-run detection caps each excursion at one run where it would have cost seventy five, which returns 148 runs a year. Live measurement takes most of the two-week wait out of each learning cycle, and understanding replaces one designed experiment a quarter with an analysis of runs already paid for. None of this counts the program that becomes possible, which is where Part 1 put the larger return.

IN THE FIELD

The Teams Moving First

Atomscale is already in use with teams in photonics, quantum computing, and power electronics. Our work spans complex oxides on silicon, including BTO, STO, ITO, III-V compounds including GaN, GaAs, InP, SiC, two-dimensional channel materials for next-generation transistors, quantum cascade lasers, and ultrathin magnets.

It began on molecular beam epitaxy, where some of the most demanding materials are grown, and has moved across deposition methods since. The results in this book come from MBE, MOCVD, and ALD, and the platform also supports PVD and sputtering, with atomic layer etch in development.

Starting early matters because the returns compound. A team that starts now has models learning from every run it makes, process knowledge that stays when engineers move on, and a next tool or site that starts partway up the learning curve. A competitor that starts a year later begins that curve a year behind, and the distance grows with every run. Our thesis made the same claim: the teams that embrace this are the ones that will lead the coming wave of materials innovation.


HOW TO START

One tool. One question.

Pick one tool and one question from the list at the start of this chapter. Send us a historical export from that tool, with the matching metrology if you have it. Within days we will show you what your data has been holding, at no cost, and you can decide on evidence from your own process whether to take the next step.

Start a free pilot

ATOMSCALE.AI/CONTACT



ABOUT US

Frontier intelligence for atomic-scale engineering.

We are materials scientists, engineers, and technologists brought together by a shared mission to accelerate technological progress by enabling atomic-scale material innovations that address the most pressing challenges of our time.

Having experienced the challenges of material scale-up firsthand, we built Atomscale to overcome the physical and informational limitations of today's atomic-scale engineering toolchains.

Atomscale was founded in 2025 by Chris Price and Jason Munro and is based in Boston, Massachusetts.

[ATOMSCALE.AI](#) · INFO@ATOMSCALE.AI



Intelligent Atomic Scale Engineering.

