QUANSCIENT

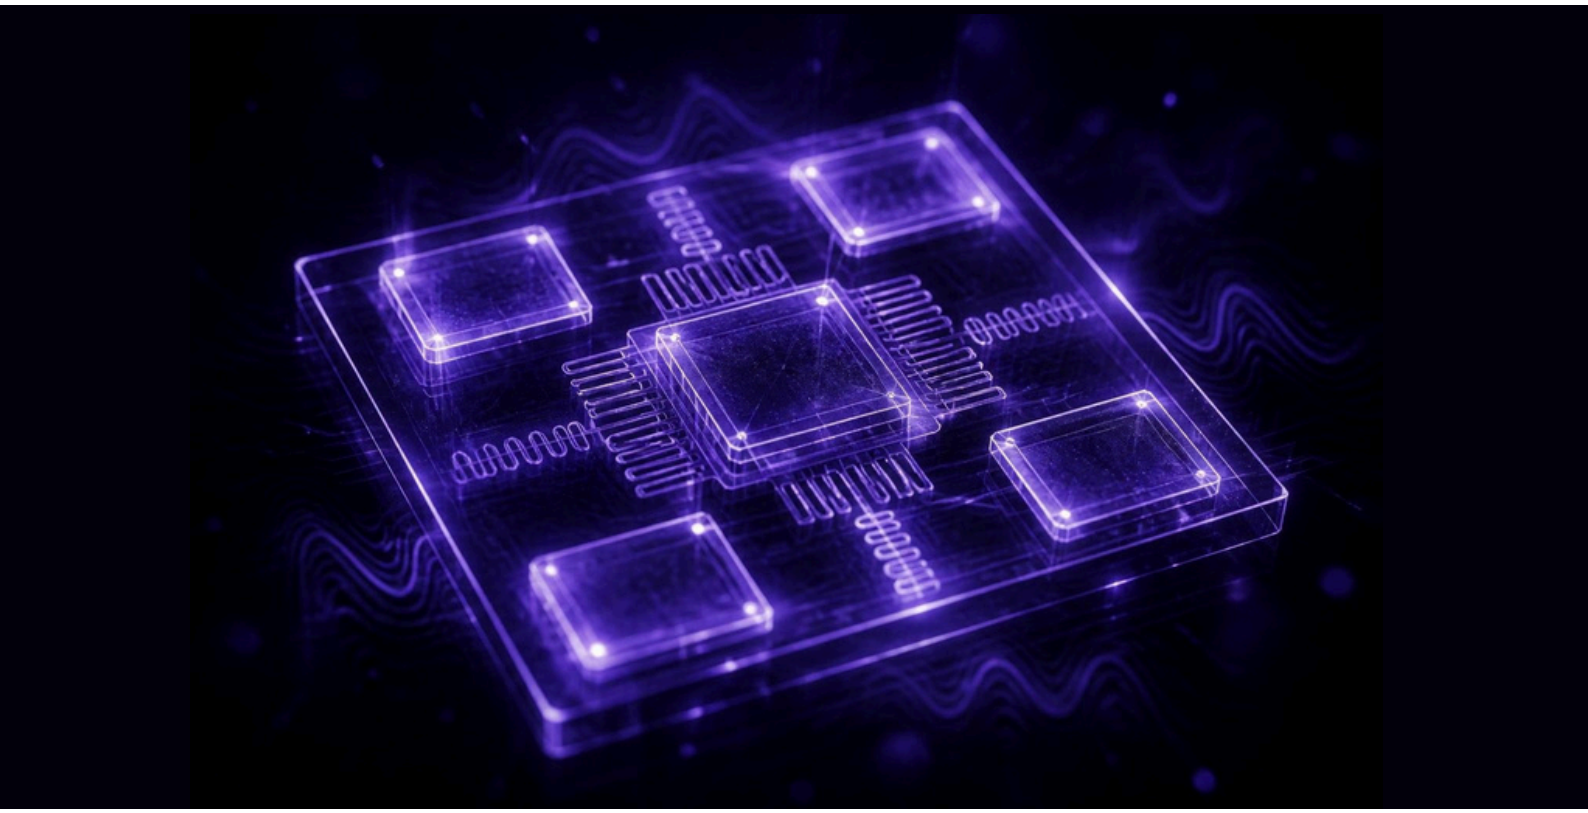

E-book | September 2026

Tuomas Eerola

An agentic workflow from requirements to simulation to reporting

A new way to solve MEMS packaging design challenges.

Contents

Introduction to Quanscient Allsolve 3

The problem in MEMS packaging 4

The solution 5

A further opportunity 8

Conclusion 9

Key takeaways 10

Get in touch 10

Introduction to Quanscient

We equip engineers to rapidly explore thousands of design options for the world's most complex challenges — to refine to an optimal solution with precision and confidence.

Learn more at quanscient.com ⟶

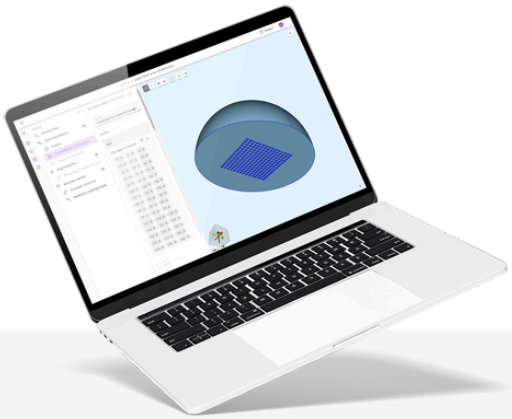

Quanscient

- Quanscient was founded in Tampere, Finland in 2021 to improve accessibility and performance in engineering simulation.
- Quanscient Allsolve, a cloud-native multiphysics simulation platform was created by combining expertise in algorithms, cloud development, software business and quantum technologies.
- Continued growth is driven by platform innovation, team expansion, digital product development on native quantum computers, and Allsolve Agent, which automates the entire workflow from product requirements to scalable simulation and reporting.

Trusted in both industry and academia

UK Atomic Energy Authority · infineon · Massachusetts Institute of Technology · Proxima Fusion · imec · PAUL SCHERRER INSTITUT PSI

Tuomas Eerola

Tuomas has extensive experience helping businesses address digital product development challenges using engineering simulations, cloud-powered high-performance computing, and AI. Since 2025, he has been leading sales at Quanscient, supporting the adoption of Quanscient Allsolve, a cloud-based multiphysics simulation platform.

The problem in MEMS packaging
Complex requirements and only 24 hours in a day

Packaging a **MEMS device** is not just mechanical protection. The shape and structure of the housing, the viscoelastic behavior of adhesive layers, the coupling between thermal expansion and mechanical stiffness, the sensitivity to air gaps and layer thickness - **all of these affect the device's final performance** at the same time.

On top of the performance targets that are defined by the final application, engineers also need to recognize the possible limits set by standards and regulations. And in practice, there is often a good reason to favor materials already used in production, because it helps with sourcing volume and inventory turnover.

In practice, packaging teams keep running into the same shape of question. Why does a part's offset drift after assembly and reflow, how long must a part sit before it is stable enough to trim, and which process tolerances actually move performance enough to be worth controlling. These are representative examples, not universal ones, but they point at something real: these decisions often rest on a handful of data points and prior experience, because a study that would actually settle the question never gets run at the scale it would need.

When there are several such variables and constraints, not just two or three, the traditional approach stops working. That approach is to simulate a few options by hand and pick the best one. Already if each variable is tested at only five levels, the number of combinations quickly reaches hundreds of thousands, or even millions. Running a million simulations by hand is not realistic. If you would run them, your software licenses and servers would be busy for a long time. And even if you could run them, there are not enough hours in the day to read through the results.

In practice, this shows up in the engineer's daily work in the wrong order. Most of the time goes into setting up the model, pulling inputs from different systems (PLM, ERP, standards documentation, files on network drives), preparing input data, and putting the results together. All of this takes time away from the analysis and design decisions that actually create value, help find the winning and profit generating product design.

The solution
Agentic AI builds and runs the DoE

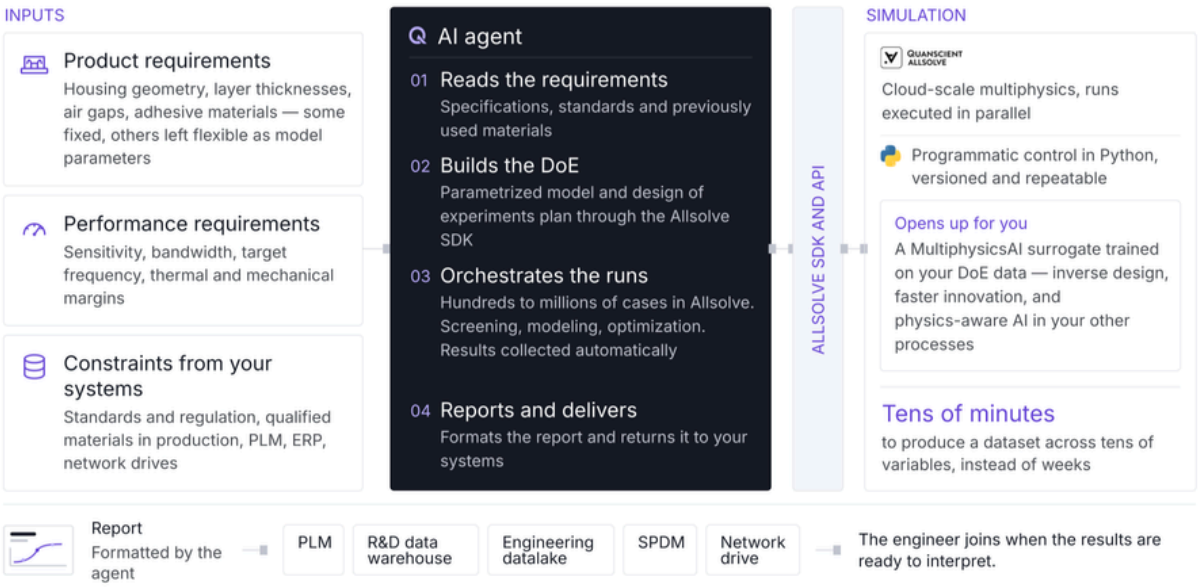

Quanscient Allsolve offers a concrete solution to this. Allsolve has a novel architecture that is designed for scalability, and it comes with a game-changing API (Application Programming Interface). An **AI agent** can read the performance requirements, application specs, standards, materials used in earlier products, and previous findings. From this, it builds a parameterized simulation model and a Design of Experiments (DoE) plan through the Allsolve SDK's (Software Development Kit) Python API.

The runs execute in parallel in the cloud. Results are collected automatically. At each step, from model setup to the finished report, there is a checkpoint a person can inspect before the next step runs. The engineer is no longer in the run loop, launching one simulation at a time. They step in at the point where the value-adding analysis and design decisions can begin.

The solution
Agentic AI builds and runs the DoE

Picture a die bonded to a package with an adhesive layer, as one example of what this looks like in practice. Under a thermal load, the anchor point can shift, and that displacement is a useful proxy for how the change will affect device performance. An agent can sweep bondline thickness, temperature, and adhesive stiffness across a run count that fits the question, anywhere from tens to hundreds, even to thousands, of transient runs. Each run returns a full relaxation curve over time, not just a single point. That is what lets the results show how much die stress moves per micron of bondline thickness, and how long a part takes to settle after a thermal excursion, including the fact that a cold part tends to settle more slowly than a hot one. None of that shows up in a single-point simulation, only in one that tracks the full curve.

The right scale depends on the case, and on how many variables and constraints are involved. Sometimes a hundred runs are enough. Sometimes you need ten thousand. Sometimes a million. Quanscient's solution scales to match, because the computation runs in the cloud instead of on one workstation. Allsolve is a cloud-native simulation platform, and its solver technology is built for scale, so even a complex simulation campaign finishes in a realistic time. In projects we have done, running a full DoE across dozens of variables has taken tens of minutes, not weeks.

The agent that we have introduced here works well when it has a clear question and enough context: the specs, the standards, the materials in use, prior results. It does not carry background knowledge of its own, so when something important is missing, it stops and asks rather than guessing. The engineer decides and signs off on the result before it is used for a design decision.

The solution
Agentic AI builds and runs the DoE

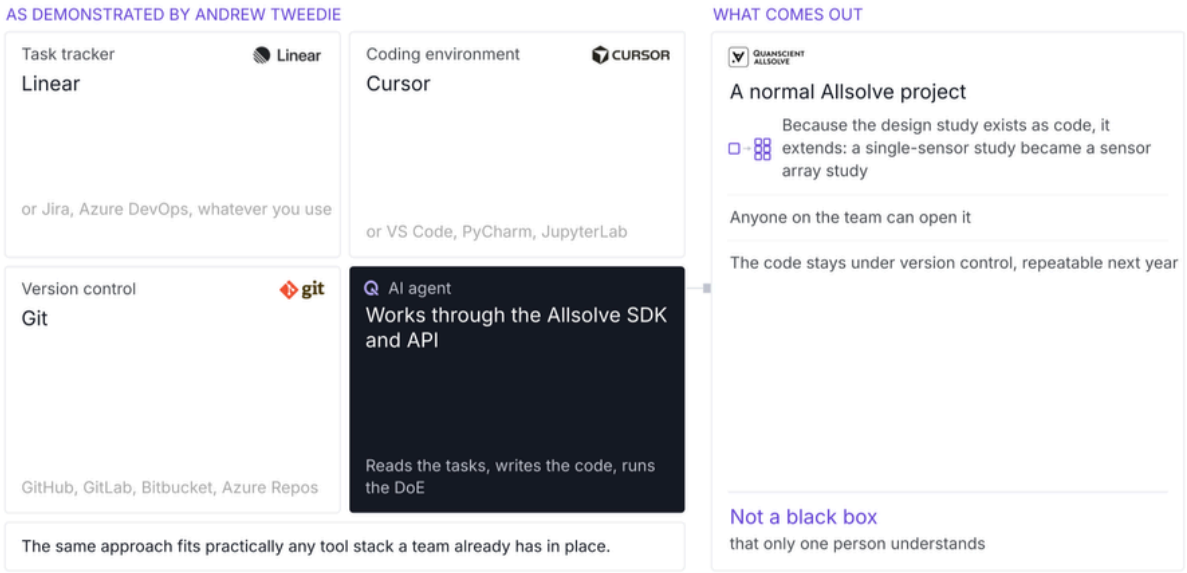

In **Andrew Tweedie's presentation**, he demonstrated this using Linear as task tracker, Cursor as coding environment, and the AI agent, and versioning in Git. This setup is just one possible combination of tools. Every organization has its own tool stack. The same approach fits practically any tool stack a team already has in place. What it produces is a normal Allsolve project that anyone on the team can open, with the code under version control, not a black box that only one person understands.

One of the measurable benefits is that it frees the engineer's time from running simulations and assembling reports to interpreting results and making value-adding decisions. A full campaign, one that would take weeks by hand or simply could not be done currently at all, finishes in hours or a day. This also moves product development into an "as-a-code" era. The whole chain is version-controlled and repeatable as code. That means the same model can be adapted for new use efficiently, or run again next year, without anyone having to remember what was done last time.

A further opportunity
MultiphysicsAI and AI-assisted discovery of new designs

With Allsolve, the data a DoE produces does not have to serve only one project. An organization can use that data to train a **MultiphysicsAI surrogate model.** Once a company has such a model, it can explore new design variants faster and at lower cost in its next similar projects. Once the model is trained, you can ask the question the other way around “I want to maximize these two performance metrics, show me the best design candidates.” The model answers by offering a Pareto front of possible designs, without a single new FEM run.

In one project we worked on, this kind of inverse search found options better than the original design. The sensor's bandwidth improved by about 35 percent and sensitivity by about 6.5 percent, while the target frequency stayed on point, all without a single extra FEM calculation after the training data was generated. That is a real gain in innovation power. The expert can pick the most interesting candidates directly and confirm them, if needed, with deeper analysis. As in this case, some of these candidates can even fall outside what intuition would suggest.

Measurable benefits is a radical time saving when exploring new design options, because each candidate no longer needs its own simulation run. The model can also surface design solutions that even an experienced engineer would not have thought to try. Because the data and the models stay the user organization's own property, every project and every DoE adds to their competitive edge, and makes the next one faster, not just this time, but for good.

MultiphysicsAI also opens up further possibilities worth exploring over time. A model like this could eventually serve people in the organization who are not familiar with FEM simulation at all, so a colleague – for instance a sales engineer outside the simulation team – could get a first, rough answer before deciding whether to bring in an expert.

“Because the data and the models stay the user organization's own property, every project and every DoE adds to their competitive edge, and makes the next one faster, not just this time, but for good.

Conclusion

The more physical effects, material choices, standards, and constraints are linked together in a design, the faster the traditional “simulate a few options by hand” approach becomes impossible. This is not only a problem in MEMS product development. The same holds for almost any complex hardware design where the number of variables grows past what a person can manage by hand. MEMS packaging is just one place where this shows up in a very concrete way, because materials, geometry, and physics are so tightly linked there.

Quanscient Allsolve solves this by combining three things into one chain: cloud-scale compute power, programmatic **(Python-based) control**, and MultiphysicsAI models. On top of these, agentic AI can build the whole path from requirements to simulation to reporting. It reads the requirements, builds the DoE plan and runs it, and ultimately helps an organization using Allsolve find better designs and serve its customers better, without every task requiring a full FEM simulation.

For an Allsolve user, this shows up in three concrete ways: a faster development cycle and shorter time to market, fewer physical prototype rounds and their cost, and the engineer's time freed from routine work for what their expertise is actually needed for: analysis and innovation.

A fair question for any packaging team to ask at this point: which of your own decisions today rest on fewer data points than you would like, and which one would you point to a workflow like this at first?

Key takeaways

- AI-assisted process shifts engineering time from setting up simulations to discovering new solutions directly from performance requirements.
- Cloud-native simulation enables thousands of multiphysics simulations to run in parallel, reducing campaign run time from weeks to hours.
- Cloud computing, AI, Python-based integration, and a powerful multiphysics solver enable efficient exploration of the design space under complex constraints.
- Once trained on DoE data, MultiphysicsAI surrogate models turn performance targets into a Pareto front of design candidates, without additional FEM runs.
- Agentic AI automates the workflow from requirements and model building to simulation, optimization, and reporting.
- Transparent, version-controlled workflows keep engineers in control while improving efficiency, reproducibility and quality.

Get in touch

Learn more and request a demo at
quanscient.com

Tuomas Eerola

VP of Sales
+358 50 336 7730
tuomas.eerola@quanscient.com

QUANSCIENT

quanscient.com

info@quanscient.com

linkedin.com/company/quanscient