

Solidigm Insights from CloudFest 2025


AI Takes Center Stage for Cloud Infrastructure

The CloudFest event, held in Europa-Park Germany every spring, is a terrific opportunity to get the pulse on cloud operator priorities and how technology requirements will shape adoption in the year ahead. The Solidigm team was on hand for the event, and we had the unique opportunity to survey the audience to better measure the state of the cloud in 2025. The following report captures the most pertinent takeaways from our data gathering across workload evolution, cloud oversight, and infrastructure demand.

2025 is a Cloud World...from Cloud to Edge

In discussing the demand for computing in the European market, operators remain bullish for continued growth of cloud instances but also see significant growth in edge computing complementing cloud operations.

“2025 will drive:”



This is likely due to the expected enterprise uptick in AI adoption with AI inference being targeted across this compute domain. In fact, while over 43% of survey respondents signaled significant growth in cloud computing ahead, another 40% clarified this position stating that both will grow in 2025.

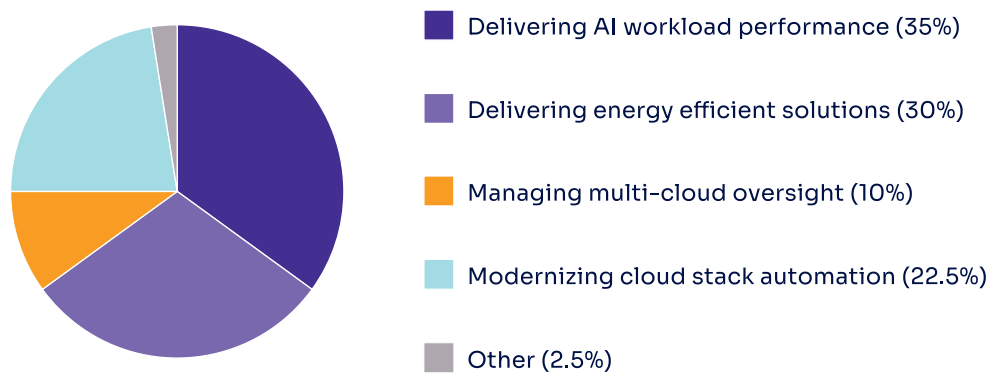
“A lot of AI has been consumed in the cloud in the last many years. But now, a lot of that is moving on-prem because cloud costs are high, and you have sensitive data that you can’t have in the cloud. So, when you don’t have access to that kind of cloud-based storage, you need that same kind of functionality in your own data center.”

- Thomas Jorgenson, Sr. Director Technology Enablement, Supermicro

AI is Challenging Cloud Performance and Efficiency

Unsurprisingly, AI dominated discussions at CloudFest with operators, customers and vendors discussing the urgency for infrastructure innovation to fuel the AI pipeline from training to inference.

“What is the most pressing requirement for cloud computing in 2025?” (CSP and End Consumer responses)



35% of operators and customers surveyed at the event signaled that delivering AI workload performance was the most pressing requirement for cloud computing in 2025, reflective of the accelerated computing capability demanded from integration of AI into workloads.

“Our unique model is we are on the full value chain, so we have the capability to build our own server, using two factories that we have, one in Europe and one in North America. And we are producing our own servers that we run in our own data centers. More than 40 data centers across the globe, plus our own network. All of this, we get the best performance-price ratio on the market for our customers. Currently, when we have changing demands, Solidigm is always there to help us find the solution, understand the storage technology and help us translate in our own, new solutions for the customers, which is great.”

- Guillaume Gojard, Product Director, OVHcloud

While AI performance continues to be a north star of focus, energy efficiency is a twin star, perhaps Alpha Centauri A and B, for many operators.

Notably, a full 30% identified efficient compute as the top challenge, and this, we believe, reflects the existential challenge of AI adoption for many of these operators, especially in Europe where EU power regulations are advanced and cost from the grid is high.

“Water cooling [...] is one among the many industrial innovations we were able to incorporate in our manufacturing processes. For example, water cooling with one glass of water, we can cool down one server for 10 hours of use. This is seven times less than the market, which is actually copied by many others. But this is a good trend for our own world and for the sustainability in general.”

- Guillaume Gojard, Product Director, OVHcloud

While Cloudfest featured companies demonstrating the latest in energy efficient technologies from liquid cooling to energy savings from storage and network technology alternatives, we expect to see more on this front at the upcoming OCP event later this month in Dublin.

The Data Pipeline Must Advance


With AI firmly established as the predominant workflow influence on cloud operator minds, one thing we wanted to dig deeper on is how operators were tackling a data pipeline to fuel AI algorithm training, fine tuning and inference. After all, 2025 promises to be a year of increased adoption of data hungry LLMs as well as a year where enterprises begin exploring data addicted agentic computing in earnest.

“What we have come out with of sorts in the past few years is we have our Petascale servers, and what that allows us to have, is a very high density of flash drives and obviously in that one, Solidigm is a key partner for us, to deliver that high-speed storage that can serve the AI models as fast as humanly possible. And then for the larger data sets typically it is stored in the slowest tier. So you have this multi-tiered storage set off for AI clusters. But one of the main things is that high density, Petascale kind of storage that we now have a huge demand for.”

- Thomas Jorgenson, Sr. Director Technology Enablement, Supermicro

Our suspicions in this space were right.

“Data pipelines:” (CSP and End Consumer responses)



- Require modernization of infrastructure to guide real time application performance (55.3%)
- Are ready for uptick in AI inference in the enterprise (26.3%)
- Require advancement to take advantage of distributed data stores (15.8%)
- Other (2.63%)

Over 55% of operators signaled that data pipelines require modernization of infrastructure to guide real time application performance, with only 26% stating that current data storage was prepared for the uptick in enterprise AI inference expected this year.

So how will data pipelines be upgraded, especially within the performance tier? The industry has its sights on SSD innovation, including the introduction of increased density drives that provide improved pipeline latency while saving a lot of energy to fuel accelerated compute.

“Size, size, size. Capacity is everything. Solidigm has the QLC product and it’s a firm favorite in the market. The 61TB, now 122TB, it’s massive and we are seeing a huge, huge pull for this product.”

- Charlie Hacker, Sales & Marketing Director, M2M Direct

We are at the beginning of the adoption curve for the latest 122TB NAND drives in market with the first solutions becoming available through broad distribution right now. Over 30% of operators have prioritized migration to SSD technology as part of their modernization approach.

Solidigm Takeaways

With CloudFest in the rearview mirror, we’re delighted to see that the cloud is thriving in Europe, and service providers in the region are at the tip of the spear in delivering efficient and performant instances to the market. With the complexity of data sovereignty in the European landscape, we expect to see more about cloud-to-edge computing solutions that navigate regulations and deliver performance to fuel AI inference for multi-national organizations. We also expect to hear customer adoption stories of new infrastructure from accelerated compute to new data pipeline adoption to competitively deliver services to the enterprise and see our industry leading 122 TB SSDs as a critical element to building out the data pipeline to serve this important market demand. For more information about our engagements at CloudFest and our entire portfolio of leading storage solutions, please visit our website.

About Solidigm

Solidigm is a leading global provider of innovative NAND flash memory solutions. Solidigm technology unlocks data’s unlimited potential for customers, enabling them to fuel human advancement. Originating from the sale of Intel’s NAND and SSD business, Solidigm became a standalone U.S. subsidiary of semiconductor leader SK hynix in December 2021. Headquartered in Rancho Cordova, California, Solidigm is powered by the inventiveness of team members in 13 locations around the world. For more information, please visit solidigm.com and follow us on Twitter and on LinkedIn. “Solidigm” is a trademark of SK hynix NAND Product Solutions Corp. (d/b/a Solidigm).

