


# „Thanks to the cloud, providing a big iron infrastructure is no longer an issue“



There are many good reasons in favor of the cloud, one being that it promises low initial investments. In addition, large amounts of computing power can be booked by the minute. For example, we can test special AI applications on short notice. Such opportunities are unthinkable with an on-premise solution, as the procurement process for an additional server alone often takes several months. What are the advantages of data management from the cloud and what aspects should you be careful about? We asked these questions in a background discussion with the employees of SCALE GmbH (Ingolstadt, Germany): Dr Martin Liebscher, Managing Director, Torsten Landschoff, Software Architect, and Eric Zuchantke, DevOps Engineer.



Martin Liebscher

**Dr Liebscher, where do your clients stand with regard to data management provided by the cloud?**

**Martin Liebscher:** Most automotive OEMs are still at the beginning of their transformation process to cloud computing, and especially to the use of public clouds. They are very open to this prospect, but still hesitant, especially with regard to aspects such as IP protection. Although the starting signal for a cloud-first strategy has already been given in many cases, there are still some fears that need to be overcome. It will be a few years before we see productive installations for CAE data management in the cloud on a large scale.

**To put it bluntly: When did SCALE go to the cloud?**

Our cloud-first strategy goes back more than six years ago. It evolved as part of the generational change of our software products. We have already published a story on this in digital AGENDA (1). To some extent, it was also a stroke of luck that the automotive OEM Audi/VW had decided on a cloud-first strategy around the same time as we were also driving our development in this direction.

**Often, cloud strategies contain phrases about protecting investments, meaning that money already invested in IT does not have to be written off when migrating to the cloud. What is your perception of the market's goals?**

**Torsten Landschoff:** Interestingly, we serve companies in Japan that would prefer to do everything on-premise. But that is more of an exception. The fact is that we have to meet a wide range of client requirements with regard to cloud connectivity.

**Martin Liebscher:** At the same time, we see that almost every client is now asking for access to the cloud.

**Torsten Landschoff:** You can also put it the other way around. For us, too, it means protecting investments if our clients want to migrate to the cloud sustainably. After all, this involves considerable effort in developing the appropriate technologies.

**Martin Liebscher:** We can name clients who are already actively operating data management in the cloud. Moreover, the vast majority of clients who are not yet in the cloud have given us a clear signal that they want to migrate to the cloud sooner or later.

**Torsten Landschoff:** Cloud technologies also have a crucial advantage. They make us independent of our clients' hardware requirements or virtualization measures. This is a big issue with an on-premises installation. The associated, sometimes extensive measures before the client can try out the system are completely eliminated with an installation in the cloud. The system is put into use with a test data set at the push of a button. There are practically no costs for testing, which makes it easier to onboard the client.



**Martin Liebscher:** It is also the strategy of some of our clients to first test a system in the cloud and then consider which parts they want to operate in a hybrid mode, for example because extremely sensitive data is to be used or because cost calculations play a role. In general, you can assume that running a simulation data management application in the cloud is cheaper if you keep an eye on the overall costs. And that's what clients expect. In the long term, however, it should be possible to store and edit all data in the cloud.

***Is a cleverly chosen cloud strategy good for gaining a competitive advantage?***

I am convinced of that and think we are very well positioned in this regard. We have been dealing with this issue for a long time and our products are largely designed for native cloud operation. We are already thinking about the next big steps, but maybe more on that later.

***What has changed compared to the previous system architecture?***

**Torsten Landschoff:** Until six years ago, our solution was a monolithic system that ran on a single server. The problem was that the server had to be dimensioned according to the number of users. Load balancing over multiple servers was not possible. As part of the revision of our system architecture, we split our system into smaller services that Kubernetes can then scale depending on the load.

***Kubernetes?***

Kubernetes is an open source management tool developed by Google, which is used for the provision, administration, and monitoring of so-called container environments. Unlike virtual machines, containers do not emulate the hardware, but act as an operating system. Virtualization is therefore achieved at a higher level without the otherwise necessary hypervisor (VM Monitor for creating and running virtual machines). We can easily run additional containers depending on the load. This enables dynamic scalability. We are therefore no longer dependent on a 'big iron', i.e. a large server. This allows us to move more and more use cases to the cloud.

***A noble goal, isn't it?***

**Martin Liebscher:** In addition to the post processing phase — where typically terabytes of data are generated — simulation data also includes model data, information that has its origin in the preprocessing phase. Here we are pursuing a hybrid approach: The client can still work with a rich client, but the data is synchronized with the cloud in the background. This gives the user a 'local' look & feel ('user experience') when creating the CAE model, for example, like they are used to.

**Torsten Landschoff:** We want to stick with the desktop client for some use cases, for example for model editing. This is also what our clients expect. In the long term, however, all data should be stored in the cloud and edited there.



Torsten Landschoff



Eric Zuchantke

***An important point here is that SCALE's cloud computing strategy changes little in terms of the user experience. Right?***

Definitely right! Because the costs of changing the user interface are enormous. Not just for us, but also for our clients, for example in terms of training costs. You have to be very careful with any kind of GUI modernization ...

***... it goes without saying that customer experience is what sets a system provider apart — or not. You certainly don't want to fall behind in this area.***

**Martin Liebscher:** That's why we offer both a desktop and a web client. There are also use cases where a local application is necessary for viewing and processing data. For other tasks, such as statistical investigations — keyword here is 'data analysis' — with AI algorithms, the cloud with its standards is the first choice. The different use cases define our transformation process for the client.

***What can you report on the electric vehicle startup Rivian Automotive, SCALE's first client in the cloud?***

Amazon is an 18-percent-shareholder in Rivian. At Rivian, we have implemented simulation data management in the cloud in its purest form. The experience has been very good because we received very good support from our partner AWS. It is always something special to work with a (now former) startup.

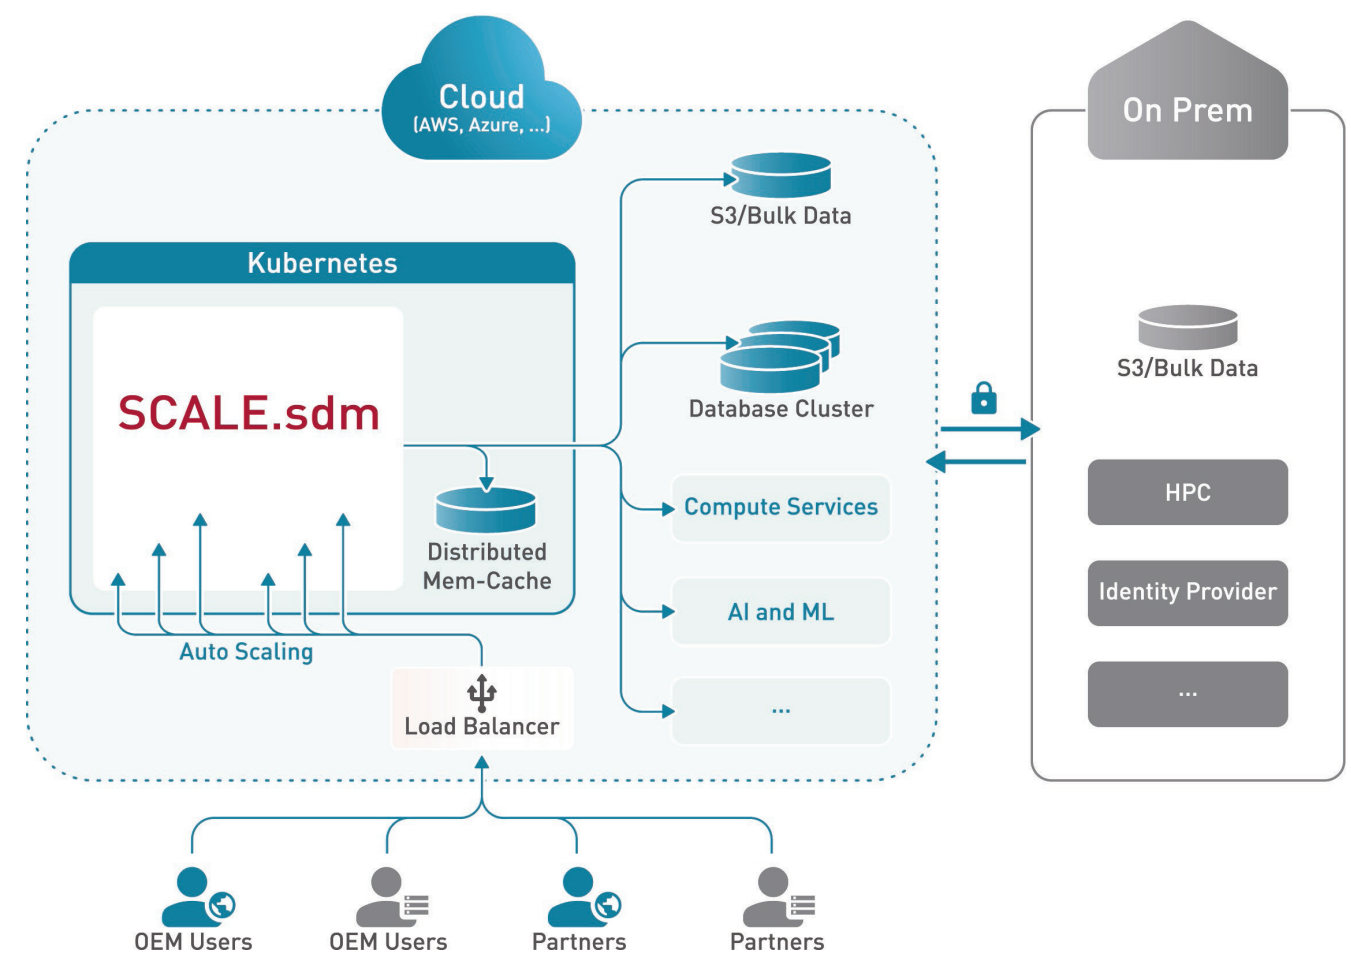
To put it briefly: The experience was very good because after a three-month proof of concept we made a clear decision in favor of SDM from the cloud. However, the deployment remains a challenge. Now, the productive business relationship is developing and it remains exciting, especially since the Volkswagen Group has invested 5 billion US dollars in the company.

**Eric Zuchantke:** For an innovation-driven startup, our cloud offering is very attractive. We deliver the data management system in its entirety from one source — including an automated installation option. The actual commissioning is very quick. The total effort covering 14 days is mainly required because we have to clarify an IDP and other company-specific IT issues for basic deployment, especially for user registration.

It is important to observe the different compliance rules. Stricter compliance rules in particular help us to sharpen our processes. We are still going through a certain learning phase. Basically, however, deployment is very similar for every client. Two scripts are used: The first initiates the infrastructure in the cloud and the second deploys our applications in the cloud.

***Does that have something to do with infrastructure-as-a-code (IaC)?***

Yes. IaC refers to the ability to provide and support the required computer infrastructure using a software code instead of manual processes and settings. With the automated IaC approach, the hardware as a whole and the associated network infrastructure are defined as code, which is then configured at AWS, for example. If you like, we also deliver the virtual hardware. This setup can of course also be adapted to the client's needs.



Cloud system architecture of SCALE

Source: SCALE 2024

And what do your clients say about it?

Clients welcome this type of deployment as something very positive because they have never encountered it in this form before.

What happens next?

Once this is complete, it's time to fine-tune the infrastructure. We check whether the microservices provided match what the client really needs. We also monitor the system in terms of utilization and the associated costs — depending on what the client is willing to invest. Remember, hyperscalers offer autoscaling. That's a good thing, because if more resources are needed, they are made available immediately. But the client also has to pay more for them. That's why certain limits are set for each deployment.

That means you always have an eye on which resources are actually being accessed.

Yes. For example, I check this at a time interval that the client specifies. There is also the option of defining alerts that are triggered when cost limits are exceeded. The same applies to alerts as an indication of failures. These types of services are subsumed under the term 'cloud governance'.

SCALE supports not only AWS, but several hyperscalers. Is that a challenge for you?

Not at all. We opted for a management system in the cloud for our system architecture, which often means not using the cloud directly, but using

Kubernetes as a middle layer in between. We can implement our software on the various hyperscalers and the user doesn't even notice which cloud it is running in. This abstracts the cloud infrastructure actually used. We didn't forget to enable high-performance cloud-to-cloud and cloud-to-on-premises connections, for example via VPN tunnels. It is important to take into account that data transfer within a cloud is free to inexpensive, but data transfer costs arise from cloud to cloud. This can make a hybrid mode unattractive.

Can you actually leave the cloud again?

This is not always easy. The way back to on-premises can be made more difficult if you access very specific services in the cloud that are not available locally.

Got it. Let me put it this way: Going to the cloud is easy. But the road back, for example due to unexpectedly high costs, can be rather difficult.

It can mean that you spend weeks and months on the migration back. Of course, the hyperscalers would like to exploit this vendor lock-in effect. At least we are on our guard here and support our clients so that they do not fall into a trap. Prices can change overnight. Our message at this point is: There are solutions for this too.

Dr Liebscher, finally: What does SCALE's cloud roadmap look like?

**Martin Liebscher:** The topic of cloud is established in our products and the SCALE cloud is publicly accessible. We have completed our tasks with regard to IP security aspects. The SaaS business model is also a target on our roadmap. This is because the clients no longer have to operate their cloud account themselves, but can concentrate on calculation and evaluation. Incidentally, our SaaS model is in line with the hyperscalers' strategy.

We also follow the PaaS approach. The German Fraunhofer Institute SCAI or companies like SIDACT, for example, implement their machine learning tools as extensions for SCALE.sdm and thus use our platform. We have disclosed the interfaces for this. The clients can commission their own — even proprietary — extensions from us or third parties; we provide documentation and support for this.

The AWS Marketplace, which is basically the same as the Play Store for Google apps, also attracts our interest. We can well imagine making SCALE.sdm available as a data management app in that store. The work on this is already well advanced. This could be of interest especially to smaller engineering service providers.

Gentlemen, thank you for the interview!

Interview: Bernhard D. Valnion

References

- (1) Valnion, B. D., „Simulationsdatenmanagement und CAE-Prozessdesign weiter denken als bisher“, d1g1tal AGENDA 2/2022, S. 18 ff., Baden-Baden

www.scale.eu