

David: Diego Tamburini is the principal industry lead for Azure Manufacturing in the Microsoft industry experience team where he focuses on developing technical content to help manufacturing companies and software developers deliver their solutions on Azure at scale. Today, we're talking with Diego about accelerating innovation in manufacturing by moving our compute to Azure. Welcome to the show, Diego.

Diego: Thank you, David. It's great to be in your podcast.

David: What is it that we're really talking about when we talk about moving compute to Azure, because it seems like so much of that stuff would need to be on site?

Diego: What we really mean is basically move the computationally intensive workloads from your hardware in your desktop to the hardware in the cloud. The reason that we will do that, and why it's so important for manufacturing, and when I say manufacturing I also include engineering and design, of course, is that engineers, they routinely use applications that are very demanding computationally. The applications such as 3D modeling tools, rendering, photo realistic visualizations, and also analysis and simulations, finite element analysis, the discreet even simulations. These are workloads and applications that are very demanding computationally, and if you do them on your hardware, even if you have a souped up machine, you're kind of limited to what the hardware you have in your laptop or your desktop.

David: Where with an HPC or a high performance computer, you could scale up pretty easily by adding cores, and memory, and things like this, right?

Diego: Exactly. When you need that and for the amount of time that you need it. No more, no less.

David: Yeah, pay for what you use.

Diego: Right.

David: Cool. Well, I understand that manufacturing is a whole industry into itself, but I'm really interested in the 3D rendering that you slightly mentioned. Tell us how that fits into the story.

Diego: The 3D rendering requires computers or hardware with GPU, basically. So the more GPU resources that you have in your machine, the faster that renderings will take place. Now, there are, in some cases, folks that rendering very complex scenarios with lights and materials that take a considerable amount of time, either for 3D models or animations and things like that. The cloud is ideal for that because we have, in Azure, we offer virtual machines that are specifically optimized for graphics. They are GPU enabled, and you can recruit as many of them as you need at a given time. So rendering is one of the workloads in design that went to the cloud first, just because it's so massively compute intensive.

David: So I can have my desktop application, say, render my model. It will go out to the cloud, render it out through the GPUs on a high performance computer, and then give me back the rendered image or rendered model.

Diego: That's exactly the way it works, yes.

David: That's pretty cool. Okay, well, manufacturing doesn't necessarily sound like a ... Oh, we've already touched that. Nevermind. I'm sorry. What are some other problems in manufacturing that Azure is helping to solve?

Diego: From the compute perspective, the one that is gaining a lot of steam, of course, is anything that has to do with IoT. IoT, the ingestion of data itself does not require a lot of compute resources, but the analytics and specifically when we're IoT data to train machine learning models, that does require a lot of compute resources. So that was one workload that basically started in the cloud. You stream data to IoT data, sensor data to the cloud, and there in the cloud is much more effective to store it, to analyze it, to train machine learning models.

David: Often, many of that ... Let me start over. Often, much of that Internet of things data that gets reported, gets reported in time series data, which in and of itself has some idiosyncrasies that require us to use even more of that HPC to ensure that we're getting things in the right order.

Diego: Correct. Yes, the natural place to ingest and analyze IoT data is the cloud. It will be very, not impossible, but very difficult and inefficient to do it elsewhere.

David: Okay, I'm going to pause. What else do you want to go into?

Diego: Probably the ways ... So the type of applications that do compute in the cloud and the ways to scale it. Probably the ways to scale it first, and then the type of applications.

David: Oh, because we haven't talked about Batch or anything.

Diego: Right, correct.

David: Okay. Just making some notes. Okay. So we talked a little bit about that desktop application that renders our 3D model, which is really cool, but what other types of apps work really well with compute in the cloud?

Diego: Hold on, let me ... Rephrase the question because what ...

David: That's the one where you told me you wanted to talk about the types of apps that do compute in the cloud.

Diego: Okay, yes. Go ahead. Okay, I'll start that. When it comes to compute in the cloud, there are basically three flavors of how applications can do compute in the cloud. Number one is the one that we touched upon a little bit earlier. You have a desktop solution. The

user interfaces with that desktop solution, and say you want to render or analyze a 3D model. The application ships their compute intensive processed to the cloud. It's processed in the cloud, and then you get the results back in the desktop solution. So that's number one.

Number two are solutions that run on a cloud virtual machine and are accessed by users via remote access solutions. That's very popular too. That's how a lot of rendering and some of finite element analysis applications run in virtual machines in the cloud. Then there are the cloud native solutions. Everything happens in the cloud, and the user interface through a browser type of interface. Compute and storage, everything happens in the cloud.

David: So that is more of a SaaS model, right, where somebody's going to come up to a pre-existing solution?

Diego: Yes, correct. The third one is the more SaaS solution. Correct.

David: We talked a little bit already about scaling up with an HPC, but what about the story of scale? Are there more ways to scale given the workloads that we're getting?

Diego: Yeah, there are two basic ways to scale, vertically or up, and horizontally or out. Vertically is basically, you have your solution in, say, a virtual machine, and you move it to a more powerful VM, a VM that has more GPUs, or more RAM, or faster storage. In Azure virtual machines, specifically, we offer them in many configurations including virtual machines that are optimized for GPU and HPC applications.

The other way is horizontally. Horizontally there are, in turn, two main ways. You can create a cluster of VMs. So when you find a specific machine configuration that you like, you scale it using virtual machine scale sets, like a cluster of virtual machines. This is a group of identically virtual machines that are load balanced and can be deployed in minutes, basically.

The other way to scale horizontally is that you spawn compute processes. This is more what we have in mind when we think about parallel computing. The way you do that on Azure is through the Azure Batch service. So with Azure Batch, you can run large scale parallel and high performance computing applications more efficiently. So it lets you schedule HPC in the cloud, and you can have access to 16 cores or 100,000 cores when you need them. Also, it gives you a consistent management experience of scheduling jobs, and whether you select Windows server or Linux compute nodes, etc. So again, vertically by getting a more souped up VM, and horizontally by clustering VMs and/or using Azure Batch.

David: One of the things I want people to make sure and do after they're done listening to the show is go download the use case that you'll find linked in the show notes that underlies this problem that Diego and I have been talking about. You can take away some more information from that use case document. Diego, where are some places people might go to look for partner solutions. As Microsoft really the infrastructure provider, we

depend on our partners to deliver full baked solutions. Where would people go to find those?

Diego: The perfect place to find a solution is on the Azure marketplace or on the Azure app source. Like I mentioned before, you'll find ... It'll probably help you if you keep in mind that you'll find these three flavors of solutions that are fully on the cloud, SaaS services for analysis or simulation. There are also solutions that give you kind of a remote access to souped up virtual machines in the cloud, and then there are the ones that basically enhance your desktop applications with analysis or rendering in the cloud. So look for these solutions on the Azure marketplace or the Azure app source.

David: Also, be sure to visit our manufacturing homepage, which we'll also have linked up in the show notes. Okay, that's a good pause. Is there anything, Diego, you wanted to touch?

Diego: I don't know how to touch the choice of cloud. So probably you can ask after I mention the three types of solutions. Can you move stuff around, I suppose? If we record something now, you can ...

David: I can.

Diego: So after I explain that there are the desktop solution, the VM solutions, and the cloud native solutions, you can ask me something about who chooses the cloud. I can say, "Sometimes the ISV, but it should actually be the customer."

David: All right. I'm going to have to mark that. Hey Diego, who typically chooses to move to cloud in a given organization?

Diego: I'm sorry, the question will be who chooses the underlying cloud for this ...

David: Oh.

Diego: Right.

David: So Diego, who chooses the underlying cloud that we're going to build our solution on?

Diego: Sometimes the software provider will make that choice, either because they host their solution on a given cloud, or they provide connectivity to a given cloud, but of course the better software providers enable their customers to make that choice. Customers may have a preference on a given cloud. Say for example they already have an enterprise agreement with Microsoft and they already have Azure credits. So they will want to be able to use those Azure credits as opposed to pay again for another cloud because the ISV didn't give them the choice.

So I will say general for ISVs, it's a good strategy to provide a choice of clouds because some customers are going to need that. For the customers, make sure that you're not paying for cloud more than one. Your organization may already have an enterprise

agreement with Microsoft, and you already have cloud credits that can be used for compute.

David: So you're not advocating for explicitly sticking to one provider, but a hybrid cloud solution might be better?

Diego: For the ISV, yes.

David: Yeah. Well Diego, I tell you what, we're just about out of time here, and I want to thank you so much for joining us on the show today. All of those resources that we mentioned are going to be linked in the show notes for the episode on our podcast pages. I can't thank you enough, Diego, for joining us today. It's always a pleasure.

Diego: Thank you, David, for hosting me in your show, and to your listeners for spending the time with us.