

David: Welcome to the Microsoft Industry Experiences Team Podcast. I'm your host David Starr and in this series you will hear from leaders across various industries discussing the impact of digital disruption and innovation, sharing, how they've used Azure to transform their business. You can find our team online at [Aka.ms/indep](https://aka.ms/indep) or on Twitter at [industryxp](https://twitter.com/industryxp). Nick Lima is a principal industry lead for insurance at Microsoft. He advocates for Microsoft Azure within the insurance industry and is a recognized thought leader in the insurance IT space. Today we're here to talk about evolving actuarial risk compute and modeling on Azure. Nick, can you tell us what we're really talking about here?

Nick: For insurance companies in general, except really a very highly regulated industry and one of the whole myriad of different regulatory bodies. So each state has its own insurance commissioners and they have their own set of regulations. Generally they're similar, there's a few states that are very different like New York and California, and there's also federal regulations on insurance as well, and then there are global insurance regulations on top of that and to sell insurance within different entities, you have to keep those regulators happy. To do that you have to be able to prove that if you sold an insurance policy an insurance policy had a claim that you had enough money to pay the claim and to do that you have to sort of model what if scenarios: what if the interest rate went up, what if the interest rate went down? If this person lived this many years, have I collected enough premiums to cover their benefit at the end of their life if it claims.

If it's for a fire or earthquake or whatever the same thing, you have to say, "All right, I have enough money so that I can justify and pay off a claim going into the future." So that's sort of the why or the what and the why is really as we move forward, those regulations keep evolving and there are new products that come in and the amount of data that is used to calculate that risk is also expanding. A couple of big things that have happened historically in Canada, [Likehat 00:02:27] was a series of calculations that were required for life insurance companies that added multiple, multiple scenarios on risk. So it really forced insurance companies to do up to 20 times as much work as they were doing before just for this one set of regulations.

And then in the US for some of the large insurance company C car, which was also a looking at risk for the entire company caused the insurance companies that were covered to quadruple the amount of work, if not more than that on top of these reserve calculations. Going forward a new regulation that goes into effect in 2020 is IFR 17. So this is for all insurance companies that's held contracts outside the US that are longer than a year. Generally that PNC, but any life insurance companies and some other non insurance entities have to do these new set of regulations which are, if not as complicated, but there's a new way of looking in calculation. So companies are finding that they have to update their existing systems for both processing data in and then reporting the information out so they're finding it easier to upgrade to a newer set of technology. So that's one of the things that are really pushing the market at this point. And then you start layering on all the additional data that's coming in.

David: If I could ask what solutions are they using today to solve for these problems?

Nick: Historically if you look back say 10 years, everyone was doing their risk calculations on premise usually on a mainframe system with a series of different software providers. Then as they're moving through time and you know, new data sources and new complicated systems, they realizing they could either keep expanding their on premise set of hardware or they could move to the cloud. With some of the new products that require stochastic calculations, which are just random samples of random samples and it can balloon up; either you have the runs take weeks and they have to be finished in a few days. So that's you have to scale up. So the big move to the cloud was how do you scale up and then you burst from an on prem into the cloud to use those additional resources.

That was the first step and then realizing, wow, I only need these resources for a short window of time. It doesn't really make sense to have this big group of hardware of on premises that I only use occasionally. Let's just only use it in the cloud. So companies are finding that it makes more sense to move to the cloud where the hardware is always updated, the security is in place and all the major software providers have options to run completely in the cloud or do a sort of a burst to the cloud type solution.

Speaker 3: Can you tell us the difference between bursting to the cloud and just moving my workloads to the cloud?

Nick: For many companies bursting was an easy way to get started in the cloud. They realized that to get a certain amount of runs done in a very short period of time was either to go through the expense and process of really building out their internal network to support the run that only needed to be there for a few days or do their normal work on prem. Then when they had to do these runs that at quarter end or year end, or that were needed to be done in a short period of time and to address these new calculations, all the data and everything would stay on prem. But then just the calculation piece would move to the cloud, do all the calculations, report back and then shut off. So an additional say 10 or 20,000 cores would be spun up and then run in the cloud and shut down and then the data would initially move up to the cloud and move back.

All in the cloud allows you to have complete flexibility. So maybe you only have a few hundred cores up all the time and then you can do your work, do your development, you can have your actuaries look at the systems and test it. But then when you need to do a run at month end or quarter end or for an audit, the first thing you need to determine is how much time do I have to do this run? And then you can scale up to meet that. So if you want it done in a few hours, maybe you ended up at a 10,000 cores or maybe you've got a two day window. So maybe you only spin up another three or 400 cores. That gives you some flexibility because what happens occasionally if there's an error with data input midway through the process was like, "Oh my God, I can't finish this in the time allotted." You shut it down, spin up an additional 10,000 cores and then spread that workload across more core so it finishes quicker and still meet your deadline.

David: Now when you talk about spreading across multiple cores, do you find more that they're using high performance computing with lots and lots of cores, or are we using Azure Batch to separate out the problem into minute pieces and solve that way?

Nick: Depends on the application they're using. Certain applications that are designed to work in batch and can scale up to a 150,000 cores. I think we have an example that's out there that we can do literally 150,000 cores they divided the workload into small pieces and spread it out across multiple regions and then in a few hours they're able to do a complete set of analysis and bring vital information back together again. It can also expand out, but you expand out to a fixed number and then the application actually distributes the workload, so it will finish within a specific set of time. So that the biggest problem when you split things out, if you split them out incorrectly and you've got one small piece it's taking four hours, you didn't split it up further, even if the little pieces took only a few minutes, that's still the long pole in the tent that's preventing it.

Some of the applications actually find those problems that take longer and then redistribute it as other pieces finished so that the footprint is always running as most efficiently as possible and can finish the shortest period of time. And that varies by the different solution providers.

David: I didn't realize they were using Batch in that way. So tell me if I've got this right, what you're describing: the customers are looking at the time that it takes for the batch run to complete and then there are tuning the algorithms to make it faster?

Nick: Correct. There's a couple of different approaches and like I said, it varies by provider. Another solution in this same space is as Azure Cycle cloud, which I've used previously before Microsoft acquired it and we were able to ... that whole scaling process even on an older set of technologies, it adds a whole nother level of control and flexibility and management through the whole process and that can be worked with applications that maybe aren't necessarily cloud ready at this poin. But using Azure Cycle cloud to manage that effectively non cloud type solution, it could work extremely efficiently and add all the control pieces that can keep your auditors happy.

Speaker 3: And Azure Cycle cloud is a resource to manage massive scale of course. But is that across VMs? Is it across containers?

Nick: It can be used for both. When I used it several years ago before Microsoft acquired it, this was across VMs, across multiple clouds and was ability to scale up and scale down. Now that Microsoft has acquired it, it can be used for massive scaling, but also in the 20 to 100,000 core. But I think where some of the real benefit can be in this smaller scaling between 10 and 5,000 cores, but it adds a level of management to the whole process so you don't have someone spin up 10,000 cores and leave them up over the weekend. It manages both of your spend, you can put spend limits on different users and then and manage whether a run has failed or not, and then bring the different components down so they don't stay up, which can be very frustrating as you spin up servers to meet a certain need and then you'll have some junior staff forget to turn things back off again. So it has some control features that can limit the spend as well as automatically scale things up and scale things down within a specific time period.

David: So I'm hearing that this technique of scaling my runs for risk analysis if do them in the cloud ultimately they end up being cheaper because, or rather I'm paying for what I'm

using as opposed to infrastructure that I'm keeping in racks and my own data center or something like that.

Nick: Definitely, from my own personal experience working at a couple of large insurance companies for like/unlike runs versus on premises hardware or private cloud, fixed private cloud versus a dynamic cloud in Azure, the difference can be up to 80 percent. The net tends to be closer to 50 because, especially actuaries see that they can get these runs done quicker the ease of use, but you end up doing a lot more runs. So the total cost savings is less, but they end up doing a lot more analysis and fine tune the results. There's a couple net wins there. You have greater confidence in the calculation so if there's an error, you do a lot of what if scenarios. So maybe I did it a run with a specific interest rate and then I did three or four and before when I did an on prem, I only had a choice to doing one.

I better pick the right set of scenarios going into this modeling process because I can only run it once on the equipment I have. But now that I'm in the cloud, I could run it 100 times and just keep making different modifications and fine tune to get the results that I'm looking for. We find that as people move to the cloud, they end up doing more runs and getting more fidelity and those results because of the flexibility. But the net save is there's still a significant net save and having that flexibility can be critical, especially if there's ever an issue with the data and you have to rerun it. I've seen this happen many times where as a process of pulling the data into the models occasionally there'll be a glitch along the way and it's not necessarily caught in the beginning.

So then do you run it with known bad data and then make an adjustment at the end or you stop everything and rerun it. With the cloud you have that flexibility. "Okay, I'm halfway through my run time. I can stop a bad run, reload everything, add more resources so it can still finish on time." That's a huge benefit. Especially if you look at the penalties that could be incurred for faulty calculations and the actuary staff's time as well. If you can make sure you get those done on time and have to do less additional adjustments after the runs are completed. It's a significant savings both from the total cost of the ownership of the run, but it's really how many hours overtime that you have an it staff and on actuaries in all these different things that end up happening if something crashes in the middle because of bad data, but you could actually fix it, rerun it, and then stay on target.

David: Nick, one of the things that's interesting about Azure is that we provide a platform for people to build their solutions upon, of course, and independent insurance companies are taking advantage of Azure of course. But so are independent software vendors who sell insurance facing solutions, where can people find those partner solutions?

Nick: Most insurance companies, they already have a solution that they're using now. And almost all those providers have a cloud options and we support those on Azure. So as people are working the first step is really work out with your existing partner for your actuarial calculations that can have a range of solutions from things that are a lift and shift into the cloud. Some of the bursting type solutions or even up to a SAS solution, which is a solution that you effectively just flip a switch and then you move the data into the cloud environment running on Azure. Then they do all the work of spinning up and

spinning down all the hardware behind the scenes and the interface that your actuaries and the people that are running the calculations stays the same as your on prem except you're doing it on the cloud.

It's all sort of behind the scenes. So that's the range is the SAS is probably the simplest way and the one that requires the most work from the IT side is really that sort of lift and shift IS solution. And there's a past solution which fits in between. That's an option as well. And each of the different providers will have those different solutions. And even if, for example, some of these older solutions that are not necessarily technically cloud ready could be almost that lift and shift, that IS solution that could leverage things like a cycle of cloud and really be a very dynamic cloud ready solution very quickly with these additional wrappers on the outside. It can make it look exactly like the environment you have, but only turn it on when you need it and then turn it off.

David: Can they also look in the Azure marketplace for existing cloud solutions for insurance companies?

Nick: Yes. That is an option. We are rapidly building out that ecosystem for risk computed and other solutions that support the entire insurance ecosystems. We have several solutions that are available now, like I said, as we build that out to really complete the entire ecosystem. That's our goal is to make sure if you have a IT solution need for an insurance company, especially in the risk side and very industry specific, we have those solutions listed at our marketplace and then how they can partner up with other solutions. For example, on the risk side, a lot of times they'll pair the calculation component for the liability piece with the asset piece, so it's that asset liability matching and a lot of times there are multiple vendors that create a different components that need to feed together for those answers.

One is strictly on derivatives and hedging calculations and one maybe just specifically on mortgages, so there's a lot of different asset pieces that feed into those asset liability matching. On the liability side is another set of calculations that maybe are just for valuations and other ones that are just for calculations and projections into the future and how do we make sure they all fit together. So we're creating a marketplace so that you can pair up those applications. Usually a lot of those people already have those different resources together that might be in different environments and they have to copy the data from one place to the other. We can centralize that in one cloud. We can streamline that entire process and really make it much more efficient from an operational perspective and really streamline those calculations so that can be more accurate.

David: So for someone who wants to get in and play with these services to see if they're going to address the problems that they have, first of all, they can obviously get a free Azure membership for themselves that will give them credits to spool up some resources in Azure. But other than that, Nick, what are some steps that people can use to get started if they want to build their own solutions on top of Azure?

Nick: One of the key things, first steps to move forward is really how do we get certain data moved to the cloud because all your tools have to run against that data. So it's generally

much more efficient to pull the data from existing sources. For example, admin systems and extracts and move that data to the cloud. There are a lot of tools, that's simply automated and you pull that in the cloud and then you can run your calculations on top of that. Generally the simplest way is to almost do like a lift and shift of what you're currently doing in an IS solution so you spin up the VMs that are comparable to the VMs that you're currently using on premises and then connect to the data that you moved up. Either you have a live feed that gets updated and then the the software runs against that that you're currently using and then your BI tools can run against that afterwards.

So generally we want people to start small. Maybe look at certain components that you're doing now on prem and then move those up incrementally. Get comfortable with the way Azure works and how the data moves up and moves back and then use tools that we're using now. And then there were a lot of different services that are available that you can expand on that in the cloud. Then as you move forward, expand that out to be more comprehensive and look at other solutions as well.

David: One of the things I'm wanting to just mention is that you can download a supporting use case paper from the industry experiences team page in the show links. Well Nick, you made insurance really interesting for me today. Thank you for being on the show.

Nick: This is great. If anybody has any questions, please feel free to reach out to me. I have a long history of working in the insurance industry and there were a lot of solutions that are available now and on Azure and we keep adding more and from a technology stack it's constantly evolving. So I think especially as these new regulations come into effect, we're trying to make sure that those solutions are available to meet those needs. So if you have any questions or if you're a partner who was looking to move more of your workload to Azure as well, please let me know.

David: All right. Thank you very much for being on the show.

Nick: Well thank you.

David: Thank you for joining us for this episode of the Microsoft Industry Experiences Team Podcast. The show that explores how industry experts are transforming businesses with azure. Visit our team at [Aka.ms/indexp](https://aka.ms/indexp) and don't forget to join us for our next episode.