

DISPELLING COMMON CLOUD MYTHS WITH MICROSOFT'S STEVE BOHLEN



In this Run and Reinvent podcast, I chat with Steve Bohlen, Technical Evangelist from Microsoft, to dispel many of the common myths around cloud. This is a condensed transcript of the discussion.

Allison Cramer: What are some of the first mental shifts people need to make in order to be able to truly take advantage of the cloud and determine if they themselves are cloud ready?

Steve Bohlen: As organizations think about moving to the cloud, I think there are a number of things that they need to think about. Some of them are technology centric, and others of them are operations and process centric. The first thing I would suggest is that in moving to the cloud, one of the first things to pay attention to is that for a lot of organizations that already doing things in virtualized environments in their own data centers, fundamentally there are elements of moving to the cloud that are not really structurally that different from what they're doing in their own data center.

As you start thinking about your cloud strategy as an organization, I usually tend to suggest that people think about these things in phases. And the first phases is taking your own workloads and running them, to the degree it's feasible, as is in the cloud as kind of a first step of getting your feet wet and understanding how some of the technology works, and how things like your governance and security model need to extend to the cloud as an organization.

And then as you kind of move up the abstraction stack, you get to the point where you're starting to

use platform as a service offering, and so on, in the cloud. And those kinds of things really require thinking carefully about what your application architecture might need to be modified to in order to start taking advantage of those things.

Allison: What would you say are some basic rules you can look at for how you evaluate which workloads you want to test in the cloud?

Steve: If we take the situation where perception is reality, or where the perception is that my on-prem data center is in some meaningful way more secure than the cloud data center then one vector of thinking about what workloads are cloud-ready is understanding how your own governance and compliance models might apply to workloads that might leave your data center.

Lots of people have data sovereignty concerns, and there are lots of ways in the cloud to solve data sovereignty, but the easiest data sovereignty solution is simply delete the data in your own on-prem data center. And so, for organizations that have really strict regulatory compliance constraints and so on, thinking about which workloads do and don't fall within those kinds of constraints is one way to think about it.

If you've got workloads that are very burstable and very elastic therefore in the resources that they would consume, then that might be a good candidate to be thinking about moving to the cloud because there, even if you have regulatory concerns and even if you may have to think about re-architecting some of your solution for that, it might very well be that the biggest value add for moving to the cloud is for those kinds of workloads precisely because they're burstable and elastic, and the cloud might make a lot of sense in those conditions.

Another way to think about it is workloads that are greenfield, right, so that you're not looking at taking your existing technology solutions and tearing them apart to get them cloud ready. But instead you're saying new workloads will move to cloud because we're going to be able to architect and engineer those solutions for the cloud natively moving forward. But old workloads might remain in the data center because reworking all that code is either too much work or too much risk versus the benefit of moving to cloud.

Allison: Do you think most people have a great handle on their workloads and the inter-dependencies, or no?

Steve: Part of the process of evaluating a solution for its sort of applicability or appropriateness for cloud is getting to the point where you've got what you think to be, with a high level of confidence, a great 360-degree view of what that application's real requirements are, so that you're making the right choice when you say this is or isn't in the end a good candidate for a cloud workload.

As a simple for instance, it's possible to have an incredibly chatty application that talks to the database very chattily and produces lots of traffic between the application and the database. And so long as you're at a very low percentage of the overall throughput that your internal data center can handle, that chattiness is not really a negative, which is to say that you are paying to support that chattiness when it's on-prem. But, in a sense, you've bought the chattiness in bulk.

When you start moving to the cloud and start paying per chat, right, every time the application talks to the database there is a cost. Now it tends to be cents or fractions of cents, certainly, but you're paying as you go. And so, the penalty for being chatty is manifest very differently. Rather than paying up front for a data center with X capacity, you're paying every time you start doing a database transaction, for example, right. And so, the chattier application, which doesn't seem to be a cost problem in an internal data center scenario, starts to become very different when those things are

metered in the cloud and you're actually being charged per transaction in that way.

Allison: I think everybody who's gone to the cloud had gotten one of those bills where they're like, "Oh."

Steve: Yes, absolutely. In order to not be surprised by that bill, there is a very necessary upfront cost modeling process that needs to be done in order to understand that the act of simply moving your workload to the cloud is not going to magically produce savings for you, necessarily. It may. You may be lucky, right.

But the reality is that in many cases you're not going to realize the sort of true economies of scale that moving to cloud are going to give you until you're actually at the point where you can begin to be thinking about modeling and designing your application based on the way the cost model in the cloud works rather than the way the cost model in a data center works in an on-prem data center.