



The Deloitte On Cloud Podcast

David Linthicum, Managing Director, Chief Cloud Strategy Officer, Deloitte Consulting LLP

Title: Getting real about cloud spend issues and how to address them

Description: According to some sources, cloud expansion and spend might soon slow significantly. While opinions on the subject might differ, there have been some serious bumps in the road recently for cloud. In this episode, David Linthicum discusses cloud issues that are worrying the C-suite: cloud complexity and ballooning costs. He also gives advice on how to tackle them. First, figure out how cloud works best for your organization, then focus on planning and deployment, and, finally, adjust expectations to meet reality.

Duration: 00:25:11

David Linthicum:

Welcome to this Deloitte On Cloud podcast knowledge short exploring a specific topic related to cloud computing. This is a short tutorial talking about the real-world concepts in the emerging world of cloud computing. I'm your host, David Linthicum, cloud computing subject matter expert, author, speaker, and

managing director with Deloitte Consulting, and this is discussing a survey report that 81 percent of IT teams are directed to reduce or halt cloud spending by the C-suite.

So, this is a bit a departure from what we normally do here. We talk about a topic—Observability, multi-cloud, things like that—and then kind of expand on the topic as a point of education. Here we're doing things in a new way. We're going to look at an event or a survey which reported some interesting things and then also comment and provide our perspective on what this means, and if you're in the world of cloud computing, what you should anticipate this is happening in the market.

So, according to a new study from Wanclouds, 81 percent of IT leaders say their executives and boards of directors have directed them to reduce or take on no additional cloud spending. This is a bit of a surprise, ultimately, because this is around a reaction, the fact that we found that cloud costs recently have been skyrocketing while the economy remains in flux, a little bit of question mark there. And, so, after years of a lot of cloud spending and attendant growth, the finding suggests that enterprises have basically called a halt to unbridled cloud spending and looking to temper some of it down. So, need to look at a few things. Number one is this really what is happening, and if it is, what should we be doing, and just kind of taking it apart a bit.

So, first and foremost, always a bit skeptical around any study that declares something extreme will happen or is happening. So, aspects of these things are typically true, and there's some truth here, and we'll get into that. The technology shifts much slower than most people really kind of understand that it does, and enterprises will take their own time in slowing down or speeding up around any kind of technological change, cloud computing included. And, so, we had move into service-oriented architecture, move into the PC, distributed computing, all these kind of micro-revolutions that occurred over the years where cloud computing is kind of a macro-revolution in turning how we're going to move and consume technology differently. And, so, even after cloud computing being a relevant market change for, say, 15 years, we're probably at about 20 to 30 percent of workloads and enterprises that are migrated already to the cloud.

So, bear in mind that even if this has some truth into it, which I'm sure it does, that things don't typically move that fast. We don't decelerate and we don't accelerate as quickly as I think some of the analysis tell us that we do. Also, you need to look for the underlying motivations and conflicts of all of these studies. Not implying anything here, but typically they're going to reach a conclusion, other studies, which is going to be germane to putting forth their message. For example, security companies won't find that security is a problem that's been solved. Obviously, it's going to be something that needs to be determined and proven to be ongoing. That doesn't mean that the study's incorrect, but they're doing their studies for their own purposes, and that's why they exist, and there's nothing wrong with that. And I do a lot of analysis of those studies and surveys and spend a lot of time on this podcast and also on my InfoWorld blog trying to explain what they mean and picking them apart like we're doing with this one.

So, in the case of the Wanclouds studies, there is some truth in what they're asserting, in my opinion. So, that's because there are a couple of things that are occurring in the market right now. First, as we mentioned on this podcast a bunch of times, enterprises are hitting a complexity wall. So, they have so many services onboarded that they have no way of operationalizing them within existing budgets and resources. So—and that was also mentioned in the study, so what's happening is things have gotten complex.

We're moving to multi-clouds, we're moving to many thousands of services under management, versus many hundreds of services under management about ten years ago, and we're trying to move forward with the same sort of operational budget, and it's just not happening. In other words, we have so much complexity and so many things that we're managing, and so much heterogeneity that's going on within the enterprise, that trying to approach it and operate it at the same budget, the same number of resources, getting the skills you need, getting the tools you need, things like that is becoming a price tag that many in the C-suite are probably unwilling to pay, and so that's what's occurring right now.

So, multi-cloud usage is becoming increasingly unwieldy and costs are difficult to manage across hybrid environments, and that was reported in the study. And we talk about these problems here a ton, so I'm not going to belabor the point, but we all know that if we lead to more advanced systems, we're choosing different services, we're leveraging more services, we're dealing with more heterogeneity, we're dealing with more distribution, we have to deal with security and operations and FinOps and all these things that need to layer into these implementations, and sometimes they can be overly unwieldy and overly burdensome based on the resources that we have.

You've got to remember with enough money and time, we can operationalize anything, but enterprises have a finite amount of money that they can spend on this stuff, and so what they're saying is, "If you're asking for us to operationalize these systems, we need 20, 30, 50 percent more resources, money, skills to make this happen," and they're kind of balking at the price tag. So, that's what's going on here. So, didn't realize how much it was going to cost, didn't have a good understanding of where this was going to go, now they're seeing the price tag from some of this stuff, and I think the C-suites and the boards of directors are looking to slow things down in terms of figuring out why the costs have inflected to the point that they've inflected.

Second, as we've recently covered, the ongoing costs of cloud computing is shocking most of the enterprises. So, you've got to remember 15 years ago when we were talking about cloud computing—not me specifically—but cloud was being promoted as something that's going to reduce costs. In other words, it's going to provide more efficiency in the environment and through economies of scale and CapEx to OpEx shift, those sorts of things. We're going to get to a point where we're going to have cost efficiency. That's going to be built into this new architecture. And, so, right now, as I mentioned earlier, we have barely 20 to 30 percent of enterprise workloads on the public cloud providers, and the bills are much higher than expected, and this is for a few reasons.

Number one, there probably isn't a lot of discipline and accountability when it comes to cloud costs, so if you are allocating things they probably don't need, and they're probably not returning those resources back into the pool, so they're not billed from them. So, there needs to be some cost accountability discipline that needs to occur, and that's what FinOps is, the ability to put financial operations, programs, and tools so we're monitoring costs, figuring out what's being spent, and are optimizing our existing cloud infrastructure to optimize costs and so to return more value back to the business. And, also, the fact is that I think the cost of cloud computing in general, not just the service cost, the cost, the fees you're paying for the service usage, but the skills and

the training and the amount of people that you need, the ability to do integration with existing legacy systems, many of that costs were not necessarily factored into the overall cost of cloud computing, and many times enterprises didn't see that coming. Now they see the bill and they're balking at the bill.

So, most of the ROI problems that I see out there for many enterprises are self-inflicted. However, the industry, I think, bears some culpability for overselling and overstating cloud cost savings, as I mentioned earlier. So, while I would love to say that cloud computing is always going to save you money – by the way, that's something I never said or anybody who understood this environment ever said. In general, when I'm called to help fix a cloud implementation gone wrong, I find a couple of root causes. Number one, first, little thought went into planning that needed to occur before the first purchase decision was made around cloud computing. How are we going to leverage and implement those environments?

And second, trying to lift and shift your way to success has rarely worked out, as people are finding. And, so, the reason is that if you're lifting and shifting, you're taking something without modification, you're moving it from the existing as-is state, which is going to be on-premise, to the to-be state, which is going to be in cloud. And if you're not taking the time to redesign every factor of those systems to take advantage of cloud-native features, then it's not going to perform as well, it's not going to be as cost-efficient, and you're missing some reasons why you're moving to cloud computing.

But of course, we saw this during the pandemic; people made a mad dash from the on-premise systems into the cloud because they viewed it as low risk because of the quarantines and lockdowns, remote work that was going on, cloud computing, which was a better model and still is a better model from kind of the next generation of the enterprises where they're taking it and how they're going to make these things work together and the ability to kind of leverage cloud as a value proposition. So, didn't modify, didn't do a lot of planning as they move applications into production, and they did so for cost savings and because they wanted to speed along the migration process, and that didn't turn out as expected.

So, we have very inefficient applications that are not necessarily taking advantage of the cloud-native services that are running on the clouds, and that's where the cost is coming from. And, also, it's just kicking the can down the road. Many instances, the applications need to be redesigned, the databases need to be redesigned and normalized, can't find a single source of truth on all those sorts of things, and in many instances, those weren't fixed. Not all. There's companies out there that are doing the right things and made modifications of the systems, did the right amount of planning, and the applications are going to be better off running in the cloud because of the refactoring and investment that occurred as they moved the applications into the cloud.

So, you have a couple of choices now. Number one, you can fix them where they are, so in other words, they're migrated to the cloud and we can refactor them there, but that typically means we look like we're migrating the applications twice—once lift and shift to get into the cloud and then second migrating into a new target platform that's going to be more efficient leveraging cloud-native systems, things like that. And, so, that's a bit of redundant work, and that's a hard pill to swallow for people who just made millions of dollars of investment in moving into the cloud. But in many instances, there's really kind of no choice here. Ultimately, if those things are inefficient and they're costing you more money, the cloud's going to be your target environment, and that's where you're going to run your production environment for the foreseeable future, probably for 10, 15, 20 years before something better comes along, and it probably won't for a while. Then you're going to have to fix the applications so they're much more efficient and running on a cloud environment, so you have to figure out how to make that investment.

So, while it's tempting to wag fingers at everybody in enterprises that got into cloud trouble because doing lift and shift when it shouldn't have been done and not doing planning, it's not productive. So, it's better to determine where your enterprise is on its trip to cloud and then figure out how to incrementally improve in both the short and long term. So, what does that mean? Well, a couple of things. Number one, you may have to go slower to go faster. So, some missteps that need to be fixed such as massive lift and shift projects that involve moving poorly designed and built software to the cloud without modifications means that lifting and shifting the software, as we mentioned earlier, that doesn't work well in that the data center will not magically solve its design problems when things are moved to the cloud. So, we can't move things to the cloud and expect things to be fixed. They need to be fixed either before you move them, during the move, or once they're there as we just discussed.

So, there's no cost improvements that can be found at the end of the journey. As a result, in many instances, this is more complex deployments that cost 30 to 40 percent more to operate in the cloud at the end of the day. So, in other words, the difference between something that's done right, the ability to refactor an application so it takes advantage of the cloud native features, and just lifting and shifting an application, getting it working, compiled, linking it up to data, and then leaving it out there in production, you're going to get an underutilization hit, inefficiency hit of about 30 to 40 percent. So, that's 30 to 40 percent more money you're going to have to spend when you run those applications, and that's the price to be paid because you're avoiding the money and the investment in fixing the applications, refactoring the applications once they reach the cloud.

So, again, these problem programs need to be fixed or refactored before they migrate. You can fix them after migration or rewritten to take advantage of cloud native features, and that's typically where the savings is. Or research might reveal that the best operational and most cost-effective solutions for potential lift and shift applications is to keep them to remain in the data center. So, in other words, if you're not going to do the things you need to do to refactor them so they're able to leverage the cloud environment, and therefore we're going to have the applications underperforming and underutilized in terms of cost efficiency, then in many cases it's probably better to leave in the data center and then migrate them at a later time when you're ready to make the investment.

So, those decisions need to be made. Sometimes they're not popular because we want to get moving and start moving things into the cloud, but in many instances, if you can't make the investment because you may not have the budgets to move the applications to improve their utilization, to improve their efficiency in the cloud by doing the refactoring and the planning that needs to occur to make that happen, maybe you shouldn't do it. Leave them where they are.

Next would be focus on planning and deployment, as we mentioned earlier. Planning kind of was key to this. People have kind of gotten away from it. Used to do a lot of planning on moving systems to any new platform and haven't necessarily done the proper amount of planning and moving them into the cloud. That doesn't mean no planning's been done. Usually there's a base amount of planning that occurs, but as far as what we're going to do with the applications, what we're going to do with the data, how we're going to reutilize it and get it into a state where it's going to be more cost-efficient and more

productive for the business and also provides the core reason that we're moving to cloud computing in the first place, the ability to change things as we need them and really become a force multiplier for the business. So, use optimized architectures instead of what seems to work or what someone else hypes. I think that's kind of the core lessons learned over the last ten years.

Containers, for instance, are solid solutions for more existing and new applications, but they should be reconsidered with all the alternatives such as using traditional non-container development and understanding that they bring additional operational complexity, which should be factored into the solution. So, we have to understand moving to containers, we're doing so for a good reason. Containers in Kubernetes container orchestration, getting the portability, the simplicity, have a very nice ecosystem of tooling, good reasons you can use containers, but there's always downsides to it. It's going to cause more operational complexity and also it's, in many instances, going to take more resources. We call it the container tax. In other words, you're going to spend more money in moving an application to containers or building a net new application in containers than if you did not leverage containers. So, in many instances, there's a reason to do that, to leverage containers. Other instances, there's not. You have to consider the tradeoff.

The goal is to move to the most optimized solution for simplicity and cost. Sometimes it involves something different than what we envisioned. We have to make tough decisions about not using technology that seems to be in the zeitgeist as popular, containers being a good instance of that. Lots of people are using containers right now, and they're using them successfully and for good reason. But in many instances, just like we gave the example that some applications should have been left behind for now, that we may have to not necessarily leverage the technology we feel we should be leveraging because we're looking at a more cost-efficient way of doing something.

We're able to find a solution that's able to return more value to the business somehow. And, sometimes, that's not what we envision, and architects—I tell this all the time when I do architecture training—you have to keep an open mind in terms of what technology to leverage and how to bring them in and the reasons to bring them in. There should be a good business validation reason for leveraging any technology. And not to pick on containers, but other things like AI and serverless computing and lots of trends that are going on right now. There's a reason to leverage those technologies and there's a reason to avoid them.

Finally, adjust expectations. Cloud computing is not the savior of poorly run IT. It's just another enabling technology that works well, and if the right amount of planning occurs before resources are committed, then you're going to be successful with cloud computing. Ten years ago, or so, the reality was often shouted down in the early hype-filled days of cloud computing. So, I remember those days, and people were talking about cost efficiency and it did nothing else. There could be no chance that cloud wasn't good at saving you money, and architects kind of knew the reality of that, that you can certainly save money with cloud, you had the potential there, but you have to make the right moves in terms of architectural planning and leveraging the right technology, everything we just talked about so far.

So, this is not "I told you so" moment that vindicates those of us who were shouted down and didn't necessarily look at the downsides of cloud computing at the time. It's clear and a present wakeup call that we need to learn from poor assumptions made in the past that we must deal with today and make better plans for the future. And, so, we're learning here. In other words, we're iterating through this. In many instances, this is the first time that many enterprises have moved to cloud and the first time professionals have engaged with it. So, we're learning as we're going. We're making mistakes, and that's going to happen. In many instances, we're looking at the reason for the pushback from the C-levels in terms of the cloud spending, pushing back on some mistakes that were made, but we're learning from them, we have to overcome them. So, this keeps moving forward no matter what, so you have to figure out some way to find an efficient route to fix technological solutions that you have in place and fix some of the plans you have in place, and really try to take the enterprise to the next level using the right technologies and then solving the problems incrementally. That's the only way to get out of here.

So, ultimately there's going to be some uncomfortable conversations with the executive team and board of directors about some of the issues that are now arising, and I have those often, too, myself. The key to success moving forward is to admit that things need to change and have plans in hand that show you're willing to do the heavy lifting, ultimately, to make those changes. So, we're going to have to make some uncomfortable decisions if you're in this spot. In other words, you're not necessarily able to progress because you're not at a state of efficiency where you have the confidence of the C-suite and the board of directors that you should be spending more money to go down this path. So, things need to be fixed and you have to figure out how to fix them. That's going to be a tough thing to do, but it's doable.

We've been through these cycles before, however, this time it'll take more time and money than other past improvements and fixes that needed to be made. Cloud complexity is a problem that many didn't see coming, and heterogeneity and dealing with the cost overruns, things like that. It was here to stay. New tools and configurations come on the cloud scene almost weekly to help us deal with complexity, but underlying problems still need to be addressed. We hurried up to get to cloud, and the lack of planning is starting to show. That's really underlying of what this means. If you hurry up to apply the endless parade of fixes that just kicks the can down the road, that's not necessarily going to solve your problem. So, this becomes a more fundamental look at how we do architecture planning and, ultimately, set a path to fix IT if these issues are currently in your enterprise.

So, it's time to take a step back, identify the problems, and do your research, and then you iterate through them. Lots of battles win the war, and so it's getting complexity a bit better, getting application utilization, getting cost accountability better, getting a FinOps program in place to at least know where things are being spent, and you're able to have more control over cost monitoring and cost utilization. Those key things are fundamental to moving down the road and then getting to the tougher things. In other words, we're going to have to go in and reengineer some of these systems and do some refactoring—not all of them, but some of them—and we're going to have to make some tough decisions around which enabling technology that we want to leverage.

And, again, looking at the cost utilization and the ability to bring value back to the business, that should be the key metric here. It's not the ability to containerize everything. I always get concerned when people have a technological goal, but it is the ability to make the business more innovative to a certain degree, the ability to make the business more agile to a certain degree, the ability to define all this value of cloud computing which we can leverage as a force multiplier, and then return that value back to the business. And that's really kind of the goal of IT. It's the goal of technology in general. People who are—we're leveraging within IT the weapons of cloud computing and other technologies and all the underlying stuff that's out there—AI, serverless,

containers, edge computing, all these sorts of things that are hyped. We still have to figure out how to measure and how to optimize the value that kind of comes back to the business, and I think ultimately that's our objective here, and we need to kind of keep that in front of us.

So, anyway, to summarize this, I think that not unexpected. We always have these kind of hangover periods after adoption of technology. Cloud computing's about 15 years old in terms of its hype in the market. It's been around for a long time, by the way, but let's say it has a market life that really kind of started 15 years ago. We're running into some problems, and they're not unexpected, costing us more than we thought it was. It's creating unforeseen problems such as complexity and talked a lot about here and how to deal with that. It's creating overruns in terms of how much it's costing.

Some things we saw coming in terms of not necessarily having accountability around cloud costs and utilization but also the cost of skills have gone up tremendously. Lots of problems and also lots of solutions out there we're able to leverage. And, so, keep in mind that the way you solve this thing is to line them up, figure out how to iterate through each of the problems and solve them, make sure they're in a priority order that you can work on, make sure you're aligning the resources correctly to make that happen, and execute. That's all we can do.

And if you're looking for assistance and help, reach out to someone at Deloitte, send me an e-mail, let me know what's going on, and we're happy to give you our advice on how to do that. And, so, just kind of keep that in mind. So, we're all in this together and these are all solvable problems. Let's get going on them.

So, if you enjoyed this podcast, make sure to like us, rate us, and subscribe. You can also check out our past episodes, including those hosted by my good friend, Mike Kavis. Find out more at deloittecloudpodcast.com. If you'd like to contact me directly, you can e-mail me at dlinthicum@deloitte.com. So, until next time, best of luck with your cloud journey. Everybody stay safe. Cheers.

Operator:

This podcast is produced by Deloitte. The views and opinions expressed by podcast speakers and guests are solely their own and do not reflect the opinions of Deloitte. This podcast provides general information only and is not intended to constitute advice or services of any kind. For additional information about Deloitte, go to Deloitte.com/about.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor.

Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

Visit the On Cloud library
www.deloitte.com/us/cloud-podcast

About Deloitte

As used in this podcast, "Deloitte" means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms. Copyright © 2022 Deloitte Development LLC. All rights reserved.