



The Deloitte On Cloud Podcast

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

Title: Looking back on 2022—what we saw, and what it means for cloud

Description: The past year has been a mixed bag for cloud. Certainly, there have been positives like the rise of supercloud and metacloud to help reduce complexity, the growth of observability to provide deeper operational insights, and an increase in cloud native development. However, cloud has also been beset by rising costs, complexity, and other challenges. David Linthicum provides his insight and analysis on all these topics—and more—in his review of 2022.

Duration: 00:23:19

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 looking back at 2022 around the world of cloud computing.

So, I always do a recap on what has occurred in the previous year. As we get toward the end of the year, and as I'm recording this, we're certainly heading quickly into 2023, and, in fact, looking forward to that year in terms of growth of cloud computing, technology in general, and really kind of an exciting time to be in this technology space as we're kind of in for lots of changes in terms of how cloud computing's going to be leveraged and what technologies are important, things like that. And we will have a forward-looking recap, in other words, a forward-looking idea of what we're going to see in 2023, and we'll bring that to you in another podcast, but for this one, we're going to look back at 2022 and look at the major items and topics and concepts that occurred, what was really memorable about this year. And, so, let's get going.

First and pretty obvious, the continued rise of multi-cloud. So, we've seen multicloud kind of rising pretty rapidly and aggressively over the last five years, and 2022 was no different. Probably it was a little bit more of an aggressive rise as people were more and more leveraging multiple public clouds to build a multicloud, and either they were coming at that from two different ways. Either it was happening by accident, which is what most enterprises are experiencing. In other words, they just started onboarding different public cloud providers, and suddenly they're in the multi-cloud game just because we needed a best of breed service that we couldn't find running on one particular provider.

And, so, we went from one brand of cloud computing to another brand of cloud computing. Suddenly we have two public clouds, sometimes three, and therefore we have a multi-cloud. Fewer, but they do exist, enterprises are planning on how they're getting to multicloud. They're putting together a strategy in place, figuring out security and governance and operations, all the things we talk about here and doing on the podcast that really should be done, and they're getting to multi-cloud on purpose. That's obviously the better way to do it, but it's going to cost money and time and go a little slower than the way that the enterprises are doing it now, just having it occur organically.

And you're going to come at it from one way or the other. Chances are, if you're in a cloud computing, 95 percent of the folks out there who leverage cloud computing aren't doing multi-cloud right now. Typically, they have a major cloud provider that could be 80 to 90 percent of the cloud workload's around that particular provider, and then they have a secondary provider, which gets them into being a multi-cloud, sometimes three, sometimes four. And as time goes on, we're going to need to leverage these best-of-breed services. We're going to need the best-of-breed in terms of AI, in terms of serverless computing, in terms of data analytics, things like that, and we don't want to be limited by a single public cloud provider. Therefore, we're going to extend our reach out to other public cloud providers, and that's going to be the driver that's going to move us into multi-cloud.

So, we also saw some unique things that occurred in the multi-cloud world this year, but first and foremost when we talked about this on the podcast before the rise was supercloud or metacloud, and what that is; it's not a product, it's not a technology stack. It is kind of a technology stack, but not in a way I think we traditionally think about it. But it's a loosely coupled set of technologies that sit above the cloud providers, not within the providers, that provides cross-cloud services. So, we can have common FinOps, financial operations, common CloudOps or cloud operations, common security systems, common data control planes, all these sorts of things.

And instead of doing it within each specific cloud provider using the ways that that cloud provider wants you to use in their walled garden, you build these systems that are loosely coupled to the public cloud providers above the cloud providers that work across clouds. They physically may be running on one particular cloud or another, but the idea is that we have this conceptual framework of this technology stack that sits above the public cloud providers that provides these common services so we don't have to implement them redundantly in the particular silos of the public cloud provider.

So, we don't have to have—if we have three public cloud providers, three different security systems all native to those particular providers, three different CloudOps systems all native to those particular providers, three different data planes all native to those particular providers, governance systems, all these things that we need. Instead, we solve the problem once, put it in a cross-cloud system, and they're calling that supercloud or metacloud because what we're doing is we're trying to remove as much redundancy as we can from our multi-cloud deployments and then push them up into this larger layer which is loosely coupled to the public cloud providers but provides a common layer across them.

So, we build security system only once, we build the CloudOps system only once, we build the FinOps systems only once, the data plane only once, and we use the same things from cloud to cloud, and so that seems to be a better idea and certainly is a better idea because it makes things simple. And you've got to remember, the big thing that we're battling with multi-cloud deployments is going to be the complexity of it, and the complexity is caused by the fact that we're answering every question with a different cloud service from a different cloud provider, and therefore we're going to too many cloud services under management that we have to operate, and that won't scale. And, so, if we go from 1,000 cloud services under management to 3,000 or 5,000 cloud services under management, we try to do so with the same operational budget, you know, something's got to give.

And metacloud really is a partial answer to that because what we're doing is we're not implementing a new cloud service each and every time we need to solve the problem, but we're finding services that are able to run across clouds. Obviously, that's going to be a better solution. So, continued rise of multi-cloud and the concurrent rise of supercloud and metacloud, and we'll see more of that in 2023 as we try to figure out what that is, look at the different technologies that fit in that stack, get some architectural best practices in how we build these systems.

And next and also another knowledge short that I did just a while back, the rise of operational observability has certainly occurred in 2022. Observability has been around for a long period of time. As I mentioned in the knowledge short, it's an engineering term, and the idea is that instead of looking at raw monitoring data that just tells us current states of things, and CPU saturation point, and temperature of a server, and all these sorts of things that are meaningful in terms of how we operate the system, but they provide us with the ability to find insights in that data.

So, it's almost like instead of looking at the raw information and monitoring a system as to what its current state is and even looking at the past history of the states and looking at the log files, things like that, but the ability to find out what all this information means. In other words, what does the trend in some sort of a performance indicator on a particular database server mean in terms of our need to reboot the server, update the server, fix the server, or do something that's going to be more proactive and us getting ahead of problems versus reactive where we're waiting for problems to occur and then fixing the problems as they occur.

And kind of the way operations is done today, on-premises as well as in the cloud, is that everything's pretty much reactive. They wait for something to happen, they call it firefighting, and they go off and fix it. Observability allows us to look and find different insights into massive amounts of information, so instead of just looking at the raw data and trying to figure out what the trend is, what could be occurring, we have these systems, typically AIOps, artificially

intelligent operations systems, that are able to take these massive amounts of data, this monitoring data, and apply different analytical systems to it so we can figure out the trends that are going on within the data, we can gather insights within the data that we normally couldn't see, and that allows us to find problems before they become problems.

And, also, it allows us to automate the fixes of the problem so we can, instead of manual intervention each and every time, notifying a human and they have to go fix the issue, we set off a set of self-healing processes that fix the issue automatically. And that's obviously something that's going to scale better. We're going to be able to provide more uptime because we're able to operate on near-perfect information if we get the observability systems up and running and working.

The other thing is observability, not just operations, sits in other areas as well. We have security observability, the ability to use observability systems to find the security trends that are leading us to understand what could be occurring in terms of a potential breach. For example, we know that saturation of the CPU could be an indication of an application that's running that may be overworking the CPU, or it could be an indication of a breach attempt occurring. And we can link all these things together so again we're being proactive, the ability to kind of find breach attempts and hack attempts before they occur, versus reactive, and that's how most people deal with security issues now. They wait for the hack or the breach to occur and then they go off and fix it, putting together as much prevention as they can to stop that from occurring, but you're not looking at the core trends and insights as much information as you can look at so you can figure out how these things are going to happen before they happen. Again, that's the power of observability.

And then finally FinOps observability, cost observability. We'll talk about that later in the podcast, but the rise of FinOps is kind of a core headline in terms of what occurred in 2022 as well because as we started building these systems, we needed some way to monitor the cost, and just like operational monitoring and security monitoring, all this cost data comes in which is huge amounts of information. It's very difficult for us to sift through and understand. Well, FinOps and cost observability, cloud cost observability allows us to figure out trends and insights into that data—cost data that's coming in from the cloud billing systems, things like that to figure out things we can do to potentially optimize those systems in terms of cost optimization.

We get more for less, and ways we can save money ongoing and moving things into the cloud. Right now, it's a bit of Wild West. They're doing cost monitoring using spreadsheets and whatever native tools that are coming with the particular cloud providers. That's something that won't scale, certainly if you get into a multi-cloud or complex cloud deployment. You're going to need these cost observability systems to get the insights you need to do to monitor costs, govern costs, and optimize the costs.

Another trend we saw in 2022, skill shortages continued from 2021, and certainly 2020. There was a mad dash into cloud computing, and suddenly people with any sort of cloud skills were in huge demand. And, so, we have about five to ten cloud jobs that are chasing one qualified candidate. And, so, not enough qualified candidates to go around based on the demand that the enterprises are seeing in the marketplace, and, so, a couple things are happening. They may be pushing off projects and putting projects on hold until they can get the talent in place. So, the skill shortage is limiting progress and success of cloud computing, ultimately, and we're never going to be a place where we have demand for skills that's aligned directly with supply.

There's always going to be some mismatch there. Usually, it's going to be more demand than skills that are available, so enterprises are figuring out how to staff up their teams, how to do training of existing personnel and moving them into cloud, and kind of creative ways to getting the skills in place that they need to take their business into cloud computing. But it continued to be an issue, and it's a core limiting factor in terms of people's progress and success of leveraging cloud for the enterprise.

Next in 2022, we saw the continued rise of cloud native application development and architecture. Huge space and huge amount of growth occurred in this area in the last two to three years, and for good reason because we're trying to move to modern applications that are very aware and able to leverage the particular features on the cloud platforms that they're on. And, so, we're moving to technology that provides us with more portability and more scalability and the ability to link into the features and capabilities of the particular cloud providers that they're hosted on. And cloud native can mean a couple of things. It can mean things that operate only on a particular cloud platform, on a particular cloud brand, or it's mostly to find something that is turn of an architecture that provides a scalability and portability that people are looking for in modern applications.

So, it's a bit of semantic infusion there. I blogged about this on InfoWorld about six months ago, and the fact of the matter is it seems like we have two different meanings there, but for all practical purposes, the core meaning is as it's defined by the Cloud Native Foundation where we're looking at particular technologies to leverage that technology to take our application game to the next level. So, that's going to be using containers—that's certainly on the rise, hugely popular right now—the use of container orchestration such as Kubernetes, and use of microservices in these architectures. And as we move into 2023, we're going to see federation of these systems, and we'll talk more about that during our podcast that looks forward in 2023, but kind of keep that in mind that cloud native application development and architecture is well supported. There's an ecosystem there. There's a number of people who get how to build these systems.

There's a consortium of vendors who focus on this area, the Cloud Native Foundation, and it's a bit confusing in terms of the semantics, but we have the capabilities and the patterns there and the architectural foresight to figure out how to modernize either existing applications, in other words, take them to a cloud native deployment, a cloud native architecture, or build new applications, get them to a cloud native deployment, cloud native architecture. But it's on the rise just because of the fact we need to have these modern applications that are better able to take advantage of cloud scale systems, be able to provide better services, able to provide better portability, and able to provide more sustainable operational value as the business leverages them. So, this occurred in 2022. We're going to see it moving forward in 2023 and likely beyond.

The final trend we're going to cover for 2022 is kind of a downside. In other words, we did some surveys and understood the core benefits of cloud computing and who was gaining with cloud and who wasn't with cloud, and many enterprises are finding that cloud costs more than they thought. And by the way, at Deloitte we did a survey, you find that there's certain organizations that are able to take the same amount of money and do a lot of creative and innovative things with cloud computing and add value to the business where lots of organizations using the same amount of money typically couldn't figure that out. And we do have a podcast on that survey, and I urge you to go back and kind of get some data points from that because it's interesting in terms of where the market's evolving and how people are leveraging this technology either as a true force multiplier or something that's not as valuable to the business.

So, we saw lots of cost data and found out that many businesses are claiming in surveys and data points and things like that that they're not getting the value from cloud computing that they thought they would get. And then we have to look at the reasons for this, and many instances these are self-inflicted, and we found this in the survey as well. So, lack of planning, they're not putting the planning in place to figure out how to leverage the cloud computing systems well and optimize the use of those systems, so it's lots of pushing applications.

We certainly did this through the pandemic. Pushing applications from existing legacy systems into the cloud and not necessarily localizing them for the—to take advantage of the cloud native capabilities that are on the particular cloud platform, so security, memory management, things like that. We're just finding a platform analog that exists on the particular cloud provider, same platform, same operating system, same development stack, recompiling the applications, attaching it to a database. If it tests, it runs, everything's good to go. The problem with that, those things are going to be way more expensive to run in the cloud. They're not optimized, they're leveraging too much CPU, they're leveraging too much storage, too much network connectivity. There may be a need to pull data into those systems, into the cloud, which is ingress, which is expensive, or egress, moving them outside of that cloud provider to some other system to support the application.

So, if you don't do a lot of planning and you don't expend some resources in doing things like building cloud native application development architectures for these systems so we can modernize them better, typically they're going to be much more inefficient in running on the target cloud platform you're moving to, and the enterprises figure this out. So, they're getting large cloud bills, and that was a bit of a concern based on the fact that they're trying to figure out how to make cloud computing migration, use of cloud resources come back to the business in terms of value, and that obviously was a deficit to value, and that continues to be a problem.

So, you have to consider the tradeoffs with moving fast and therefore building inefficient and under-optimized systems that exist on the target platforms that they get to are moving in a more purposeful way. We're getting to these systems with the right architectures, optimizing these systems, these applications for the cloud platforms that they're on, and therefore optimizing the amount of money that we're spending on these systems because they're going to operate at less cost. Keep that in mind.

Second would be too much complexity for multi-cloud was something that many people didn't anticipate. So, we already talked about why and how people are moving to multi-cloud today. The complexity, or basically the use of services. You may go from 10 different databases to 30 different databases, 5 different application development stacks to 40 application development stacks. All these things kind of come along with you leveraging additional cloud brands to solve particular types of problems, basically leveraging best of breed technology to solve those issues.

So, this leads to complexity, and complexity costs money. We have to operate complexity, we have to secure these complex systems, these complex deployments. We have to provide FinOps around these various systems and make sure we're accounting for the different cost monitoring systems. So, the more services we have out there that are running, the more it's going to cost us. Many enterprises didn't anticipate these costs, or they see the cost overruns coming in. And, certainly, from the surveys that we've done and for the surveys that the other analysts have done that we see this kind of phenomenon in the marketplace right now. So, we've got a cost issue. We've got to figure out how to solve the cost issue.

The other couple of issues would be cloud cost discipline. In other words, accountability, making sure everybody's participating in the optimization of the costs, we're not just allocating resources willy-nilly – there has to be a purpose to that. We have to put governance around who's able to allocate resources and for what purpose because too many times resources are allocated in a public cloud provider, storage system, database, things like that. They're left up and running too long, they're bigger than they need to be, they're not managed as well as they should be, and therefore the cost overruns are there.

So, a couple of things here. We have discipline issues that we have to take care of, we have accountability for who's producing and allocating those resources and managing those resources effectively, and then putting FinOps programs around them, so we have visibility into those resources, what they're doing, what they're there for, why they exist, who owns them, all these sorts of things where we have this overarching system – typically spans multi-cloud, this would be a component of the supercloud metacloud, by the way – and is able to have a centralized way to manage, monitor, and govern cloud costs. And that doesn't seem to be in wide deployment today. And, so, we're seeing the cost overruns come back, but there really is no visibility into why those are occurring. They know we're getting a large bill, very much like you getting a large electricity bill or water bill at your home and having to look at what the root cause of those bill overruns are. Same thing with cloud computing.

And finally, lack of cloud cost observability. We already talked about the notion of observability but taking this to FinOps is going to take things to the next level. So, again, it's not just looking at monitoring the cost data. We're getting billed for particular things that are coming off, resources that we're leveraging, storage, databases, things like that, but what it actually means and what insights are in there to understand. So, getting to a point where we're not only understanding the cloud costs that are coming in but the ability to look at insights in terms of how to improve and optimize these cloud costs moving forward.

And doing so by taking it out of the hands of the humans and not having them look at it and then try to figure out what cloud configuration is going to be a better optimized configuration, whether we should use reserved instances by resources ahead of usage, all these sorts of options that are available to us as people who are managing cloud costs or just the ability to kind of do optimization of these systems. So, we have an automated system that's able to do this for us, and so we have a FinOps system in place that leverages cost observability, it's able to find out why we're spending what we're spending and ways to reduce spending and still provide the cloud resources that we need and kind of take our business to the next level and being better optimized.

So, keep in mind that we have 2023 cloud outlook coming. We're going to sit down and pontificate about what's going on next year after we talk about what went on this year. And lots of great content coming up in 2023 in terms of the On Cloud podcast, and guests and topics and lots of things you've been asking me for, so make sure you keep that subscribe button hit and let us know how things are going.

Anyway, 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. 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](https://www.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.