



## FTC Cloud Computing – Taking Stock and Looking Ahead Meeting | May 11, 2023

Chair Lina Khan:

Good afternoon, everybody, welcome. Thanks so much for joining today's discussion that the FTC is hosting on cloud computing. I'm thrilled to welcome our panelists, whose breadth of expertise will help make sure that the FTC is fully grasping the intricacies of cloud computing, and is best positioned to approach its policy and enforcement work accordingly. I'd like to give a big thanks to our Office of Technology and our terrific CTO, Stephanie Nguyen, for spearheading the FTC's inquiry into cloud computing business practices. And I'm also very grateful to Hillary Greene, who is our Special Counsel for Competition Policy, who helped put together today's panel. So, in collaboration with the FTC's Bureau of Consumer Protection and Bureau of Competition, our Office of Technology staff have done a tremendous job in developing key questions that prompt a wide array of stakeholders and get at our root concerns relating to competition and security in cloud computing.

Staying at the cutting edge of key technologies and new business developments has long been a core part of the FTC's mandate. Congress expressly designed the FTC with tools that would enable us to fully understand how new technologies, how emerging technologies were working, to make sure that our traditional authorities could keep pace. Today's inquiry is squarely part of that tradition. Cloud computing, as we all know, is now a central part of our economy. It's the engine that is driving the proliferation of data-driven services from video conferencing to streaming to artificial intelligence, and it now often serves as the scaffolding on which companies are storing data, rapidly processing information, and scaling their businesses. Disruptive technological innovations, too, are often built upon cloud computing infrastructure, and as businesses race to innovate and commercialize new AI tools, they, in particular, will continue to rely on cloud providers to develop and deploy their products on a massive scale.

We've also seen, however, that as large swaths of the economy become more reliant on the cloud, we've seen only a small handful of providers are currently equipped to provide this key infrastructure. And we want to make sure at the FTC, that we're fully understanding, what are the factors that are leading the market to be so concentrated in the hands of a few companies? And also, what are some of the downstream risks that might stem from that, that need to be on our radar, both on the competition side, as well as on the consumer protection and data security side? So some of the things that we're looking at, for example, are the high barriers to entry. We know that it can be extremely costly for companies to build the underlying infrastructure on which cloud operates, and that maintaining a network of physical data centers can be a huge investment. Among other things, new entrants would need the real estate to physically stand up the data center, as well as the equipment to effectively safeguard and operate the technology, and so, these high barriers to entry mean that it's much more difficult for new firms to really enter this business. We've also seen how the biggest cloud computing

providers are all large tech companies that have been able to substantially invest in their products over the last few decades, and these companies, in particular, have benefited from network effects that a large existing customer base provides them. We've seen one of the downstream risks of consolidation in this market has been a resiliency issue where, if you have a single provider that is hosting business for a large sector of the economy, if that one provider goes down, a big chunk of the businesses that depend on it will also be impaired.

And over the last decade, we've seen time and time again, how cloud outages can have a cascading effect across markets and across our economy, and so, understanding, in particular, this resiliency and fragility implication as something that is top of mind. We're also keen to understand the degree to which data security issues might arise, again, given that we have a handful of players that are charged with operating and overseeing this infrastructure, and how this concentration may compound some of the security risks. Already, the FTC has brought several enforcement actions related to security issues in the cloud, including numerous breaches where stored consumer data ended up being exposed. And as we learn more about the associated risks, I'm confident that we're going to have the expertise to continue to be nimble in our approach.

One way in which competition and consumer protection intersect in this way is the question of whether additional competition in this area would incentivize firms to better compete on providing greater security and providing greater privacy, and so, that's an intersection that we're looking at, as well. Just to step back, really excited, again, about this event. I think we've all seen how cloud computing is providing such a foundation for a growing part of our economy, and we at the FTC just want to make sure that we are fully abreast of some of the challenges that might arise from the current market structure and some of the practices that we see. So thanks so much, again, everybody, for joining us, and thanks so much, again, to our brilliant panelists, whose insights I'm sure will make for a fantastic discussion. And with that, I will turn it back over to Hillary.

Hillary Greene:

Great, thank you so much Chair Khan. I begin with the disclaimer that all my statements are my own opinions, they do not represent the position of the commission or any commissioner. As Chair Khan indicated, better understanding cloud computing practices, particularly as it relates to competition and security, is of enormous importance. And it is those very issues that lie at the heart of the FTC staff's recent solicitation for public comments on the business practices of cloud computing providers. That document, called a request for information, or RFI, went live March 22, and the agency will be accepting comments through June 21.

This panel complements the RFI by bringing together five exceptional individuals with diverse expertise and positions for a discussion that touches on a number of key themes associated with the RFI questions. Now, it goes without saying, these individuals do not represent the agency views, nor are they a representative sample of possible views. What they are, however, are very experienced and very thoughtful experts who will offer valuable points of reference as part of much of a larger and ongoing discussion. But hold tight, we're going to be introducing our panelists in just a second, but first, I'm very excited to introduce Alex Gaynor. He is the Deputy Chief Technologist of the FTC, and by background, a software resilience engineer, and we're going to turn to Alex now to provide a brief tech primer associated with cloud computing. Alex, take it away.

Alex Gaynor:

Thank you for the introduction. I'll start by also noting that these thoughts are my own, and don't necessarily represent the views of the commission or any commissioner. As I was putting together these

remarks, I realized I was coming up on 15 years of being a software engineer, which means my career has roughly covered a time span that started with a default being for a company to purchase its own servers, place them in its own data centers, and use them to run its own software. Now many companies no longer own or operate any of their own physical servers, and instead they've migrated to cloud services, where they rely on a service provider to offer them virtual servers, and sometimes infrastructure that's even further abstracted from hardware.

A feature of my career, having covered this transition, is that I clearly remember the point in which I and many of my peers and colleagues hated the term, "cloud." We thought it was vague. We thought it didn't have a precise definition and that meant it was destined to be a marketing buzzword, not something practitioners ever talked about. Like many changes in language, the word stuck, whether we liked it or not, but it still leaves us with the question of, what exactly is a cloud? It's common to organize cloud services into three high level buckets: infrastructure as a service, platform as a service, and software as a service. Infrastructure as a service products offer customers raw building blocks, such as servers, networking, and storage. Platform as a service are more integrated offerings which trade off less flexibility with doing more out of the box. And finally, software as a service offerings are complete tools, rather than simply being building blocks.

It's important to recognize that while these categories are useful, they're more of a spectrum than hard and fast distinctions. And to further complicate matters, many cloud service providers have offerings from all of these categories at once. These categories help us organize within cloud services, but they don't help us understand what a cloud is in the first place. For example, is a private cloud where developers create servers on demand, but the company still owns and runs physical infrastructure, really a cloud? It's a tough question, and the answer really depends on what's important to you about a cloud. I'm a software engineer, so the best I can do is tell you what's been important to me about them. Rather than offering a formal market definition, I want to instead talk about three elements that help me distinguish cloud from non-cloud services.

First, in the cloud, resources are elastic. It's possible to provision new servers, and then spin them down with the snap of a finger. No longer are users reliant on placing orders for hardware and waiting for them to be delivered. For example, if you have a product that's more popular on the weekend than on weekdays, you can spin servers up on Friday afternoon, and spin them down Monday morning. Second, you're billed by how much you use, not based on what you physically purchase. You can see how this is intimately linked to resource elasticity. If you use a server for three days or three minutes, you want to pay for that much usage only. The utility of being able to spin up and down would be dramatically decreased without billing by usage. And finally, cloud services have APIs, which means you can write software to manage your infrastructure. To keep our example going, you might write a program that handles the scaling up on Fridays and down on Mondays, rather than doing it by hand each week.

Focusing on these three elements combined with the growth of platform as a service and software as a service, are useful in understanding how, to an engineer, the cloud is different from what came before it. There's no question that use of cloud infrastructure is becoming ubiquitous and deeply intertwined with many of the other trends we see in technology. It's quite likely that any AI chatbots you've played with recently were running on a cloud. Along with the cloud's ubiquity, has come many security breaches, which have cloud-specific features. For example, a fact pattern we've alleged in several cases is that one of the factors contributing to a company's poor security and data breaches is that they mismanaged and misconfigured their cloud systems. And while it's beyond question that companies have a responsibility to secure consumer's data effectively, whether in the cloud or not, it can also be instructive to ask, are there factors that make it more or less likely for cloud users to misconfigure their environments in the first place?

At the Commission's December open meeting, I spoke about how the field of safety engineering teaches us that effective security programs don't accept human error as a final answer to why a vulnerability occurred, they used it as a jumping off point to ask deeper questions about how systems are designed. Because of their widespread use, when a cloud provider designs part of their system to be secure by default, all their customers benefit. But when a cloud provider designs its systems in a way that make it easier to misconfigure something than they do to configure it securely, incidents can follow.

Further, the set of offerings from some cloud providers are so expansive that knowing how to use them and secure them effectively is a challenge requiring dedicated expertise. I'd be remiss if I didn't also take the time to share that I worked for the FTC's Office of Technology, created earlier this year. The role of our office is to build on the FTC's distinguished history of leveraging hands-on technical experts in its work. Events like these, with the goal of connecting the trends in technology with both our competition and consumer protection missions, are central to what our office is here to do. And with that, I'll return the virtual microphone to Hillary.

Hillary Greene:

Thank you, Alex, for starting us off in such fine form. I appreciate it. Now let's turn to our discussion—a brief word about the format before I introduce our panelists. The discussion will be divided into five topics, industry landscape, competitive landscape, open source software, security, and artificial intelligence. Each panelist has been asked to anchor one of the discussions by providing a short statement on their topic, then the other panelists will join in with their own comments.

Turning to our esteemed panelists, namely, the folks who are going to do all this work, I will provide just brief introductions. In alphabetical order, our panelists are Salil Deshpande. He is the founder of the Bain-backed Uncorrelated Ventures, which focuses on open source and infrastructure software. With over 18 years as a venture investor, he has invested 500 million plus into more than 50 companies. Salil was on the Forbes Midas list for 100 best performing venture investors worldwide, in nearly every year for the past 10 years.

Frederic Jenny is Professor Emeritus of Economics at ESSEC Business School in Paris. He has written extensively about trade competition and economic development. Prior positions include non-executive director of the Office of Fair Trading in the UK, judge on the French Supreme Court, and vice chair of the French Competition Authority.

Abby Kearns is a veteran tech executive, as well as board director and angel investor. Her career has spanned Fortune 500 companies and startups alike, including Puppet, Cloud Foundry Foundation, Pivotal Software, Verizon, and Totality. Abby currently serves as a board director for StackPath and Lightbend.

Frank Nagle is an assistant professor in the strategy unit at Harvard Business School, where he is also co-director of the HBS Linux Foundation Core Infrastructure Initiative. He serves on the advisory boards of multiple big data analytics startups. Prior to his academic career, he worked at startups in large companies in the information security and tech-consulting industries. And fun fact, he developed and taught a two-week course that all FBI cyber agents had to pass before entering the field.

And last, but certainly not least, we have Steven Weber. He is a Professor of the Graduate School at UC Berkeley School of Information, and a partner at Breakwater Strategy in Washington DC. Steve founded the Berkeley Center for Long-Term Cybersecurity and served as its director for seven years. His research teaching, public interest, and advisory work focuses on the political economy of knowledge-intensive industries. In such brief bios, it's impossible to fully capture the relationship of the panelists and their research to the cloud industry participants broadly defined. There is more information online at the

events page, and the panelists have also kindly agreed to just briefly elaborate here for the benefit of our audience. So what I'm going to do is to just go in alphabetical order, and so, we'll start with Salil.

Salil Deshpande:

Hey, everyone. So Hillary, you want us to elaborate on our bios?

Hillary Greene:

Yeah, this was the disclosure information that we talked about.

Salil Deshpande:

Ah, yes, yes, yes. So Hillary covered the bio sufficiently, and as disclosures, I am a venture investor, and I'm invested in companies, several that benefit from the cloud's behavior and practices, and others that are harmed by their behavior and practices. And aside from the venture funds, I'm also invested in US large cap growth funds that have Amazon, Google, and Microsoft as core positions.

Hillary Greene:

Great. And we next have ... I'll do my alphabet quickly here ... it's Frederic.

Frederic Jenny:

Good afternoon, everybody. I have written a report which was commissioned by CISPE, which is a trade organization of European cloud infrastructure providers, on the state of competition on the market for cloud infrastructure. And a lot of what I'm going to say today is comes out of this research, which had two dimension research on the one side, and a survey of customers on the other.

Hillary Greene:

Great, thank you. Abby?

Abby Kearns:

Similar to Salil, I'm an active angel investor and advisor, so obviously, I have a strong bias to companies that are building in this space. And I'm also an investor in the public caps market, which is obviously very tech heavy, given that's my experience.

Hillary Greene:

Frank.

Frank Nagle:

Thanks so much for having me. Just in terms of disclosures, some of my research is sponsored by the Linux Foundation, GitHub, and also, the Omidyar Network, where I'm a senior fellow, although none of those organizations have any sort of final say over the research itself.

Hillary Greene:

Thank you, and Steve.

Steven Weber:

Also, thanks. On disclosures, in my capacity at UC Berkeley, I've received research funding in the form of unrestricted gifts from a number of technology companies that might be relevant in today's discussion. All the sources of those research gifts that have funded my work over the last eight years are actually listed on the website of the Center for Long-Term Cybersecurity at Berkeley. In my capacity as an advisor and consultant, which I've done, really since 1990, through both personal capacity and a bunch of firms. And currently, at the advisory firm, Breakwater Strategy, I should say, we do work in the technology sector writ large, we do not currently work for companies whose interests are directly impacted by any of the things I'll talk about today.

Hillary Greene:

Okay, great. Well, thank you, guys, and now, let's get started. Okay, Abby, I'm going to turn to you to start us off and help situate us, perhaps by providing an overview of the industry landscape for cloud computing, perhaps flagging main players in the market, and describing a little bit about what dynamics have contributed to the current marketplace structure. Abby, take it away.

Abby Kearns:

Well, thank you for having me, and I'm going to build on what Alex started, which was really a generalized description of cloud, and how cloud was used today. But cloud, the way we think about it, has been around since AWS launched in 2006, if you can believe it. So it's been some time that we've been talking about cloud, but to really simplify cloud, in my mind, I think of it, it's just somebody else's data center. So as Lina pointed out, there's a tremendous amount of capital expenditure require to build out these footprints, with servers, infrastructure, data centers.

And we've been building on that for years and years and years, but it wasn't until about 2016 that we really started to see significant traction in workloads moving to the cloud. And it was in 2016 that we had a confluence of a couple of different things. One, cloud native architecture, so architecting and application and your infrastructure, to actually take advantage of those things that Alex described, in terms of elasticity, resilience. The technology really started to evolve more quickly, starting in 2016, and we started to see, particularly in open source, we saw projects like Kubernetes start to get traction. We started to see a lot of different part of the cloud native stack really start to evolve in the open source, that really, fundamentally changed the way we thought, designed, and built around the cloud.

And that also happened at the same time we started to see meaningful traction, and what we've termed in the enterprise as digital transformation. Now, that's another overloaded term that we've been talking about for years, but it was really where companies realized that software was becoming a stronger core part of their business, as well as the way they competed in the market. And so, through all of that, we find ourselves today, all of us are using something that runs on the cloud every day—be it our iPhone, be it our photos, be it services that we use, how we order food, to how we check in for our flights and buy tickets. We're constantly leveraging technology that runs in the cloud. But if we take a step back, we're seeing more and more services run there because it's easier, it's more efficient, and oftentimes, it can be more secure than running in your own data centers, which is where a lot of companies really started, running their own data centers, and having to run and manage all of this infrastructure, and bear that cost.

And so, cloud, for a lot of reasons, has made it easier to build, run, and manage applications that are really targeting specific services that we're building for customers. To put some facts around this, AWS was started in 2006. Today AWS pulls in about \$80 billion in revenue a year. Gartner says that the worldwide public spending on cloud, on public cloud, is going to hit about \$600 billion this year. So this

is a very big market that's getting bigger every year, it's not getting smaller. So there's no fixed pie that companies are competing for, it's actually just getting bigger and bigger as more and more companies are leveraging the cloud. And when we talk about leveraging that cloud infrastructure, we're not talking simply about taking workloads that we're running in a data center and moving them to the cloud. We're actually talking about new software that's being developed in the cloud to run these unique services and these unique capabilities.

So it's actually, that pie is getting bigger every year. Obviously, we're seeing a ton of that momentum with AI driving new applications, new usage patterns, new capabilities. But, at the end of the day, we're starting to take more and more advantage of that. And then, one thing I actually wanted to mention is, in 2016, we started talking about the concept of multi-cloud. So we've talked about single points of failure here, we've talked about resiliency, but multi-cloud really became a strategy that more and more organizations started taking advantage of to really either take advantage of unique, native services in the hyperscale cloud providers, or to actually have ultimate optionality. And I should probably end, since I've kind of rambled here for a minute, with who those providers are. And today, the top three are Amazon's AWS, Microsoft's Azure, Google's Google Cloud platform, and with a fourth provider being Oracle's Oracle Cloud infrastructure. I know I had five minutes, I was trying to keep it to five minutes.

Hillary Greene:

I sound like a tyrant, but thank you. And now we're going to turn to Frederic. Please, what I'd love for you to do, sir, is with that foundation, what types of anti-competitive or at least competitively suspect practices might one expect to observe? To what extent do you actually see them? What justifications have been offered? Can you start us off in that inquiry? Thank you.

Frederic Jenny:

Well, thank you very much for having me. I'm going to talk about one particular aspect of this sector in the sense that business firms are certainly users of cloud services as well as suppliers of those, and I'm going to talk about the competition issues with respect to professional services that are offered by the cloud. And from this point of view, I would like to start with the idea that it's important to keep in mind that there are complementarities between three elements, the information system that the firm generates because it provides data, the infrastructure service for storing the data of the firm and getting the computational support, and the softwares that it uses to analyze the data to support the management task of the firm. What this means is that there is a certain technical complementarity between software and infrastructure cloud services. Now on the supply side of the infrastructure service, you have firms, as we have said, which provide both software and infrastructure facilities. Microsoft would be one with Office 365 on one side and Microsoft Azure, and those are integrated providers. And you also have firms which are only providing infrastructure services, which I would call the naked providers. Now, when an integrated provider supplies complementary products or services, let's say A, B, and C, and has a strong market position on one of those complementary services, let's say A, but faces more intense competition on the other services, B and C, it may be tempted to leverage its market power on A in order to acquire market power on service B and C. And one of the possible reasons for doing that is to protect itself against the risk that another firm would acquire market power on the services B or C and then challenge its own market power on service A.

The cloud infrastructure service market is very concentrated with an increasing level of concentration certainly between 2015 and 2021, and the increase in concentration has favored the integrated providers, the providers of both software and infrastructure services. And the three leaders, as has been said, are Microsoft, Google, and Amazon, but roughly speaking, they have 70% market share of the

cloud infrastructure service. Now, it is worth asking whether this highly concentrated situation and increasing concentration is an evolution that result from anti-competitive or unfair practices by the integrated providers or whether the result from the superior efficiency of the integrated providers. Now, from an analytical point of view, what are the anti-competitive practices which a dominant firm on a software service, for example, may have an incentive to implement?

I'll give you three examples. The first one is tying cloud infrastructure with software. I tell my clients, "If you want my software, you have to store your data on my cloud." This would be an example of tying. The second is bundling. That's creating a technical, or legal, or economic link between the two services offered. For example, if you use my cloud infrastructure, I will charge you less for the software service on which I am dominant or my software service will work better with your data. The third possibility, besides tying and bundling, is predation. And there the dominant firm on a software service can subsidize the sale of infrastructure service through the profits which are made on the sale of software where the firm has a dominant position or a market power so that the price for the infrastructure service it offers is below its variable cost, which makes it impossible for naked providers to match the offer of the integrated provider.

Now, we were asked by CISPE to talk to customers and ask them, "What are the practices that you see as buyers of those services?" What did they tell us? Well, first they felt that some of the integrated providers were resorting to exclusionary strategies. One of the best example of this is that in 2019, Microsoft imposed on the users of its productivity software, Office 365, that they would have to repurchase their licenses if they store their data on a third party cloud ... Sorry, somehow I've been muted, so let me go back for a second on this. In 2019, Microsoft imposed on the users of its productivity software, Office 365, that they would have to repurchase their licenses if they store their data on the third party cloud infrastructure. There was no technical reason which was given for this practice, but it certainly had the effect of deterring the users of Office 365, and there are numerous, we know that.

Hillary Greene:

Can you also make it more broader in terms of explaining how these points might apply more broadly than into just that setting? I don't know if that's possible.

Frederic Jenny:

What do you mean by more broader? I mean I think whenever we in the clouds are talking about complementary services, some of which may lead to a dominant position, those can be leveraged by the holders of the dominant position to other services, and that creates a possibility of a competition problem. And I was talking about Microsoft, but I could equally well talk about some of the other examples which were given to us. Oracle, for example, as a strong position on a software for database management and, from what we've been told, Oracle imposed restrictions on its software licensing which can lead to a tenfold increase in the licensing cost of its software if this software is used on a third party infrastructure provider.

There are other example with SAP, I'm not going to go into this arm. But second, there is the possibility of strategy that try to tie or mix bundling, and those seems to be not infrequent. A good example of this is Microsoft's decision to tie Microsoft Team with the other offering in Office 365, and we know that there are competition issues. Other cases are cases where Microsoft, when it loses a tender for infrastructure, will come back to the potential customers offering to include Azure for free in a bundle with the upcoming renewal of Microsoft licenses. Finally, there's the possibility to reduce the



compatibility of the software so that it will increase the cost of rival providers of infrastructure. And there are also example of this.

Hillary Greene:

I'm going to jump in here and thank you so much for that because you've now set the table with laying out issues about tying, bundling, predation. You've also flagged questions about when conduct is strategic in terms of it might be technical versus is it a business decision. And with a amply-laid table with so much information from both Abby and Frederic, and I thank you both for that, what I want to do is I know that our fellow panelists are chomping at the bit to get in because I've chatted with them separately. And so let me open this up to more broad. Who would like to go first? Okay, Steve, I will turn to you, if that's okay.

Steven Weber:

Okay, I'll start. I'll just tie a few things together that I think Abby and Frederic said, and double down on them. Look, I mean we know there are real economies of scale in the native provider landscape that are probably just almost facts of nature at this moment, that are technological realities. Server farms, data centers, compute, et cetera, it's just a fact of life that it's expensive, and it becomes cheaper as you have more and more of these things in one place. I think, Hillary, you've asked us to talk as well about what kinds of data and information should the FTC be looking at and collecting.

That's relevant to the question in this case of whether or not there are such economies of scale for integrated providers, and I would just double down on what Professor Jenny said with regard to a very close examination of licensing terms. I think licensing terms tell us a lot about the ways in which people are trying to parse the difference between technical constraints and business model constraints. And so I think a very close comparison and examination of licensing terms in different services across different cloud providers almost a necessity of this investigation.

Hillary Greene:

Okay. Salil, and then I'll turn to Frank.

Salil Deshpande:

Well, the bundling and tying that Frederic mentioned is widespread. If you make a shopping mall analogy, these clouds are like shopping malls, but they're not just the shopping malls, they're also the stores in the malls and the products in the stores. There are other stores in the malls, but for the most popular products, the stores and the products belong to the malls or the clouds. And that's not a coincidence because the clouds monitor and observe which products are popular and copy them, often by using some of the techniques I'm going to talk about in the open source topic that I'm anchoring.

And there are a number of detailed instances of bundling that we can get into that are surprising. For example, it's usually Amazon that trailblazers these practices and the other clouds often follow suit, sometimes not, but often they do. For example, Amazon charges for cross-zone transfers. Amazon has 16 availability zones in the US and when you're using their Amazon services, they don't charge for cross-zone transfers, but if you're using any other service, they charge for cross-zone transfers. In the mall analogy, imagine if the mall had 16 different parts and you got charged to go from one area of the mall to another area of the mall. Unless you were buying Amazon products, then you didn't get charged to go from one area of the mall to another area of the mall. And there are lots of little examples like this.

Hillary Greene:

Right. And so I'm going to turn to Frank, and before I do that I also want to put something else on the table. I think that part of the framework that's been set up is an important one, which seems to be the software as service having implications for infrastructure as service. And what I'd also like and invite people to consider is the extent to which infrastructure is service, what are ways that it can be entrenched through certain activities? How should we be thinking about it? I know that one of the issues that's come up with folks has been egress fees generally, and so I invite people to also opine on that. But we're going to turn to Frank and then if other people want to jump back in, please let me know. Frank.

Frank Nagle:

Sure, yeah. Before we move to that, I was going to pick up on a thread that I think Abby started but others have talked about, and that's the concept of multi-homing. Adding to Steve's list of important things, it's not just about the contracts, but it's about the tech itself too, and understanding why would somebody use one versus another or why might people be using both? And there was a Statista survey a couple thousand companies from 2022 that showed 80% of companies use AWS, but 80% of companies also use Azure, and then something like 40, 50% also use Google Cloud. And so nowadays, I think this was five years ago, this wasn't as much the case, but now everybody's using everything.

And so when we think about that from a competition standpoint, we have to at least ask ourselves the question: has this become a commodity? I can go to any gas station and put gas in my car because it's been commodified. I don't think we're at the point where cloud has been commodified, but certainly one could start to think about how commodified is this? And back to Hillary, the direction you were going, is that different in different layers? Certainly the layer that it'd be most likely to be commoditized would be infrastructure, less so when we're talking about platform or software as a service.

Hillary Greene:

Steve, I think you were going to jump in and head ... Nope, you're good? Okay.

Frederic Jenny:

Just-

Hillary Greene:

Yes.

Frederic Jenny:

... if I may, just a very short comment. A, I believe that those practices should be investigated. B, I think that we've got to be also sophisticated. There are quite a number of possibilities because of the fact that those services may be or may have to be used interoperable among each other that there may be technical reasons for which there are limitations, for example, on the license, on the use. But it's also possible that on top of these elements, which in itself makes the migration possibly difficult, there are also anti-competitive practices to create moats around there. And I think that the role of competition authority is really to try to distinguish what is technically justified from what is an additional obstacle which has no reason except to try to develop market power.

Hillary Greene:

Great, thank you. Salil, I know that a concept that you've brought up is gravity of data, and so I'm wondering if you could explain a little bit about how that comes into play here.

Salil Deshpande:

Yeah. Data has a lot of gravity, and through a combination of all of these techniques and practices, some of which obviously should be examined and others that might be okay, we end up having all our data in these clouds. And that incentivizes you naturally to have more applications in that same cloud because it is expensive to egress the data away from that cloud to somewhere else. So the cloud's charge for the egress of data. Now, you could argue that is fair because it does cost them something to serve up the data, but I think it just exacerbates the problem caused by the combination of all the other practices that if those practices succeed, your apps and your data end up in these clouds, and then you're stuck there.

Hillary Greene:

There you go. Abby, do you have anything you want to add here or not.

Abby Kearns:

Always. Right, data gravity is a problem, and I think if you get too heavily weighted in with your data and it gets too big, it's real expensive to move out. That's a known issue, but it's also the same, the technology exists for you to write your applications to have portability across any cloud. Platform as a service, which was mentioned here, strong reliance on APIs, that gives you that total optionality in terms of where you move your workloads, but obviously that data becomes a persistent problem. However, you can also natively write your applications in a way that also locks you in, but I'd like to think that people are better educated now than they were maybe four or five years ago on the risk they're taking by those architectural choices.

And we're also seeing a rise of a lot of innovation happening on the data side to give people a lot more flexibility and fluidity with their data choices. Now we're not there yet, but I do think that we're starting to see that transition where people are like, "Okay, I'd like to take full advantage of cloud, that full elasticity, but I'm going to have to rethink the way that I think about my data." And not to add further murkiness to the conversation, but I do think edge compute that we're starting to see that pick up speed is also going to help really aggressively make the shift as people start to think about moving data in a much more distributed fashion, not just in one data center, but how do they make it much more ubiquitous so they can have access for their applications that are running all over the place.

Hillary Greene:

Great. And so with that backdrop regarding the industry and competitive landscapes, I want to now turn to our three additional topics. Each of them is important in their own right, but we're going to explore them in connection to cloud computing. And with that, I want to turn to Salil. I know you've given a lot of thought to open source software as it relates to the cloud—and so why is open source software important within this context? What competitive challenges have you observed? Can you help us think about it?

Salil Deshpande:

Yeah. Open source has been an important go-to-market path for software companies. And in short, the path consists of making great software available for free along with the source code with a permissive license and allowing consumers to download it, try it, integrate it into their systems, all without having

to talk to salespeople. And if it meets their needs, then to purchase the commercial version, which might have more features or better support. And this is what makes open source software sustainable because that monetization adds money back into the system and it's used for further R&D to make better software. But with the advent of clouds, again, Amazon trailblazing has exhibited three particularly aggressive or offensive behaviors towards open source.

First one being that it takes open source code produced by others, runs it as a commercial service, and gives nothing back to the commercial entity that produces and maintains that open source, thereby intercepting this monetization of the open source. And it's done that with an innumerable number of projects. Second, it sometimes forks projects, and forking means making a copy of the main repository and then maintaining the copy as the main repository, forcibly wresting control away from the commercial entity that produced and maintained the open source projects, and it did that in the case of a company called Elastic. And third, it sometimes hijacks open source APIs and places them on top of its own proprietary solutions, siphoning off customers from the open source project to its own proprietary solution. And it did that to a company called MongoDB.

And what this has led to, and it'll only be more so in the future, is less open source, more restrictive licenses, or software that is nominally open source but all the interesting features are in closed source, not open source. And that slows down innovation, it hurts innovation. I certainly advise all my portfolio companies now to not have permissive open source software licenses just to prevent the clouds from doing what they do, whereas 10 years ago I advocated the opposite. I'm certainly in the eye of the storm where I can see this hurting innovation. Ironically, although consumers initially benefit because Amazon and the other clouds are able to take something for free and offer it for cheap, in the long term it may be harming consumers because it's slowing down the cycle.

Hillary Greene:

Great, thank you. And so I know that this is something that lots of other folks have given some thought to. Frank, can I turn to you for additional comments?

Frank Nagle:

Sure, yeah. No, I agree with pretty much everything Salil said and certainly I think, back to this multi-homing nature, open source could be a great equalizer in some ways or a great reducer of some of the competitive barriers that we were talking about. Because if you're using an open source package on Google, you can easily use that, in some ways, same open source package on Azure or on AWS. But I do think the way that the big players, in particular, but even smaller cloud players, interact with open source can be a bit predatory sometimes. And certainly some are bigger offenders than others. And I think the examples that Salil gave of AWS in particular are very salient. Look, all of them have done something like that, but certainly AWS is bigger, and so if all the cloud infrastructure is built on open source, and so if we are encouraging or allowing behaviors that are diminishing the open source ecosystem, then the whole thing's going to crumble under its own weight, if we allow it.

Hillary Greene:

Okay. Steve.

Steven Weber:

Yeah, I want to just inject a slight note of skepticism here, maybe just a counterpoint. I think we all know the open source license regimes have evolved quite a lot over the years, people have gotten very sophisticated, and lots of different licenses are out there. And there are open source licenses like the

BSD license that do permit commercialization, so just observing that's happening in and of itself, I think, doesn't constitute evidence that there's malfeasance. I think to Salil's point, it's really important that innovators think really hard about the nature of the license that they slap onto their work. It's a critical, critical piece of their business model, but it's a critical piece of how the innovation ecosystem evolves as a whole. And forking, as well, I would say the same thing. Forking can be illegitimate. There are also circumstances under which forking is a legitimate means of improving the basic elements of a source code. So the devil's here in the details, I do take very seriously the warnings that have been put forward, but I don't want to us to come away from this thinking that simply commercialization of appropriately licensed open source code is somehow in the face of it an illegitimate thing to do.

Hillary Greene:

Other thoughts? Yes, Salil.

Salil Deshpande:

Yeah. Addressing the counterpoint, I agree. I agree with the counterpoint. And at no point did I say these were illegitimate. I just said they were aggressive. So the license actually permits Amazon to do all of these things. So they're not doing anything illegal. They're following the letter of the law. Taken in combination, these are aggressive actions and may make the cycle unsustainable, but they may be perfectly legitimate and legal.

Hillary Greene:

Other thoughts. Abby? Yes.

Abby Kearns:

Yeah, you knew I'd have thoughts on open source. I agree with everything everyone has said and I think what's interesting and why does open source matter. I did want to take a step back on behalf of the audience to say that open source is in 96% of the world's software. And I wanted to point that stat out. Credit to Sarah Novotny at Microsoft's open source for that stat. But it's an important thing to recognize that it's in everything that we use. And open source has fundamentally changed in the last 10 years, to Salil's point. The way we think about licensing and it's the fundamental foundation of everything that we think about in terms of cloud, all the software that's running it. That infrastructure layer is all largely open source. And so it's become something that is critical to every business.

And so it's really changed the dynamic of the way we think about open source structures, foundations, what do you open source, licensing. That is all really dramatically changed in the last 10 years because it's now a fundamental building block for so many things, that I think we're at an inflection point where it stands that we should say, "Okay, what are we open sourcing and why? And what is the value to the larger market around that?" Because right now, a lot of software is open source, but how we manage that, how we think about that, and I'm sure this is going to be addressed also in the security section, what does security look like for our software supply chains and things like that. I just think we're at an inflection point where the way we think about open source is also changing at the same time.

Salil Deshpande:

One last point on that actually. I've written a couple of op-eds in TechCrunch, where I talked about some of what I just said, but the call to action was not to Amazon. The op-ed said that Amazon's not doing anything technically wrong or illegal. The onus is on us as the open source community to come up with licenses that are right, that accomplish what we need.

Hillary Greene:

Great. And with that teed up nicely, Frank, it seems like the panel has spoken. Security will come next. And so what does security mean in cloud computing? How might the market structure affects security outcomes? The issue was raised early on in terms of these big cascading problems. So what is single point of failure? What does it mean here?

Frank Nagle:

Sure. Yeah. And I'll give my own bias that open source and security and all these questions, we could have spent this whole hour and a half on open source and could have spent a whole hour and a half on open source security, but I'll try to stay a little broader than just open source and think about security more broadly in the context of the cloud. And usually, when we talk about these things, it's useful to have some basic definitions.

So one that we use in cybersecurity all the time is the acronym CIA, which of course is an entertaining acronym in itself, but in this context, it means confidentiality, integrity, and availability. So confidentiality, we're talking about only the correct people should have access to whatever it is, usually data. Of course, this overlaps with privacy, which we may talk about a little bit down the road. Integrity is thinking about that the data should be accurate and what it's supposed to be and not changing by itself or changing without permission. And lastly, availability. So this is thinking about the data or whatever it is, is up and running when we need it, with limited delays.

And so all three of these are super important in any context, including the cloud. But when we think about cloud, it's probably this last one, availability, that ends up being the most relevant to the competitive environment and to the most salient that hits the headlines when something goes wrong. So just as a few examples, last year, Google had a few big outages and disrupted their own applications like Gmail. Gmail was offline for a little bit, but also applications running on top of it like Spotify or Discord.

Similarly, AWS, shortly before that, had two big outages, two weeks apart. The first impacted the east coast of the US and shut down services like Netflix and Disney+ and Robinhood. And then a few weeks later, the West Coast was hit with something similar. Again, Netflix, but also Slack and the Ring Doorbell cameras all of a sudden didn't work and DoorDash didn't work as well for a few hours. So more recently, and we don't want to leave Microsoft out, earlier this year, they had a global outage that prevented access to tools like Teams and Outlook.

So to Hillary's point, and I think Chair Khan kind of alluded to this earlier, but we have these potentials for single points of failure, where if one system goes offline, then all these knock on effects and we have all these other services and parts of the economy that just don't work. Now, one defense against that is what we were talking about before in terms of multi-homing and certainly, this availability and single point of failure is part of what's led many companies to multi-home on the cloud. But it's not always, as discussed, it's not always simple to do that.

And so when we think about this in the context of cybersecurity, we can think about not only all of these things I just mentioned were mistakes or configuration errors that the cloud providers made themselves, but we can also put on our bad guy hat and think about who would we go attack if we wanted to cause the most damage or steal the most information? It's going to be one of these single point of failures. And so we can think about attackers being very focused on the cloud providers. And if they got into a cloud provider and got into the backend, they might be able to get all of that cloud provider's customers data.

Now, the way that the cloud is architected, and we heard this a little bit earlier, using containers and separate access or separate customers, it would be quite hard for an attacker to actually pull this off unless there was some serious bug in the backend. But when we think about cybersecurity, we often talk about risk as the intersection of likelihood and impact. And so although the likelihood of this maybe fairly small, the impact would be huge. When we think about competition and market concentration, these single points of failure represent a potential weakness in the overall system.

Now, on the flip side, we talked a little bit about economies of scale and how to provide these cheap services, these companies need to be at this scale, but we can also think about some other benefits like something, for lack of a better phrase, a single point of patching. So if we think about email for example, if a security vulnerability is found in Gmail, Google just develops a patch and loads it to all their servers, and then instantly all Gmail users are protected. But if you use Outlook on your desktop, Microsoft has to develop a patch and then everybody has to install it individually. And so there's trade-offs with this centralization and with this single point of failure. We also have this single point of patching that can be potentially a good thing.

And so when we think about trade-offs in this space, often the most salient are thinking about who the customers of the cloud providers are. So big companies have large IT budgets and can spend a lot on security on their own, but smaller companies, they don't. And so when we think about what benefits whom, often when we think about this investment of cybersecurity by the cloud providers, the biggest beneficiaries are the small companies. And so when we move into thinking about who's responsible for cybersecurity, we see the different cloud providers taking slightly different tack. So AWS says, "We're responsible for the cloud itself. You are responsible for everything that's in the cloud, you the customer." Meanwhile, Google and Microsoft have a slightly more tiered approach depending on if they're talking about SaaS or PaaS or IaaS.

So again, to bring everything together here in the context of the RFI and the open questions, I think that cybersecurity is going to be super important to the cloud providers, but it's unclear how much customers are actually willing to pay for better security. This has been a problem with cybersecurity free cloud, that we don't know how much people actually value security. Everybody says they value it, but are they willing to spend more money to actually get more security? And so even beyond that, thinking about how deep an insight a given customer has into the security practices of a cloud provider is still a little bit unclear as well. And so thinking about competition or the lack of competition and how that may influence whether a cloud provider spends a dollar on cybersecurity or more server capacity, I think is going to be a big place for more information to be gathered.

And then finally, thinking about just how different industries may think about security differently, finance or healthcare may put a super high value on cybersecurity, whereas just for sake of example, retailer manufacturing may have a lower value of on cyber. We have to think about how do these different industries on the customer side influence the importance of cybersecurity and investment as well. So that's it, and I look forward to the discussion.

Hillary Greene:

Right. And so who will chime in first? Steve, go for it.

Steven Weber:

Yeah, I'm happy to start. So I think Frank made a number of really, really important points. I'll just double down on the one that actually Abby had made early as well. Look, one of the most compelling reasons for companies to move to the cloud in the beginning was actually to achieve better security. Rather than have my own little CISO here, I might actually like the idea of hiring Amazon CISO to manage

my security. The single point of patch is extremely important. But I think we've said and we've noticed that it does create this very large and potentially uniform attack surface.

Now, the last thing you said, Frank, I want to double down on that and add to it—the point that we don't know how much customers are actually willing to pay for that security in that environment. That I think is where the main point of leverage needs to be. When I think about the history of banking, and you walk into a bank branch and you see the marble pillars, and remember that people used to leave the vault door opens, you could see how thick the metal was in the vault door. During a period of insecurity that was a little bit like what we experience in the digital environment today, banks competed over security. Cloud providers today do not really compete over security. I don't know that the federal government, through a policy initiative, can force them to do so, but I know that customers can do so if they ask for it.

And so I think that this, again, the responsibility for enhancing that competitive environment around what cloud providers can do in security, whether it be through configurations, whether it be through investments, I think Professor Jenny brought up configurations, whether it be through investments, whether it be through other kinds of collaborative practices, that is where we should look for greater leverage. There's really nothing more important because we're all going to be moving to the cloud in that regard.

Hillary Greene:

Is there leverage, Steve, to be brought to bear as the US as a procurer?

Steven Weber:

Well, yeah. I mean, look, big customers have more leverage in the market to use their power as procurers to demand different practices and different shared responsibility with regards to security. And if I were pulling levers, that is probably the lever I would want to pull first and foremost.

Hillary Greene:

Others? Abby, Salil?

Abby Kearns:

I have some thoughts. I think it's a great point. I do think right now, one, I will say anecdotally, people will only spend the bare minimum amount for security that they have to. Very few companies are willing to just overspend. And I want to also add to that, that I do think evolution in compliance, be it PCI, HIPPA, all of those are also driving core at least foundational good practices for security. So they're not great, it's not a hundred percent, but it at least forces companies to have some thought to how they're managing and thinking about data, privacy, access, et cetera.

Now, I do want to add a counterpoint to the idea that cloud providers should have greater security. Do I think they should take security seriously? Absolutely. But what I don't think it does is let users off the hook. Personally, if I expose my passcode in a public GitHub repo, that's on me. That's not Amazon's fault. And this happens probably more than y'all probably realize that this happens, and this happens quite frequently. And so you really have to think about, both from a compliance and security standpoint, what are the fundamentals and the framework and the guardrails, but also how do we do a better job also with educating users and educating companies on how to think about security?



I think CISA's done a great job over the last year of really starting to build that drumbeat of what we need to think about and the importance of security. We're starting to talk a lot more about the software supply chain, which Frank, when we think about open source, also software supply chain and security in that is super important with open source because it's in everything and there's really no strong focus on security from that angle right now. And so I think there's a lot. As Frank said, we could probably talk about this for hours and not even really address all the points, but I think those are also interesting tidbits that we should take into consideration because it's not just one simple fix.

Hillary Greene:

Others? Nope, okay. And so now we are on time and we are turning to Steve who is going to discuss a topic that's on everyone's mind right now, which is artificial intelligence and its importance within the cloud context. Steve, what is AI? Perhaps Steve, what are large language models? How do we need to understand it in the cloud context, and what are the roles of competition and innovation within this context as well? Take it away.

Steven Weber:

Yeah, thanks Hillary. Well, I'm going to start by saying, I don't love the term artificial intelligence. I actually prefer the term machine learning, or in this case, talking about specifically large language models, which are predictive models that, as everyone has experienced now, are able to mimic human language. But I'm going to use the term artificial intelligence because everybody else uses, and I've lost the battle to wipe it from the face of the earth.

AI is now the leading application shaping competition and increasingly security in the cloud. I think it is the application or class of applications that is going to shape the conversation that we have of this sort next year and for the next couple of years. And I think we've all seen that the applications can do things right now that seem quite incredible, but they are still just applications that run in the cloud. And so from that perspective, they're just like cloud gaming or software as a service offerings or back office productivity suites that run on cloud.

The economics of large language models and the economics of the cloud are intertwined, and they're going to stay that way for a while. So this new vision of everything as a service in which cloud ecosystems on platforms are the platforms in which we do our work and in which innovation occurs, that is all right now about LLMs, large language models. And we've seen this burst of LLM driven innovation and products over the last six months. I mean, it's nothing short of spectacular.

Reasonable people can disagree on whether this is really fundamental evolution or revolutionary stuff, but we're seeing lots of new ideas and lots of new companies doing really interesting and really value creating things. In that context, I would point people to pay a little less attention maybe to the ChatGPTs of the world and a little more attention to the domain specific large language models which don't get the same kinds of headlines but are doing really interesting things like a triage machine for emergency department admissions and things of that sort.

I think there's a lot of upside and some downside risk as well. But my number one objective for thinking about the cloud ecosystem here, when it comes to machine learning and AI, is that its support and facilitate this kind of really rapid and widespread innovation, particularly among new and small companies. In the same way, really, earlier manifestations of the cloud as elastic compute did for earlier phases. We used to say at the Berkeley Information School that with last generation cloud services, you could have a master's student project. With a couple of months of work and a credit card, you could start a company because basically you could buy all your infrastructure from AWS.

I remember back in the 1990s, we called this kind of innovation ecosystem a Cambrian explosion, lots and lots of new companies. And many of them go extinct and that's actually just fine from an innovation perspective. It's the explosion of innovation that we want. Cloud as a platform for LLMs should support that kind of thing. And I think actually just parenthetically, maybe we'll talk about it, I actually think that kind of multifaceted innovation is really good for managing some of our concerns around the downsides of algorithmic bias, fairness, deep fakes, and so on and so forth. But let me put that in a box for now.

So where are the risks to this vision? That's the question I ask myself about the cloud. And one way I like to think about this is to, maybe this is a very 19th century metaphor, to envision the cloud as an AI factory, like a factory where AI gets built. So what are the inputs that the factory needs? Needs massive compute power, needs lots of data and data engineering capability, reinforcement learning from users and so on and so forth.

But the reason I mentioned those three things is because they are all subject to significant economies of scale and they reinforce some of the natural monopoly tendencies in cloud that have come up a number of different times today. In that situation, nobody wants to compete over pure elastic factory capacity. And I think it's fair to say, and it's come up in some of the conversation, that the instinctual business practices of at least some of the hyperscale cloud providers reinforce that tendency as well as vertical integration.

And this is the point that I really want to end with because it's the most important. Vertical integration that segments the ecosystem can start to look like a pretty attractive alternative, and I think we see signs of that in today's AI environment. The number one issue for me here is the nature of some partnerships between small AI model innovators and hyperscale cloud providers. And so to be clear about that, look, I think it's absolutely prescient and smart business for big tech firms with amazing cloud businesses to invest in AI startups. On the other hand, it concerns me when those investments come in the form of cloud credits that are tied to a particular cloud provider and its specific architecture rather than the more conventional investments we might see from a venture investor. That kind of subsidy has the effect of tying a particular company's set of AI models to a particular cloud provider, which means that those two technologies overtime co-evolved together, in a way that I think Professor Jenny mentioned. And the further we go down that road, the harder it is to unbundle, the harder it is for the AI startup to run its models effectively and efficiently on a competing cloud platform. And I see a lot of downside associated with that trajectory.

Again, going back to the 1990s, we called it "walled gardens." It predisposes towards business strategies of silos and bundling, and it reminds me very much of something that we used to talk about, again, in the nineties. People may have forgotten this, but there was a concern, a fear that we were heading toward a fragmented world where there would be a Netscape web and an Explorer web and maybe an Apple web, and actually they wouldn't be able to interconnect. That didn't happen, thankfully, because of a mix of government policy, progress in open source, and the underlying open protocols that drove the web.

But those things aren't really in place right now for the AI machine learning world. So I think it's a valid concern and I think it's not too early to be talking about it. It would not be a good outcome. An AI ecosystems of walled gardens, I'm going to say would be a terrible outcome for society, for economy. It would give a slower, more fragmented innovation, higher prices for end users, and I think create a real risk that we get stuck in a local maximum or suboptimal tech trajectory with gen AI. It's very, very

Very early days in this world. There are other forms of vertical integration that we're starting to see signals of as well. Some of the hyperscale cloud providers building their own silicon may make a lot of sense for efficiency, but again, it tends to fragment the system into walled gardens where a new innovative small AI application company might have to choose, am I going to throw my lot in with this

cloud provider, that cloud provider or third cloud provider and not actually be able to move across them or essentially design for competition among them.

Let me stop. There are a number of different objections, a number different developments which might undermine this trajectory. Some of them are taking place in the open source LLM world, which is moving forward quite quickly and we may want to talk about that. But I just want to emphasize it's very early days, there is a potential Cambrian explosion of innovation. Cloud architecture, cloud business models should support that, not get in the way of that.

Hillary Greene:

Right. And with that, I know we have lots of people that want to get in. Abby, I know that we've talked about AI, should I start with you?

Abby Kearns:

Others go first.

Hillary Greene:

Okay. Who'd like to jump in?

Salil Deshpande:

Well, I just wanted to jump in and say that this bundling concern that Steven brought up is really a cross-cutting concern. It's a meta concern, not specific to AI. Obviously AI is just going to exacerbate it. But there are lots of clever ways in which the clouds do bundling. Some of it is technical where one service only works with other services of their own or users are strongly encouraged to use other services. But others are just economic and discounting. For example, Amazon has the enterprise discount program or EDP where you get can commit to a million dollars upfront, get major discounts, but only on Amazon products. Not even third party products purchased through the Amazon Marketplace, until recently. They changed that recently but with limitations up to a certain amount and then you only get 50% of the discounts that you get on Amazon products.

So I think using a combination of those techniques, I would say bundling is probably a more important issue than even the issue that I anchored, open source. Bundling is separate from open source and I think that that's more serious and it's cross-cutting across all the topics that we talked about.

Hillary Greene:

Right. Frédéric. You're muted.

Frederic Jenny:

I tend to agree with the fact that bundling is a really important issue and I would say that competition economics and competition authorities don't have very clear ideas about what are the limits, what are the benefits, what are the limits. For a long, long time we have kind of ignored the scope economies entirely. And also as has been said, we are in a world where consumers have a hard time expressing their preferences, which is also quite different from the non-digital sector.

Now as a result, we have made some decisions. For example, Microsoft was asked to unbundle the media player from the service that it was giving and this led to a catastrophic result or I mean a very negative result. Nobody wanted the unbundled version. And yet it was made in the name of an idea that consumer wanted to have more choice. Well it turned out that they don't always want to have more

choice and that therefore bundles can be a good thing. And I think that we are lacking, really, a good methodology to figure out when bundling is just an entire competitive tactic and can lead to restriction innovation, and when it is in fact something which meets the desires of consumers. So I think that antitrust has to work on this to try to find a better methodology.

Hillary Greene:

Great. Other folks? Since we're starting to get into things that are what I'll call cross-cutting, one of the things that has come up repeatedly is also this notion of when there is self-preferencing and what are the complicated calculus around understanding things that are self-preferencing. Another theme that has emerged, that I would love also if we could touch on a little more, is questions associated with lock-in and to what extent do we see that emerging within this space. Another concept that I'd love to flag for potential comments would be things relating to switching costs. These are things that have been present in the conversation that we had, but I want to sort of put a bow on it and put a name to it. Frank, please.

Frank Nagle:

Sure. Yeah. So I think, as you just pointed out, a bunch of the things we have been talking about are related to that. So the ability to move data is, if it's easy, then it's not locking people in and switching costs are low, but if it's costly then it is locking people in. And I think back to my gas station analogy, gas stations, again because it's a commodity the only way they have to keep you coming back is kind of location or rewards programs. And I don't think we would argue that rewards programs are anti-competitive, probably not. But are some of the other ways that we can think about lock-in happening in the cloud, bundling would be one and data transfer fees would be another.

I think the question is if we believe that the services are becoming more commoditized, is it going to be necessary for the cloud providers to do these things that toe the line to artificially increasing switching costs or artificially increasing lock-in that we don't like to see? I think that's one of the questions that we're at a point where that that's already starting to happen and it may increase, and that's concerning.

Hillary Greene:

Yeah. You said we're already sort of at a point where, what is happening? I'm sorry.

Frank Nagle:

Where the vendors are doing things to kind of artificially create lock-in and artificially increase switching costs. So to the point on data transfer fees, right, I agree it's not free for them to allow that to happen. But are they charging more than it probably costs them in terms of actual infrastructure to allow for data transferring? Probably. And that may be an interesting question for folks to opine on through the RFI.

Hillary Greene:

Right. And I want to get some more people's comments. But another word that I like to sprinkle like fairy dust in every conversation I've had with you guys is this question of innovation. What are the consequences for the conversations that we're having for innovation going forward? This is something that again has sort of bubbled up in our conversations, but I want to put a finer point on it—if that's something we want to make sure we are protecting, the ability to innovate going forward in competition for innovation. I see Frédéric has his hand raised, but I also wanted to flag something that Salil and I had talked about and invite him and others to bring it up, which is: there's the formal understanding of

antitrust, but also there is the layperson's understanding of antitrust and their own sensibility about competition. And I know that's something that you all have spoken to me about as well, so I thought I could invite folks to bring that up as well. Salil.

Salil Deshpande:

Well, my concern has been that if you look at each of these tactics and behaviors in isolation, you could argue that none of them is a problem. And if you look at each cloud in isolation, you could argue that none of them is a monopoly. But if you take the actions together and you take some of the clouds together, then you see the problem emerge. I think there is a problem. I can't quite describe the problem to my own satisfaction. There are occasions where these clouds seem to be using their privileged positions as the provider of network or infrastructure in their computing environments that allow them to engage in rent-seeking behavior for other software companies. And I'm wondering if our regulations are up to par address this new reality.

Hillary Greene:

Great. Other thoughts?

Salil Deshpande:

Yeah, I'd love to hear what the other panelists think of what I just said. Whether that concern even articulated in that way makes any sense.

Hillary Greene:

Okay. Salil has put something on the table.

Abby Kearns:

I would say I think it does make sense Salil. I just think it's a fine line between building value added services and capabilities and bundling those as solution and things that are actually shifting impacting the market. And I think that's a really, really tight line. Same with antitrust, there's a really fine line between market making and forcing people out. And I think that's probably in my head the most ill-defined right now in terms of where that line is because I do think the market has made certain choices about the cloud. I mean it's not like we haven't had other players that have built clouds. IBM wasn't that long ago. Verizon wasn't that long ago. These are companies that spent billions of dollars to also compete in the cloud and the market really chose those top three with a distant fourth with Oracle. That's just in the US if you think globally, Alibaba obviously has a massive footprint too.

And so I think for me, where it's not so clear is where that line is between where the market decides who the winner is based on where customers are gravitating, based on capabilities, natively integrated services, those solutions that they're bundling, versus what is actually really anti-competitive. And that's where it's not super clear in my head right now because I do think the market is making these choices.

Hillary Greene:

Okay. Steve, was that a hand?

Steven Weber:

It was. And I don't have an answer to Salil's very good question either. But I do want to say and maybe repeat that point that I think we ought to be looking at this right now from the perspective of these new

small AI application or large language model companies and from their perspective, do they see a competitive marketplace for the money that they need to spend and the services that they need to be able to access on compute and so on and so forth. We saw an estimate, I think it was from Gartner, that like 80% of the capital raised for AI startups is being spent on compute for training and deployment. Huge amounts of resource intense expenditure there. If the market isn't competitive from their perspective—they are in some sense the leading edge customer right now that I would pay a lot of attention to when I thought about: is the market competitive or am I being boxed in by different kinds of business model policies?

Hillary Greene:

Right. And I see we've flowed into our sort of closing remarks and so I'm going to go around the room. I see Frank.

Frank Nagle:

Sure. I was just going to add on to that thinking. So my mother used to work for the phone company, so natural monopolies are always part of what I think about and I see some heads nodding. So I guess that's one of the questions here, whether due to network effects or economies of scale, which are different but related, is there some sort of natural monopoly or natural oligopoly that these companies have to be huge to be able to do what they do? And I think that to me is a very unresolved question and has the kind of spillover effects to everything that we've been talking about.

At the end of the day, going back to where Frédéric was leading with kind of innovation, I fully believe that the cloud and the crowd in the context of open source are going to be what drive a lot of innovation, especially in the AI space, but in many other spaces as well. And so understanding whether this market is competitive enough or whether there's too much concentration I think is that the utmost importance to the future of innovation, especially in some spaces, but in most spaces as well.

Hillary Greene:

Frédéric?

Frederic Jenny:

I wanted to come back a little bit on the question of innovation because, A, I think it's really important not to prevent innovation. I think everybody agrees with this. But I think that, at least in competition, we have two problems. One of them is to define exactly what is an innovation in the digital sector or in the cloud. Is every new service or any new combination of service an innovation, is that what we mean by innovation or is there a more finer limit? And the second one, which is more general, is that as economists we don't know much about the relationship between competition and innovation. So we really have to learn a lot if we want because we are in a particularly dynamic sector.

The second thing I wanted to say was that in the interviews that we've conducted, we have found that there's a great deal of confusion in the mind of people between fairness and competition when it comes to relationship in the cloud. And I think that the reason is that there is no set business model that everybody can refer to and which delimitates, really, what is unfair, what is anti-competitive in a relatively neat way. Because there are many different business models that can be used, people are quite confused about what is the norm or what should the norm be, whether we think about fairness or we think about competition. So there's a tendency to, at least on the part of the users, even business users, to mix the two. What is unfair may not always be anti-competitive and vice versa. So I think that there is a need to bring more clarity there.

Hillary Greene:

Great, thank you. Others, closing thoughts? Salil?

Salil Deshpande:

I effectively gave my closing remarks while posing that question to the panel. The only thing I'd add is we need to take a closer look at what the different markets are or delineate markets so that we can look at whether dominance or market power in one area is being used to gain advantage in another. For example, I would say that the base layer of what the clouds offer, which is compute, networking, and storage, is different than the next layer up, which is say databases. Those, I think, are different markets. But to the layman they all seem like infrastructure, they all seem like the same thing. So unless we get a clearer picture on the delineation of these markets, we won't be able to answer the question of whether market powers in one is being used to gain leverage in the other.

Hillary Greene:

Thank you. Abby?

Abby Kearns:

Yeah, I agree with Salil's point and I think it's important to really parse out what we mean. But I'd say my final thoughts on cloud is it's not really a location anymore. What we're talking about is more of an operating model and it's become a bit more ubiquitous in the way we think about technology, but it's also closely aligned to the points that were made, which is the business models associated with them. And I do think that that probably merits more discussion as we start to parse out what does this all mean, rather than having a strong myopic focus on just the technology piece. But I think taking it in its whole would probably be time better spent as we think about, okay, what are the anti-competitive practices that exist across those?

Hillary Greene:

I think we're to Steve.

Steven Weber:

Is that to me Hillary?

Hillary Greene:

Yes, Steve.

Steven Weber:

Yeah. So I'll just add one quick thought, thanks everyone again for this stimulating conversation. Not going to solve any problems, I'm just going to throw another one on the table for us to think about as we leave, which is that, at least in my world of large language models and so on and so forth, the technology's evolving so quickly that even in the last couple weeks there's been a new and very vibrant conversation about running these models on the edge, smaller models rather than larger models coming out of the cloud, open source models leaked and then later essentially sort of legitimated the leak coming from Meta. Any conclusions that we come to today might be different six months from now. We're going to have to have regulatory oversight, customer behavior and government policy thinking

that is quite adaptive and very fast. And that may sound impossible, but that's what the market is demanding from us right now.

Hillary Greene:

And Frank, I didn't know if you had anything else.

Frank Nagle:

That was my-

Hillary Greene:

You're good. Okay, great. Well, things are moving quickly. One of the conversations that I've had with all of you in different ways has been a question of how quickly things are moving and the need to be mindful of that. Other things that have come up in all of my conversations with you have been, there are no easy answers and yet one still needs to make progress in thinking about things, notwithstanding the fact that this is not an easy area to parse. There's so much to say, on not only the topics that we've covered. As just one example I know that multi-homing has come up here. There's also been lots of conversations with others where they talk about is multi-homing feasible? Not feasible? Is it for some really a fantasy or is it practical? That's just sort of one example of how there's so much more to say on the topics that we've covered, as well as the topics that we couldn't address.

This panel could easily, as everyone said, could go for many more hours on any one of these topics. But do not despair I say. Why do I say that? Because this panel will come to a close, but the discussion that we're having will continue. This is the beginning of a conversation and, in particular, it is our great hope that folks will go to the request for information, which is on the FTC webpage, it closes on June 21st, and they will continue that conversation there as well. I know sort of the public writ large will join me in thanking our extraordinary panelists for a wonderful conversation. I wish you all could hear the applause because I do know that it's there. And I'd also like to acknowledge the proverbial village of people here scattered throughout the FTC that it took all of their effort to make this event happen. So I want to also thank everybody on the home front here within the agency.

Towards that end, what I can say is, take care everyone and, wait for it, keep your head in the clouds because you knew I needed to use that at some point. Thank you all again so much for your participation. The conversation will continue. Thanks.