

**EPISODE 1287**

[INTRODUCTION]

**[00:00:00] JM:** The cloud has delivered amazing benefits, like on-demand infrastructure that is easy to use, with pay as you go subscription plans and effortless scaling of applications. This flexibility minimizes the growing pains for businesses, and explains why today's startups and established companies are both building apps on the cloud. However, the costs of using the cloud stack up once companies reach large scale. Dropbox, for example, shifted away from the cloud in 2016 to opt for a custom-built infrastructure in colocation facilities. This was called Dropbox Magic Pocket. We covered it in a few previous episodes.

In that exercise, they saved \$75 million over two years, and increased their gross margins from 33% to 67%. And it's not just Dropbox who's attempting to make this off-cloud migration. It turns out that some major companies are spending 75% to 80% of cost of revenue on their cloud bills. Consider how much of a company 75% to 80% of revenue leads to. So should large companies shift away from the cloud like Dropbox did? Is it too late for some companies to untangle themselves from it? In this episode, we talked with Martin Casado of Andreessen Horowitz. We discussed the costs of cloud for small and large companies, and the financial implications of cloud infrastructure at scale.

Our first book is coming soon. *Move Fast* is a book about how Facebook builds software. It comes out July 6, and it's something we're pretty proud of. We've spent about two and a half years on this book. And it's been a great exploration of how one of the most successful companies in the world builds software. In the process of writing *Move Fast*, I was reinforced with regard to the idea that I want to build a software company. And I have a new idea that I'm starting to build. The difference between this company and the previous software companies that I've started is I need to let go of some of the responsibilities of Software Engineering Daily. We're going to be starting to transition to having more voices on Software Engineering Daily. And in the long run, I think this will be much better for the business, because we'll have a deeper, more diverse voice about what the world of software entails.

If you are interested in becoming a host, please email me, [jeff@softwareengineeringdaily.com](mailto:jeff@softwareengineeringdaily.com). This is a paid opportunity. And it's also a great opportunity for learning, and access, and growing your personal brand. Speaking of personal brand, we are starting a YouTube channel as well. We'll start to air choice interviews that we've done in-person at a studio. And these are high-quality videos that we're going to be uploading to YouTube. And you can subscribe to those videos at YouTube and find the Software Daily YouTube channel.

Thank you for listening. Thank you for reading. I hope you check out Move Fast. And very soon, thanks for watching Software Daily.

[INTERVIEW]

**[00:03:24] JM:** Martin, welcome back to the show.

**[00:03:25] MC:** Happy to be here.

**[00:03:27] JM:** There's a phrase that I have heard in the software business. And that phrase is it's easier to figure out a way to make more money than it is to figure out how to save money. Do you agree with that phrase?

**[00:03:39] MC:** I actually do. Yeah, I do agree with that. That's a good phrase.

**[00:03:43] JM:** And to what extent do you think that applies to the world of cloud?

**[00:03:48] MC:** Well, it's interesting. So all companies go through – You and I have talked about this before. But all companies I think go through three stages, right? They go through the product stage, they go through the sales stage, and then they go through the operation stage, roughly. And these stages – Or the first stage is finding product market fit. And that I think is just the hardest, it is the art. You're trying to build something that people want. And, frankly, most companies don't get out of that stage. But once you do, and you've kind of gone through that saga, you've built this DNA around a company of innovation and growth and building something new. And so then you move to the sales phase where you actually sell that.

And so I just think many companies, when they cut their teeth, they cut their teeth in this way of how do I add new value and how do I grow? And the best medium for that is absolutely the cloud. How can I set up infrastructure very quickly? How can I get access to services already running? How can I focus on my core competency, which is finding product market fit? I think the cloud is very, very good about that. But to your previous question, so you've got these companies that spend often – It could be a decade learning all the muscles for growth. The question is what happens when things saturate and they're no longer growing? What do you do?

**[00:05:11] JM:** A quote from your recent article, “If you're operating at scale, the cost of cloud can at least double your infrastructure bill.” So I can understand that operating a cloud at scale can double my infrastructure bill relative to running on-prem infrastructure. But I thought that this saved operational expense in exchange. So isn't that just fine? Like don't compensate for my increased capital expenditure with savings and operational expenditures?

**[00:05:44] MC:** Actually, the numbers are all in. So that's not the case. Forget the on-prem or not on-prem, because that actually wasn't the focus of the analysis. I think we should talk about that probably in just a minute. The analysis says, it turns out that if you reduce your cloud spend by 50%, you get this massive uplift in share price. And for the companies we focused on, which are public software companies, where cloud as a part of cogs, which is very important in this conversation. This can be on the order of billions of dollars. So what can you do with that billions of dollars? And this is all in. This includes operations, R&D. I mean, strictly, you've got higher share price. And that flows through to cash. You've got more cash flow. You've got more access to capital, either by selling shares with equity or taking debt against the new price. And, of course, you could set in growth if you wanted to. It's not correct to say that now you're doing more **[inaudible 00:06:45]**. And so does that balance it out? Because the analysis actually includes that in the uplift. I mean, this is strictly a gain in the economics. Now what you can say is, “But the company doesn't focus on this.” And so even though the numbers make sense, like this is a distraction. Like that's a valid kind of line of inquiry we can go into. But numbers are the numbers. I mean, like it's strictly all in costs. You save a ton of money if you reduce your cloud costs.

**[00:07:15] JM:** Can you define the term cloud repatriation?

**[00:07:19] MC:** Repatriation is where you take workloads that were in the cloud and then you move them somewhere else. Most likely, this means to some other hosted infrastructure that's lower level. Let's say you've got like some Equinix, or something like that, where it's probably somebody else is hosting the servers and doing the racking and stacking. I don't think anybody's talking about that. But it's somebody else's kind of bare metal, or maybe there's a few VMs running on it, or something like that, some minimal software layer, and then you move workloads off of the big three clouds primarily on to kind of a lower level infrastructure.

**[00:08:01] JM:** first example that comes to mind is Dropbox. And there's been a lot of coverage of the Dropbox move off of AWS, or more like partially off of AWS. The Dropbox case study is incredibly work-intensive, very difficult. How is that applicable to a broader set of companies?

**[00:08:26] MC:** Right. Dropbox, it turns out that S3 isn't very good for a certain type of workload where you've got a bunch of small files. And so that's why even though – So Dropbox did repatriate. And this is a very well-known anecdote. They **[inaudible 00:08:44]** and then they decided to move the kind of metadata layer back on the cloud, which makes a ton of sense, right? So they decided that's the breakdown.

It actually turns out, these names aren't public. So I can't go into it. There're other very large companies that it's leaked enough that most people will know or via Google search too have done the same type of thing. They've decided that S3 storage is too expensive. So it moves storage intensive workloads off. So this isn't unique to Dropbox. There're a number of companies that have done that.

But here's what's even more interesting, which is the reason that we put Dropbox in the article is just because the numbers are public. And we just want to be illustrative. And it had nothing to do with the analysis, actually had nothing. So if you look at the number – So let me just say what the analysis said. The analysis says that if you looked at public software companies today, 76% of which IPO'd in the last five years. If you just look at those, we looked at 50 of those. If you can reduce the cloud cost by 50%, what happens? It just turns out that you've got \$100 billion across those 50. That's being depressed by the cloud. Of those 50, we actually spoke with a handful. Let's say – I don't know exactly. Let's say eight of that we spoke with on that order,

every single one of them that we spoke to had a repatriation effort in place or they're planning on it.

And so even though Dropbox is a tired example, and it may not pertain to the general software, though it does certainly pertain to storage in terms workloads. This is a broad industry conversation that's happening, at least according to the conversations that we had. And the people that we talked to all said, "Listen, we calculate the savings is 2X to 3X for running our own infrastructure."

**[00:10:23] JM:** Maybe you can help share in some anonymized fashion what they talk about in those conversations, because to me, whether it's Dropbox, or any vanilla SaaS product, doing any kind of repatriation sounds really, really hard.

**[00:10:41] MC:** Yeah. So I think it's really important to also characterize the type of companies that we're talking about, right? And so we're talking about large public SaaS companies. And they're in the core business of offering a SaaS service, right? That's what they are. And these are deck of billion dollar companies that already have platform teams with hundreds of people. And so you're absolutely right, that repatriation is hard. But this isn't some startup. I mean, these are very sophisticated companies that are already running very sophisticated operations.

And so at some point, it comes down to economics. And so I think what's very illustrative is what happened when the clouds had this exact moment. So I think that this is actually – I think the article is a bit of a glimpse into the future. And I think what's illustrated is what the actual cloud providers did.

So I remember back in, let's say, 2004, 2005. Well, that's a little bit late. Let's say 2006, 2007. When the cloud providers were like – They did the same analysis that we just did. So we're talking Google, and Microsoft, and AWS. They're like, "Actually, if you look at our cogs, it turns out a significant part of it is the OEMs," right? It's Cisco, or the networking gear. It's whoever's providing our servers. That's making a significant percentage of our cost of goods. What do we do about it?

Now, it would seem absolutely ridiculous to build your own servers and switches. That sounds, to your point, very hard. It sounds, to your point, something that is not their core competency and a distraction. I mean, way crazier than a SaaS company building a data center. It's crazier. Like what does Google know about building network switches? I mean, this is hardware. It's totally different. And yet, because the economics were so compelling, that's exactly what they did. Facebook did the Open Compute Alliance, if I remember the name correctly, I mean, Google notoriously redid servers. They did a lot of the networking. They published about it. AWS did the same thing.

And so it turns out that once the economics are compelling enough, you have a case to do things that are not core. In the case of the cloud providers, they built switches and servers. In case of the SaaS providers, at some point – I don't know when that point is. At some point, it makes sense to reduce cogs by doing it yourself.

Now, the last thing I'll say is I used to run a \$600 million data center software business, right? I ran all of networking security for VMware. And all of that was on-prem, the majority of that. And on-prem data centers right now, I mean, this is \$100 billion business today. So there're a lot of people that are doing it and doing it just fine. I mean, the clouds actually don't have a monopoly on data center talent yet. So even though things are trending in that direction, I mean, there's a lot of proof points out there of very large data centers that are being run outside of the cloud.

**[00:13:29] JM:** And let's say you were running one of these large public SaaS companies that was doing a repatriation, how would you architect that repatriation? How would you plan it at a practical level?

**[00:13:44] MC:** So I'm not sure I would ever argue for repatriation. And we certainly, in the article, did not argue for repatriation at all. I mean, it was mostly an analytical device. Here's what I would recommend. The first one is to realize that cloud costs over time have a significant impact to share price. And that means free cash flow, right? I mean, that's access to capital. That's access to growth, right? And so you should plan for that.

I think a pretty straightforward way is you actually make cloud costs first order KPI or primitive, right? So we've talked to founders and business leaders who would actually spiff engineer. They

would reduce cloud cost through optimization. We talked to others. These are these are large, significant, successful companies that, as part of like internal dashboarding, would actually have the cloud cost as part of that. We had others that instituted a hybrid policy early on. And so I think there're a lot of steps that you can do to reduce cloud costs. That's that was the point of the article. It wasn't to argue for against repatriation. I have no idea that it makes sense for any given company to repatriate.

What I do know – And I don't have a stance on that. What I do know is, of the companies we spoke to, and we spoke to many, many either had or we're planning to repatriate. So it's clearly a conversation that's happening. Honestly, I'm not in the position to describe how one should go about that. I just know – This is more of a report. I just know that it's a conversation that's happening. And, financially, actually at some point makes sense.

**[00:15:25] JM:** As an investor, I'm sure you have seen the multitude of cloud cost optimization companies. Why hasn't there been – So like, at first glance, to me, cloud cost optimization looks like one of those areas that's sort of like logging where you have – Eventually you have some kind of Datadog winner, or “winner”, that takes a like maybe a plurality of market share. But you haven't really seen a breakout winner in cost optimization. Why is that?

**[00:16:00] MC:** This is such a great question. I actually think this gets to kind of the heart of all of this stuff, which is, okay – So I think that an analysis like this is really looking ahead five years or so, right? I mean, we're honestly still in early innings for the cloud. So the cloud not only is growing asymptotically, right? If you actually look at the growth curves, I mean, like the second derivative is definitely positive. I mean, it's super linear. But the clouds are getting very good at optimizing themselves as well. Costs for the cloud have dropped a whole bunch. And so that benefits to customers.

And so if you look at it right now, I mean, from a customer standpoint, costs are dropping, growth is going. There're a lot of services there. And so, listen, Datadog is fantastic. And there are a number of optimization companies. I mean, there really are competing with kind of a growing market where the platform's themselves are doing a fantastic job of lowering price, right?

The question happens is, what happens when things mature a bit more? Now what the clouds have been very good is at maintaining 30% margins. And so how do you get those? Well, you get them by, whatever, charging for new services, and maybe having higher prices for small customers. I mean, there's just a lot of room. Or you are more sophisticated than the rest of the industry, because you do do your own servers, or because you do have – Like there's a lot of ways to get to that 30% margin.

But, again, if you look at the OEMs, and what happened with the cloud service providers, it's instructive. At some point, at some point, things mature a bit. And at some point, there are viable alternatives, right? Actually, I was talking with George Fraser, who is the founder Fivetrans. He had this kind of great point, which I think is right, which is if you have an oligopoly, where you can maintain margins at 30%, as long as you've got a differentiated product, and you're doing great, and you're growing, that's totally fine. But at some point, the industry catches up. And that always happens.

And at that point, if somebody stands up an infrastructure that is similar, but has lower margins, and that's the goal, which actually happens in the pharmacy industry once you've got – You can actually create generics. Then, all of a sudden, that provides a viable alternative for people to move to. But that only happens once you actually start to see kind of the differentiation and innovation slow down.

And so this is kind of a long, rambling response, but it answers your specific questions. And I think it answers the follow-on question. So to your specific question, I think the reason is that no optimization company has become significant is because I think the clouds are doing a great job of lowering costs themselves through a bunch of internal optimization. And kudos to them. They've done a great job of that. However, they have maintained 30% margins, which suggests to me that as soon as that gap slows down, like they can't drop those prices, that alternatives happen, which in a competitive market, they should happen. That's the big question. So let's go four years in the future and let's say that the big three are clinging to 30% margins. They can't drop prices much more. There are competitive solutions out there that are at 15% margins. Will the big three reduce their own margins, or will they allow the workloads to move off? I don't know the answer to that.

I do know that in the case of the OEMs and the clouds, the clouds just decided to move off of traditional hardware. Maybe that will happen. Maybe the SaaS vendors are like, "Screw this. I'm not paying your 30% margins, big cloud company. I'm going to go to this other thing and pay them 50% margins because we're sophisticated enough, and we're fed up enough." Or maybe the big three will be like, "Listen, we're happy with less margins," and they'll stay on. But nobody knows the answer to that. But I do think that the finances are becoming sufficiently compelling that this eventuality will happen, like this question will happen. We're going into this paradox, this question.

**[00:19:53] JM:** The cloud cost optimization domain, do you see it? Like how investable is it? And how much like a consulting business is it? Is there a standardized way to attack cloud cost optimization?

**[00:20:07] MC:** Yeah. No. That's a great question. When it comes down to it, if you're really being grandiose, which I can't help but – We're talking about like a trillion dollars in value. I mean, in the paper, like we're very conservative with every number that we put out. I mean, it's not hard to say like there's actually a trillion dollars in market cap that's being captured by the big five, right? It's a massive, massive market.

And by the way, if you look at it, what is the combined market cap of like Microsoft, Google, and AWS? It's about \$5 trillion, right? So it's \$5 trillion. And so to say that 500 billion to a trillion is going from you know the rest of the industry into them isn't a stretch. I mean, that's 10% to 25%, right? So then there's a question, it's like, "Okay, is there an opportunity for startups to go after that 500 billion to 1 trillion? I think the answer is, "Absolutely." Whether this is surfacing cloud metrics, whether this is doing actual cloud optimization itself, whether it's providing pass for repatriation. I mean, I think there're a number of things you can do. There're kind of marketplaces, etc., etc., etc.. And so I definitely think this is investable.

But like I told you, I believe it's early days. I mean, I think this comes to a head in a few years. I just think that now we're starting to see the real impact. And I want to reiterate. Of the 50 companies really look at, 76% IPO'd in the last five years. This is a very new phenomenon, right? Of these 76% that IPO'd in the last five years, we're talking about billions of dollars of

market cap being depressed because of the use of cloud, right? I mean, listen, if you play this out in the future, something has to give.

**[00:21:50] JM:** So given the tone of the article that you're clarifying here, that this is not really about repatriation. It's not about everybody should go buy a cost optimization tool off the shelf. You're more appealing to the idea that, from the get go, and throughout the lifecycle of a company, you should be considering how to reduce costs on the cloud front, or how to avoid costs on the cloud front. So does that mean should there be a centralized cost optimization team in your company? Or should you have a KPI for every team to reduce cost by 10%? What should that look like?

**[00:22:30] MC:** Great question. Yeah, yeah. Yeah, it's funny. I mean, I've written a lot of pieces over the years. And this one has, I think, been probably kind of misquoted, like misread, misrepresented a lot. But like I also think it goes to show you that this actually is a very real issue, and people are really having the conversations. And it actually reminds me very much of the early days of software-defined networking, which is funny. In the early days of software defined networking, by the way, this is in 2006, 2007, it's almost the same. We were like, "You know what? Wouldn't it be better if instead of buying these like 60% margin switches from Cisco, you just built your own, especially if you're a data center?" And people are like, "No, that's crazy? Why would you ever do that? You don't know anything about building switches. You don't know anything about building routers," right? Yadi-yadi-yada. And then, of course, now like they did that, and it's very commonplace. This feels very similar.

And so the question that you're asking is what do you do if you believe the analysis? And the problem is – Or not the problem. I'm not sure if people are internalizing or believe the analysis yet. But if you assumed the analysis is correct, if you assumed that actually cloud is going to be 50% of cogs going forward, there's this big existential question for a company of what you do. And I think there is an array of options. And I think the one that you – I think there're two things that you should consider. I think the one that you just articulated as a great one, which is, yes, cloud costs have to be a core metric that is tracked, and it has to be good throughout the organization, right? We're actually also seeing this with analytics generally in the business, right? Like one big trend we see in BI is we're pushing it to all parts of the organization,

including product and engineering. I think the same thing needs to happen with cloud costs, for sure. So you really understand the impact of writing code.

But I think another implication is to realize that, architecturally, the business case may happen, may happen that you have to find an alternate means of doing infrastructure. And you have to answer the question of what that means. It may mean don't do anything. It may mean prepare for to happen. I don't know.

And so just to step back very quickly, I think the point of the article is to say we may hit some point in the future, I don't know when, where this is an industry and we have a shift that happens. I think there's some near term things companies can do to deal with it. But how that ends up playing out, I think it's still pretty early to say.

**[00:24:57] JM:** So in the analysis you did, did you have any conversations with the layer two cloud providers that you're familiar with? Like, obviously, Netlify is a portfolio company. Just be curious as to how they think about this set of problems, because they effectively make their own margin on top of the cloud margin.

**[00:25:19] MC:** Yeah. We talked primarily to big SaaS companies. And so that was the focus of it. The rough sketch was what I've said. You've got these big SaaS players, that if you look at their cogs, more than 50% of them is cloud. If you look at what that does to the share price, it depresses it by billions of dollars. Now what? That's what it was. And then we spoke to a bunch of them, they're like, "Well, we repatriate some stuff. We know it's an issue. We don't really know what to do." Etc., etc., etc. And so it's really just to highlight that eventuality.

I think you've got a great question, is what happens to the other infrastructure layers? I do think that one real eventuality is you do have alternate providers pop up that have different margin structures, because you no longer have an oligopoly. And the market actually takes care of this. So you actually have competition. You start to run margins. You have more alternatives. And I think those can look at any shape, right? I mean, we see startups that are like, "Oh, we do kind of OLTP as a service." And do they use the clouds? No. I mean, they often run their own infrastructure, or they use kind of a lower level hosting infrastructure. We see companies that are like next generation CDNs. Same type of thing. And so I actually think the clouds could be

dis-intermediated by a layer that doesn't impact cogs as much. But, again, now we're all guessing is what is going to happen because of this massive economic pressure that's been caused by the clouds?

**[00:26:47] JM:** Does something have to happen? I mean, because all the companies that you're looking at, I mean, still, like from a fundamental perspective, look pretty strong even with this massive cloud expense. Does something actually have to change here? Where is the pressure?

**[00:27:04] MC:** I think this is the question. As long as everybody grows indefinitely – Okay, let's just look at it from the view of the stock market. So the stock market looks at a company and it says, “You're growing nicely. I don't care about your cogs.” Is there any pressure to change? No. The stock market doesn't care. Your share price is great. And if your share price is great, you can get more cash, you can comp people better. You can get more debt. That flows through to cash. So that's great.

The question happens is what happens when growth slow and it starts impacting your multiples? And which we showed in the analysis is that, listen, if 50% of your cogs is the cloud, it's going to impact your multiples. And now you've got a question as a business leader, which is, “Okay. Well, am I going to do something about it? Yes or no?”

Now, 50% of cost – So half the companies we looked at, we estimated that cloud is 50% of their – That's an enormous amount, right. And so as this matures – Again, this is a relatively new phenomenon. As this matures, I do think there's going to be pressure on that. And I do think some subset are going to want to deal with the problem. This is already bearing out in the conversations that we have, which is like, these companies are aware of their business. They know it's an issue, and they're planning around this issue.

I do think that we're going to see something happen. And the question is, what is that? Is it the cloud providers are going to drop their margins? Maybe. Is it going to be there's going to be viable alternatives that Baidu steps in? Maybe. Is it going to be repatriation? Maybe. I haven't a clue what it's actually going to look like. But my sense is there's going to be enough pressure to cause something.

And again, remember, when this happened last time, when the cloud providers looked at their margin structure in the 2000s, they started building switching routers, man. Cogs are important, enough so that people will go build hardware, right? I just don't see why you wouldn't think that you're going to see alternatives for the clouds for the same reason.

**[00:29:02] JM:** Today, do you see battles between the cloud providers for big contracts? Like does Oracle Cloud fight versus AWS for a big contract with the SaaS company? Does DigitalOcean fight with Google for a big contractor? Does that just not happen because there's just – Yeah, does that occur or not?

**[00:29:24] MC:** I mean, for sure, there's competition for workloads. 100%. But I think there're two considerations we need to make. The first one is the premium products are the big three. I mean, I just think they've got differentiated products. They're the most mature. They've got the most users. They've got the most experience. And so I think that it's hard to compare like AWS to a different cloud, right?

And then, of course, the skills that a lot of people understand AWS, or GCP, or Azure, right. And so I actually think that they've got an oligopoly. So I think the big three, sure, they've got competition, but I don't think it's apples to apples. I think they've got an oligopoly because they were there first and they invest a lot in it, and they've got differentiated products, and they also have got people trained, and they've got a good moat. So then the question is, “Well, what about the oligopoly?” So you've got the big three. Why are they dropping margins?

So oligopoly dynamics economically are just different than free to market dynamics, right? If we can maintain our margins – If I'm AWS, because I know Google's pricing and I know Azure's pricing. But I think I can beat them on features, like why would I drop my pricing and like impact the stock price? And so you just don't see the same level of market efficiencies when you're dealing with oligopolies as when you're not.

And so I believe that it is oligopoly situation, which means that you do have three kind of differentiated products. And I do believe that that's the economic situation, which is why they've been able to maintain the margin structure that they have.

**[00:30:58] JM:** Did Kubernetes change the cost structures of running infrastructure?

**[00:31:05] MC:** I don't think so. I mean, maybe it made it more expensive. But I don't know. I don't know it's done anything. Listen, I think it's a great architectural principle. I think it's a great tool. I think it's relatively low-level thing. And I don't know if it's actually changed this dynamic at all.

**[00:31:23] JM:** Yeah. What about portability?

**[00:31:26] MC:** I don't think it's sufficient to provide portability. And I just think that that's such a tough problem that – I think it's a potentially useful tool in a very complex and big problem, to put it that way. So I think it's useful in that sense. And I think that the right team and architecture can use it in a positive way. And it definitely adds value in that way. But I don't think it's its own solution. I think that there's a much, much broader consideration if you're going to try and get the true kind of portability between clouds.

**[00:31:55] JM:** So, I have a variety of other questions that are sort of tangentially related to this subject. But we can come back to the subject. This one's a little bit far flung. But do you have any perspective on when the crypto world genuinely overlaps with the infrastructure world? Do you have any perspective on when we get the crypto cloud?

**[00:32:15] MC:** I mean, just a little bit about my background, which is, I mean, like I was basically distributed systems guy in the early 2000s, come networking guy. And then I built the company that did basically build software for data centers. Then I ran a relatively large business that does offer for data centers. And so I come from that world.

When it comes to crypto, there's a bunch of use cases where it's fantastic, right? Like the store of value, things like NFT, where you make trust assumptions. There're specific trust assumptions that you don't find elsewhere, right? Like money. I mean, like everybody can be a sociopath with money, and you want to protect from civil attacks. And that makes sense. And so I think crypto's fantastic at these types of use cases when it comes to identity.

When it comes to core, core infrastructure, however, I think one of two things needs to happen. Neither the crypto algorithms need to evolve to be competitive with a tightly coupled system. Tightly coupled systems are incredibly efficient, which is why you've got three basically big clouds and like peer-to-peer kind of that didn't work in that. And it's been tried many, many times. So either you have to kind of evolve the core algorithms to be sufficiently fast, or you've got to rewrite the workloads that you're running to be purely distributed. And not everything is purely distributed.

And so my sense is what's going to happen is crypto is going to chew-off the use cases that make the most sense. And there's a bunch of them that do. I think it's a very, very positive trend. And then I think rather than – My guess is rather than crypto all of a sudden being able to take on AWS, I think people will rewrite their apps to be more crypto native. And that's a long, slow process, right? I think that's a very different thing, right? Because you could believe like, “Okay, someone's going to come up with like the absolute right crypto protocol. So that when we all run it, it's identical to big crypto cloud in the sky, and then we can run our existing apps on it.” I think that's very unlikely. I think it's much more like here's our crypto protocol, and like you can write an app for it. And there's a bunch of apps it will support, but you have to rewrite your app. That's going to happen more and more. I mean, I can't tell you how quickly, but I would bet on that future much more than the big crypto datacenter is going to kind of pop up at some point.

**[00:34:38] JM:** What about, more broadly, applications of crypto to enterprise software? And just bear with me, because this sounds kind of ridiculous today, I think. I mean, so you think about like consumerization of the enterprise. That's happened, right? And who would have thought when Facebook was coming out that these kinds of principles would be applied to enterprise software very successfully. If you think about crypto, crypto is the atomization of incentives. And large organizations are all about incentives. Are there applications of tokens or crypto protocols that could have applications to enterprise software?

**[00:35:23] MC:** Yeah, I just want to say something that's not obvious here, because I just feel like this conversation has become almost a cliché.

**[00:35:30] JM:** Not investment advice? Is that –

**[00:35:32] MC:** No. I just feel like everybody says the same thing. What about supply chains?

**[00:35:38] JM:** I'm talking about supply chains? Definitely not talking about supply chains.

**[00:35:42] MC:** Here's the thing. I think the right way to evaluate crypto is on the trust assumptions. And there's been very few like truly federated systems that have penetrated the enterprise historically, right? So if you think about what federated systems? It turns out BGP is federated, right? It's almost counter-game theoretic, right? Like BGP. Everybody runs their BGP router, and they peer, and they trust each other, etc., etc. So there're some precedents, but it's pretty, pretty historical.

I think the reason why you haven't seen federated systems actually emerge a lot in the b2b context is because you actually have a structured relationship around payment already, right?

For example, if you're going to buy something from me, I'm going to have like your driver's license, and like we're going to exchange money, and like we've met each other. That's kind of a much more structured relationship there. So I think, strictly, in the b2b context, where it's businesses buying from businesses, I think there's a lot of laws and regulations and a slow sale cycle, etc., where I think crypto is less relevant in the near term. Where I think it's more relevant to your point is b2c, where businesses are trying to establish a relationship with anybody out there. And there, it makes a lot of sense, because not only can companies, whether they're brands or whatever, establish their own tokens and their own networks. You see a lot of companies jumping into the whole NFT thing, right? Like that makes a lot of sense.

I think in the b2c context, it definitely makes sense. In the b2b context, there may be use cases, but I think they're far less prevalent. And I think because you have existing procurement relationships in place, it just makes it less of an immediate need. So I focus pretty much on the core infrastructure. So I don't deal a ton with this. But if I were, I would focus more on the b2c use cases, I believe.

**[00:37:33] JM:** Alright. Well, closer to your specialization. You have done a lot of data infrastructure deals. And it almost seems like the data infrastructure space was kind of an emergent opportunity. And you guys captured a lot of the big deals, and almost to a point where

you have a real portfolio synergy in a gravity in the portfolio where you can presumably have some interesting information sharing. I just love to know, in this increasingly competitive venture world, how does portfolio synergy allow you to win deals that you wouldn't otherwise win?

**[00:38:17] MC:** So to win a deal, there're two big advantages to portfolio synergy. On the winning deal front, like the way that we think about investments is we look at spaces, and we do a lot of work in this space, right? And a space as large as data has got a ton of sub-spaces. Even like if you split data into like, basically analytics and AI/ML, you can have venture firms that just focus on either of those.

And so one advantage is simply that you really get to learn the spaces. And I think that's value to the entrepreneur, because if you invest, you can bring that knowledge both the background, but also the understanding of how the space has evolved. And so I think that provides value.

Another thing, and this is a little bit unique to Andreessen Horowitz, but I think it's becoming more common in the industry is, we have Andreessen Horowitz, as people know, as a platform, and we provide services to our portfolio companies. One of those services is basically a market network. So we can introduce them to a number of customers. And this has been hugely successful, right?

Let me talk to my company, Nicira. I mean, I think the first 10 significant customers came from Andreessen Horowitz's market development network. And so, for those, the more that we can kind of group like for like companies, I think the more that they can work together in a customer account, if that makes sense. And so there's almost secondary effect of it. Like, for example, we say we'll do a data week where we'll have all of our data companies meet with companies that are interested in this. It actually gives them kind of a better qualified conversation, because this is an active conversation in a certain space that's happening, right? Rather than just hear some random intro. I mean, I think it actually matures the market kind of in a bit of a different way. So I think those are the two primary advantages. We actually understand the space. And then when it comes to actually providing services for them, we can take advantage of the momentum.

**[00:40:16] JM:** Going deeper on the data infrastructure space, the Spark versus Snowflake ecosystem, are these disparate ecosystems? Are they mutually exclusive? Or are they overlapping?

**[00:40:30] MC:** Well, that is the big question. It's a great question. But that is the big question, right? So we roughly break data into two sectors. There's the analytic sector, which is more of a traditional sector, right? That includes like BI, and data warehouse, and reporting. And the primary use case there is kind of human-aided decision making. You kind of assume there's as a human being in the loop. The big use cases are like dashboarding and reporting. And then the other one is this more AI/ML world, where often more compute intensive jobs, maybe you're creating models, you make those models in production, you're dealing with unstructured data. It's kind of this new AI/ML is becoming a primitive in applications.

And so, it's pretty clear that these markets are separate and even complementary, right? We did this this big study where we talked to a bunch of practitioners and we found that a ton use both like Spark and Snowflake, or they'll have a data warehouse and they'll have like a data lake. And so I believe, today, these are two markets. Snowflake is much more on the analytics. Spark and Databricks is much more kind of unstructured data, operational AI/ML.

That said, it's pretty clear we're starting to see convergence happening. Meaning, let's talk about like the lake house construct, right? Which is if I can have a SQL layer that can talk to a bunch of different data stores, do I even need to have a data warehouse? And while it's still early days, I do think that we're starting to see that happen right now. And so the future is uncertain. But it seems to be that these are converging. But today, if you're to do a survey, I think you'll find very, very distinct markets.

**[00:42:03] JM:** How did the pandemic affect effect go to market strategies for infrastructure companies?

**[00:42:12] MC:** I think it's been actually a net positive for most companies, just because the infrastructure was really required to keep businesses online. If companies weren't talking about digitization before, they were because of the pandemic. And making that happen is just very

much an infrastructure thing. And so I think it was a net positive in that. There was a lot of pull from the customer trying to make sure that everything was up and running and online.

But the one caveat, which is the companies that relied on evangelical to direct sales and category creation environments, that was more tough, because those require trust building. They require multiple meetings. They require deep conversations. And that's just harder to do over Zoom and to do remotely. And I think the companies are trying to figure out how to do that.

What we saw across the portfolio is mostly reasonable to positive performance with a few that took a few quarters to figure out kind of the new method of selling. I would say, by now, it's all been positive. Like I think that it's been a net positive. Like sales are more efficient actually over Zoom. The companies that were dabbling with bottom up do a better job of it now. Infrastructure is probably arguably more relevant than it was before the pandemic, because it has to be remote. And so I would say that the pandemic, which was a horrible health crisis and public crisis, etc., but in general actually helped push infrastructure in the direction it was already heading.

**[00:43:44] JM:** As far as deal dynamics, since everything went to remote, did the cycle time of deals getting done get compressed?

**[00:43:56] MC:** You're talking about venture deals? Or are you talking about –

**[00:43:57] JM:** Yes, venture deals. Venture deals. Sorry. Not vendor deals.

**[00:44:00] MC:** Yeah, I don't have a broad answer for that. I think here's what did happen, which is I think you've removed a lot of standard channels for investors to source new deals or establish new relationships. And so I think you saw higher velocity with known deals. I mean, I think there's a lot more of like let's call investor X and talk about the portfolio. There's a lot more of going through like the list of companies you've already talked to. And so I think there was a lot more of like picking over the existing deals. And so I think that accelerated.

But I think, broadly, I wouldn't say things have accelerated. Does that make sense? There're a lot of deals being done, but I think it kind of forced investors to focus on what they knew, just

because there's a lot of trust building that's required and investing. And so, again, I don't want to be categorical. A lot of new deals were done as well. But I think that everybody has seen kind of the pace of deals. But a lot of these are deals that we all have known talked to, and this and that's. So I think that absolutely happened.

**[00:45:05] JM:** Was there anything that changed dramatically for you for how you think about investing or being a venture investor more specifically in the time of the pandemic?

**[00:45:16] MC:** I think my kind of macro view on what happened with a pandemic is there's a lot of – We're in this three decades march towards infrastructure everywhere and using the Internet and being remote. And we've painted a lot of lip service as a society. And it kind of forced us along an existing trend. And so I would just say that it's one thing to think about what it means to be remote. And it's another thing to be remote. And so I found that there was a lot of things that I thought were necessary that are clearly not necessary, as far as investing being present, this and that.

I think the bottom line is like I think we've all done the mental exercise in the last few years of, “What would it be if I did everything purely remote?” And you got to pressure test those ideas. And so I think, ultimately, I think it moved along a trend that we're all already on. And for myself, things that I thought that I would need to be in person to, like I got disabused of some of that. Like there're things that I thought I needed to be in-person that I don't.

On the other hand, there's things I kind of thought I could outsource, which I got disabused of that as well. If you're exploring a new area and you're working with multiple people, I think it's actually very important to have free-flow conversation. But, honestly, as someone who was traveling probably 50% time, being able to slow down on that and really focus on process was a huge wind. And so that was a big, big help too. And so the answer is yes, but maybe not in a significant – Not in a first order way. It was all kinds of second order stuff.

**[00:46:50] JM:** Gotcha. Okay, well, just looping back to the impetus for this conversation to close off. There are so many things to do, as a large company, if you're trying to do “digital transformation”. And most of those things to my mind required deploying more infrastructure, and buying more cloud, and getting CI/CD software, and getting lots of things that that cost

money. How should a large enterprise that is already in the midst of doing a digital transformation and spending a lot apply your plea to start thinking about saving money?

**[00:47:32] MC:** Oh, okay, great. So our analysis was not about the enterprise. And I don't even know if it applies. It probably doesn't even apply. So our analysis was very, specifically, if you're a software company, if you're a SaaS company, and cloud is a significant portion of your cogs, which is most SaaS companies, here's the impact to you. But it necessarily requires that the cloud be part of your cogs, like it's your product. Your product is a software product, and the cost of goods is cloud. And it requires you to be a public company, because this was like how the public market views your multiples. If you don't have those two things, the analysis is just not relevant at all.

So as far as the enterprise is concerned, I mean, it's a very, very different calculus. And I think that war is currently being battled, and these are very sophisticated teams, and I think they're probably converging on whatever is the natural outcome is probably the right outcome right now. Like that movement has been largely from on-prem to cloud has been a tough battle. I think it's the right one. I think the majority of their workload should be in cloud, that I strongly believe in.

But if you're a SaaS company and your product is software –Listen, if your product is cereal, you shouldn't be in cloud, right? It's not a part of your cogs, right? Like wheat is a part of your cogs. Like if you're if you're making car parts, you shouldn't be using the cloud, because the cloud is not part of your cogs. If you're building a SaaS app and 80% – We talk to companies where 80% of their cogs is cloud. 80%. This is a first class concern, man. This is 80% of your cogs. We're talking like potentially billions of dollars of market cap. Now it's interesting. And then what do you do? I mean, it's not super clear. But what is clear to me is we're going to see, I think, a reckoning in the next five year. Something has to happen.

**[00:49:23] JM:** Martin, thanks for coming back on the show.

**[00:49:25] MC:** Yeah, that was great.

[END]