

Company Participants

Thomas Siebel - Chairman and Chief Executive Officer

Conference Call Participants

Mike Cikos - Needham & Company

Mike Cikos

Great. Thank you to everyone for joining us today. I'm Mike Cikos, the lead analyst covering infrastructure software here at Needham as part of our tech conference.

I'm pleased to announce that we have with us C3.ai's Chairperson and CEO, Tom Siebel. Tom, thank you very much for joining us today.

Thomas Siebel

Morning, Mike.

Mike Cikos

Awesome. So this will be a 40-minute fireside. I have questions prepared on my side, but I want this to be as interactive as possible. So to the audience who's listening, you should have some sort of chat function in front of you. Please feel free to submit questions and we'll make sure we get to those while we have Tom here.

With that out of the way though, Tom, just high level for some folks who are maybe dusting off the story or just newer to it. Can you please give us an intro to C3 and how the business has evolved?

Thomas Siebel

At C3.ai, we've been at work now for 15 years and have spent over a couple of billion dollars building a software platform that allows organizations to design, develop, provision, and operate enterprise AI applications. So understand we began when we ideated this before AWS existed. We began before Azure, before the Google Cloud, before the GPU existed.

And so we're a little bit ahead of our time and we spent over a decade building this platform that provides all the services necessary to build these applications. And then we've used that platform to build today over 90 AI enterprise applications that address the value chains of consumer packaged goods, manufacturing, defense, intelligence, utilities, oil and gas, aerospace, what have you.

So as we power into the 21st century and the rest of the world has discovered AI, I think beginning about November of 2022, we're in a pretty good position to take advantage of this growing and very exciting marketplace.

Mike Cikos

That's great. And probably bleeding in right after that, but obviously Gen AI hits the airwaves. Can you talk about how this impacts or evolves the market opportunity that C3 is going after?

Thomas Siebel

I think the general predictions for enterprise AI applications get to about a \$600 billion to \$700 billion software market. Generative AI more than doubles that. So I think the predictions for generative AI applications are order of like \$1.2 trillion. So generative AI changes everything.

It dramatically accelerates the market. It fundamentally changes the nature of the human computer interaction model as it relates to these enterprise AI applications. And it is a very rapidly growing high intensity

marketplace in which we're very active.

I think we have 30 turnkey enterprise AI applications in the marketplace today for manufacturing, for aerospace, for oil and gas, for supply chain, demand chain, for Workday, for ServiceNow, for Salesforce and what have you. So that's a very large and growing market opportunity for us.

Mike Cikos

That's great. And one of the things that I wanted to talk with you as well is about the pilots, but maybe if we just take a step back rather than just digging right into pilots, can you talk about the transition in the revenue model and where these pilots, what are they expected to do versus what the revenue model looked like for C3 before that?

Thomas Siebel

Some quarters ago, as you're aware, I think about eight quarters ago, we changed our pricing model from subscription-based to consumption-based. Subscription-based being where people would pay like \$20 million, \$30 million, \$40 million, \$50 million in advance for the use of our products to a consumption-based model that made our products much easier to adopt where we would bring production-level applications live in say six months with a pilot for \$0.5 million. And then if they like it, keep it. And then they pay per CPU hour. So it's very much a value-based, usage-based model. It makes our products much easier to adopt.

The effect of it was, as we predicted, I mean, that the growth rate of the company had to diminish, okay. And it did diminish. It went down to below zero, I think. And then as we tiered in with a much larger number of transactions that, oh, by the way, a consumption model, a pricing model is revenue neutral to us over say 10 quarters versus the subscription model. It's just the revenue is now over or the payments coming in over time rather than the payment commitment coming in up front. But now we've seen, in recent quarters, we've seen an acceleration of our revenue growth rates from, I think, they went negative at 1.20% to 4% to 7%. I think last quarter, as I recall, the order of 16% revenue growth.

Mike Cikos

Great. Great. And I think that follows a similar theme we've seen in broader infrastructure software. The idea that, for right or wrong, the hyperscalers and their usage model is essentially trained organizations to think about infrastructure as a utility, right? Paid by the drink.

Thomas Siebel

Yes. Actual numbers I see, it went from Q4 of last year was zero revenue growth. Q1, 11%, Q2, 17%. And last quarter is 18% revenue growth, which would put us at or near the top deck aisle of the software universe. And so we're seeing a significant acceleration in revenue growth as a result of the transition now into consumption-based pricing. So it is working out exactly according to plan.

Mike Cikos

Right. And I'll say 18% revenue growth in the most recent quarter, 23% on the subscription side, which is exactly what you're driving towards. So I get it. I think with the most recent quarter, when you guys reported earnings in February, the company was citing momentum with pilots. And so can you talk about, again, we're a couple of quarters now into this transition, but how has the pilot generation gone versus your expectations?

The rate at which customers are engaging in these pilots, just because, again, the lower price point helps reduce some of that friction, and I would think accelerates velocity in the business, right? So how is that pilot going to do...

Thomas Siebel

There's not many transactions. We closed last quarter, but it was probably on the order of 40 or to 50. And back in the old days, it used to be on the order of, let's say 2019, 2020, when we went public, it might be an order of 10 transactions in a quarter. So we're seeing soon an order of magnitude more transactions in a quarter.

The pilots make it very easy for our customers to adopt our technology. The generative AI pilot, as I think, would bring a production application live in about 12 weeks for a quarter million dollars. A enterprise AI application would bring a production application. I mean a tough production application.

Think supply chain optimization. Think predictive maintenance for production line in a pharmaceutical company. We bring it live in six months for \$0.5 million. And then we're seeing about 70% to 80% of those pilots convert into production, which is about what we anticipated. And it's dramatically increased the number of transactions that we do every quarter.

Mike Cikos

And you are still seeing customers make multi-year commitments to C3 here. So I'm just curious like, what is driving that customer choice for those longer term commitments? It feels like it's a commitment to C3, right? So are there incentives to drive that, or is it just a function of, hey, I can capture a better price on the vCPU by going longer term?

Thomas Siebel

It is. It's a little bit counterintuitive because the idea is, if you like the application, keep it and pay \$0.20, \$0.30, \$0.40, \$0.50 per CPU hour and just pay for it by the drink. But many customers, after completing the pilot, decide that what they really want is price certainty.

And so now they're willing to engage in one, two, three-year committed transactions with a minimum commitment to us over that period of time, the trade-off being price certainty for them. And that actually came as a little bit of a surprise to us, but it is a significant number of the transactions are kind of multi-year commitments.

Mike Cikos

That's great to hear. And on the topic of Gen AI, I'd be curious. I know how new this all is, but I think one of the things, this is not a comment to C3, but just broadly, investors have kind of struggled with the idea that there's still a large amount of exploration taking place, whereas some of these projects are not yet moving into production.

It feels like it's a little bit different for you guys because you are delivering some of those turnkey applications, right? So are you starting to see an inflection, or have you already seen an inflection in Gen AI projects moving into production at this point?

Thomas Siebel

Yeah, Gen AI is huge. And it's everything from, let's say law firms, a leading law firm in the United States, we trained their enterprise learning model on the - you'll be able to relate to this, on the corpus of sec.gov. So we loaded every S1, every 10-K, every 10-Q to train the enterprise learning model. So when this law firm wants to write the next S1 for whoever comes public next, I don't know, Databricks or Anthropic or whoever it is, okay, they put in the name, address, the financials, the risk factors, hit the carriage [ph] return, and it generates the first draft of the S1 in an hour. Well, hey, that used to take seven associates two weeks.

So the, you know, other applications that are just, you know, kind of, they make all the sense in the world. You know, where we have, I think, the city of Riverside, or the county of Riverside, where we have all the county regulations, building regulations, statutes, taxes, come on, nobody understands that stuff. Okay.

And you load that into even the people who run the county don't understand it, really. And you can put, you know, all of the, all of that information into an enterprise learning model, ask any question about zoning, taxes, parking, whatever it might be, property setbacks, it gives you the right answer the first time.

Another example would be Affordable Care Act, Obamacare. These states that administer Obamacare, I mean, goodness, the rules, regulations, and documents associated with Obamacare are biblical in scale. Okay. And, you know, thousands and thousands of PDFs, websites, and documents about how to enroll, where to enroll. And the way that works today is you call a call center, in many states like California, the average call center,

call length is like 41 minutes, okay, with a call center agent, which I think you're on hold for probably 38 of those 41. And the average answer is probably 90% wrong.

Well, now we can load that corpus of information about the Affordable Care Act nationally, and the intersection of the Affordable Care Act, say with New Jersey or California, into an enterprise learning model. And you can ask any question about the Affordable Care Act. What's covered? What's not covered? What's in plan? What's not in plan? How do I get my grandmother in Fresno, you know, in the plan? What if I need dental work? And it gives you the right answer the first time. Even better, it operates in 131 languages.

So all of the documents in the state of California or the state of New Jersey, for example, are in any language you want, as long as it's English. Well, now you can ask the question in Spanish, you can ask the question in Mandarin, you can ask the question in Italian, and you get the answer in Spanish, Mandarin or Italian. So it's, you know, the applications of generative AI are really very interesting.

One of the things I'm particularly excited about is that in addition to those kind of no-brainer use cases, high-value use cases, I mean, let's look at Health and Human Services, Social Security, and Veterans Administration, okay? The budgets there are over \$2 trillion, okay? I believe they have over \$200 billion this year in the federal budget to increase the levels of customer satisfaction. Natural application for generative AI.

Mike Cikos

Yeah, yeah. I think one of the things that we're still trying to figure out on our side, and it varies where you're deploying it, right? But for example, there are companies for co-development where developers are seeing gen-AI result in a 50% to 60% productivity improvements. Or I think Bill McDermott over at ServiceNow has described gen-AI as a force multiplier.

So, curious to the extent that your customers have come back to you, you guys have done your own independent studies, but how is it you're quantifying or are you able to quantify what these productivity improvements from gen-AI are at this point? Or is it still - it depends on use case, it depends on application, it's still a little bit early to be going down that route now?

Thomas Siebel

Well, let's take the case with the law firm. I mean, it used to take them, you know, five associates, two weeks to generate the S1. Now we get the first draft of the S1 in an hour. Do the math, okay? Or it used to take 41 minutes, you know, for the Affordable Care Act to give the citizen the wrong answer. Now in two minutes, you give the right answer. Do the math.

I mean, so this is productively increased order of two orders of magnitude. So, I mean, it's huge. Or we're using one of our largest applications that we've talked about before is predictive maintenance for the United States Air Force, where we loaded all the data sets for F-15, F-16, F-18, F-35, KC-135, et cetera, into a unified image, including the telemetry.

And a B-1 bomber alone has, each one of these airframes has 42,000 sensors emitting telemetry in like 800 cycles. So this is a large data set. And what we're doing there, this is what they call readiness in the military. In the private sector, we call it predictive maintenance. We're identifying system failure and subsystem failure before it happens, maybe 50 to 100 flight hours, auxiliary power unit, flap actuator, igniter in the afterburner, whatever it might be. Well, if we can identify it before it fails, we can have the personnel and the materiel converge with the airframe, again, fix it that night in Munich, and it doesn't fail.

So we're increasing availability of aircraft in the United States Air Force by order of 25%. This is the largest AI, one of the largest AI deployments of any kind on earth, and the only AI system of record that I'm aware of at all of Department of Defense.

Now, the interface of this, like all the enterprise applications that we've seen, is pretty technical, and it's designed for somebody who is a technical expert in aircraft operation to operate. Now, when we put generative AI on top of that, the user interface becomes the mosaic browser. Ask any question in English and get the answer right now. So even the chairman of the Joint Chiefs of Staff can use this, and I'm not suggesting that he isn't a bright guy, because this guy is a very bright guy. But he's not used to, now he has that user interface that everybody in the world knows how to use. It's mosaic. I mean, you know it as the Google interface, but it isn't.

It came out of the University of Illinois in 1993, and it's mosaic. You type in any question, what are my readiness levels for F-35 squadrons in Central Europe? It gives you the answer right now, okay. Tells you where they are, it takes you to ground truth, it doesn't hallucinate, it doesn't de-exfiltrate, and there's no IP problem. So we've solved all of the problems associated with generative AI.

If I could just take one second for those of you who are interested in generative AI, I've been looking for a good book on this subject for the last year, and it's dribble, okay? There is one. Get it on the internet, okay? You can get it on the internet for it's written by the name Stephen Wolfram, who is a distinguished physicist and member of the academy. It's called, you know, what is ChatGP [ph] doing and how does it work? In 91 pages, okay, you will know more than anybody in your organization about how these large language models work, what they do, what they don't do, and to what extent is their voodoo involved. And there is voodoo involved.

But anybody, I'm actually printing 10,000 copies of this book just to give away, but I encourage anybody who's interested in this subject, and everybody on this call should be interested in the subject, download the book. If you don't like it, send me an email, I'll give you a money back guarantee.

Mike Cikos

Thank you. We get a fireside with Tom Siebel and homework to do as well. But no, if I could just come back to the revenue model, because I do want to touch on federal, I'll come back to that, but for the revenue model as well, as you guys are building these pilots, right, there's two benefits to the C3 model.

The first is customers bringing over more workloads, and then the second is increased consumption on existing workloads, right? So how do you think about the mix of those two pieces as far as drivers of growth? I feel like in the near term, it's probably more from additional workloads or applications over the longer term. It's just increased consumption of those applications. Once you have like a larger volume of those applications on the platform, is that fair?

Thomas Siebel

Yes. Let me say the real benefit of the C3 model is something different than that. Okay.

Mike Cikos

Go right ahead. Go ahead...

Thomas Siebel

It's the fact that this is a platform play. Now, of these 90 applications that we've built, predictive maintenance for offshore oil rigs, okay, fraud detection in financial institutions, supply chain optimization at cargo, these sound like very different applications, but in fact, 90% of the code that's running in all these use cases is identical across all of our customer base. So it is a software platform play, and that is a big deal.

Now, where is the growth coming from? The growth comes from new customers. The growth comes from additional customers of existing use cases like the predictive maintenance application for the Air Force, where we have, I think, a thousand users today, and I think we cover 22 airframes today, and we're going to 40 airframes, so that we'll get additional users there. Then we'll move, you can imagine that application will move from the Air Force to the Navy, Navy's got a lot of airplanes, okay, as does the Army, so that will move across the services. So that's additional use cases of the same application, additional users.

At the same time, we're talking about Department of Defense about, we're engaged with the Department of Defense, okay, in entirely different use cases, contested logistics for Transcom, contested logistics for DLA, generative AI for NRO, and so there's additional use cases, additional customers in use cases that we have, and you know, then growth through additional markets. I think we operate in probably nine industries today, utilities, oil and gas, manufacturing, aerospace, defense, intelligence, what have you, but soon travel, transportation, government services, holy moly, how big has state and local become? I mean, we started state and local, what, two years ago? I think last year, last quarter, I announced it was something like maybe 39% of our bookings. I could be wrong, but I think that's about right. So, I mean, that's a huge business.

Mike Cikos

And, all right. So that's exactly where I wanted to go next. Like, what's interesting, last quarter, again, you guys said federal sector revenue is up 100% year-on-year, bookings up 85% year-on-year, right? And the bookings from state and local plus 29%, federal A&D is plus 25%.

One of the things, and maybe this is no knock on federal or state and local or A&D, but they tend to be thought of as being behind maybe some faster moving other pockets or sectors within commercial. Are you surprised with how quickly federal is moving with AI? Or how do you explain the - because they are investing heavily in this right now. What's your take on that?

Thomas Siebel

I think as it relates to defense and intelligence, we're dealing with an existential issue. I mean, these people are preparing for war. I think there's a lot of concern in a number of theaters, I think most particularly China. And we are in an AI arms race with China.

So as it relates to these new technologies and the new generation of warfare, be it subsurface, be it space, be it cyber, be it whatever, be it aircraft, be it surface warfare, the infantry, AI is at the heart of all these systems. And so the people who run the services and the Pentagon, they know this. So they're investing in a huge way in AI and generative AI to make sure that we can defend the free world. And so there's a very real sense of urgency there. And the investments are substantial.

If we look at now DoD budget, I think is order, or excuse me, US federal budget is order of \$7 trillion. It used to be like 3, but last few years has come up to 7. Okay. And, yeah, and DoD is relatively rounding here. I mean, that's budget is order of, I think, \$850 billion or \$900 billion. I mean, all the money is in health and human services, health and human services, \$2 trillion alone. So it's all in social security, health and human services, veterans, what have you.

We look in that area, in those areas alone, I believe they have order of \$200 billion in the budget just to increase customer service. Okay. For health and human services, CMS, social security, veterans administration. Mike, that's going to be all about generative AI. So yeah, they're invested in - we look at what we look at veterans administration, which is something of a operational disaster. And, you know, they want to reinvent and digitally reinvent, okay, the veterans administration. Hey, that's all about supply chain, demand chain management problem, which is an AI application.

So, you know, I think governments are committed, states are committed, okay, to adopt AI to provide higher quality of services to their constituents, Department of Defense and the intelligence community has a mandate to defend the free world and AI is very much part of that. So we're seeing a lot of growth. This goes all the way down to state, local, law enforcement, code administration, what have you.

Mike Cikos

Great, great. And again, just given the success you guys have demonstrated there, I'd be interested, like, how do those sales cycles or negotiations differ from, not typical, but from private sector, right? Because again, you guys have been there now for multiple years and you are demonstrating strong success with government, federal, whether it's state or local. So how does that differ as far as sales cycles or negotiations?

Thomas Siebel

We have to pay our dues to get the door. I think we started working about 2013 or 2014, working with the federal government. I think it might've taken us until 2018 to close our first contract, okay. And then another and another and another, but just like the private, so it is, it's a little tougher to get started. Okay.

But once you demonstrate success, people know who you are and the sales cycles can happen pretty quickly. And so with our model where we're not, you know, we're not trying to go in the door winning, you know, \$100 million, \$200 million, \$300 million, \$400 million RFPs like, you know, like, you know, Northrop Grumman and Lockheed, we're starting off small and growing it, growing it, growing it. And I think that, you know, you will get to the applications like the United States Air Force, that's become a pretty big business.

Mike Cikos

Great. If I just shift gears for a second, a little bit off the beaten path, but I know you were actually one of the first people I think that was out there talking about data center constraints, right? Which, which is increasingly become a central concern among a lot of the AI players.

I guess, would be curious to get an update on your thoughts, but what is being done to combat some of those data center constraints? Like at what scale does this just become a more serious headwind or is it already a headwind?

Thomas Siebel

You know, Mike, I think we've moved beyond that a little bit. And I don't think the - the perception of today is there's a constraint on data center and a constraint on Silicon. I don't, and yes, and there is a constraint on the - at the - as it relates to Nvidia and how fast they could get TSMC to build GPUs for them.

But I think the real constraint is not that, the real constraint is availability of power and the grid infrastructure is maxed out and the people who operate the grid in the United States have been focused on things other than providing more energy, you know, at lower cost.

And so the base load is kind of constrained. So if you want to build a data center, okay, in the Silicon Valley today, for example, the only thing that prevents you from doing that is availability of power. The utility operators here cannot provide you power for a data center. And this is true all across the United States, right?

I think the real constraint is not going to be data centers. It's not going to be Silicon. Power is a very real constraint. Unless we get after that, you know, this is going to be, you know, it's going to be a real problem.

Mike Cikos

Yeah. And...

Thomas Siebel

Not only for AI, I mean, really in the big picture, who cares? I mean, we need to run hospitals. Okay. We need to run, you know, we need to run police departments. We need to run the military. We need to run trucks. And the power issues in the United States and Europe are becoming really quite critical.

Mike Cikos

Yeah. That's actually where I was going to go next. It's not just a US phenomenon. Like I think Ireland has a moratorium essentially on - they're not letting you build new data centers in Ireland until '27 or '28, I think it is. So I hear your point.

Thomas Siebel

I mean, you shut down all your nuclear reactors, you shut down all your coal reactors, and you think you're going to power Germany with solar power. Hey, you ever been to Germany? Okay. How many months does the sun shine? Okay. Not 12. Okay. So let's get a grip here. And so you're right. It's not simply US. It's, you know, non-rational policy that is going to become an economic reality. It is an economic reality today.

Mike Cikos

I think one of the other things, if I just come back to the model here, I know that C3 is expected some short-term pressure on operating margin, partially due to investments being made to upgrade customers to the latest platform. You guys put a lot to this Version 8.3 [ph] Can you explain what is required to have those customers get upgraded to V8.3 on their part?

Thomas Siebel

Yeah, we are in the process of moving all of our customers over onto Version 8. And Version 8 is, I don't know, four or five years worth of work of engineering. And it's an order of magnitude faster. It includes generative. It is a - in every way, a superior application.

That being said, it was a major rewrite of our application foundation. And the applications are not backward compatible. So we need to re-factor the existing applications that our customer sites to run on Version 8. And it takes a little time and we're making a lot of progress. And we're investing in all the customers to make sure they get on to Version 8 quickly. We're on track and it's all good. But it is not - it is not, you know, Version 7 applications are not - Version 8 is not backward compatible in Version 7, like, R3 [ph] was not backward compatible to R2 at SAP. And this is just the reality of sometimes when you make a dramatically improved product, it doesn't have backward compatibility. And so we're making some investments and it has some impact on gross margin. That being said, I think the financials look pretty healthy to me.

Mike Cikos

Right. On that upgrade path. So is the expectation that all customers will upgrade? Like are you - do you have a timeline for when you could potentially deprecate some of those older versions that you had been previously servicing in Version 7?

Thomas Siebel

Without having the specifics in front of me, it would be amazing to me if this fiscal year that began May 1, if all of our customers will convert to Version 8, it would be non-rational for them not to. I mean, it's just much more cost effective, it's much more powerful, it's faster, it's easier to use, it has a lot more capability. And you can expect this year that 100% of our customers will convert to Version 8.

Mike Cikos

And with Version 8 too, I guess maybe a little bit more forward looking. But is the expectation that, hey, if whenever we come out with Version 9, whenever that might be to build in backwards compatibility or. No, is that just the nature of how this is built...

Thomas Siebel

At all costs, we really try to get backwards compatibility with every upgrade and we're not planning anything in our horizon right now that's not backward compatible with the previous release.

Mike Cikos

Got it. If I just think about the management team over at C3, obviously we have the new, well, not new. Your Chief Accounting Officer, Mr Lath, is now the CFO. I think he assumed the role in early March, right. So the question would be how the first two months have been since he assumed that role, and then what do you expect him to bring to the position?

Thomas Siebel

Well, Hitesh is a star. I mean, he spent 23 years in public accounting at Ernst & Young. He was responsible. You know, he was the senior auditor in charge of the Oracle account as an example, so the guy's been in and around the software business. He is an expert in accounting. He is surrounding himself with absolutely an expert team. He was the perfect guy to transition into the CFO role as he was our Chief Accounting Officer and he knew the organization.

Hitesh is a star. He is surrounding himself with very talented people. And I'll be surprised if he doesn't build kind of the most powerful and productive F&A organization in the software industry. So Hitesh is a great addition to the management team.

Mike Cikos

Got it. And just rounding out the management team here, but a little bit different. I know that the company recently appointed Alan Murray to the Board of Directors as well. So twofold question one, with Mr. Murray joining, did you actually increase the size of the Board or was that a replacement? And the second is, what was the attractiveness of adding Alan to the Board of Directors?

Thomas Siebel

Alan - we increased the size of the Board with the addition of Alan. And I mean, Alan is, to the extent, you know Alan, he is a very bright guy. He is very well read, and he is extraordinarily well respected by virtually every CEO in the world. So we are extraordinarily fortunate to have Alan as joining the Board of Directors of C3.ai.

As you know, we do have a very distinguished board, and I think it's, you know, more distinguished with the addition of Alan. So we're - that's, I mean, he's a really neat [ph] human being, a very bright person. Extraordinary, well respected, and I'm certain he will bring a lot to the table.

Mike Cikos

That's great. We probably have time for maybe one or two more questions. Just as a reminder to folks tuned in, please feel free to submit questions. I think, at least on my side, if I could just shift over to the go to market.

Obviously, you guys have been fine tuning this model now for a couple of quarters. But just as a reminder, as we're coming into a new year now, those sales reps, what is the compensation basis when thinking about pilots and conversions? Are they compensated based on driving pilots and then consumption thereafter? Is that a fair way to think about how they're being incentivized to drive business to C3?

Thomas Siebel

Sales reps are compensated on revenue. It's very simple. It's like an old 1980s, 1990s enterprise software model. So it's very simple. They're compensated on revenue, revenue from pilots, revenue from consumption, revenue from production applications.

As an investor, final thought, if we want to look at C3 and we look at the value stack of AI, which I think most people agree is not ephemeral and most people agree today that this is a large, addressable market opportunity. So at the bottom of this stack, we have silicon. Above that we have infrastructure. Above that we have foundation models. And on top of that we have enterprise AI applications. C3 AI place at the top of this stack, enterprise AI applications.

Now, just like the PC market, just like the mini computer market, the early stages of the market, most of the value was in the silicon and the infrastructure. Think about the PC that you had in 1990, okay. Most of what you paid for was hardware and the CPM or the Windows operating system that ran on top of it. And you might have \$200 million or \$300 of software running on top of that computer from BusyCal [ph] or wherever you got, okay.

Now, that PC on your desk, it cost your company or cost you in terms of appreciate [ph] about \$200 a year for that hardware. It cost you about \$200 a year for the infrastructure. And by the time you add up all the applications that you're running on it, be it Bloomberg, SAP, CRM, what have you, you have \$8,000 a year in applications.

Well, this same game is going to play out in AI as we move into a trillion dollar marketplace. The bulk of the value is going to be in the applications. The silicon will get commoditized. It always gets commoditized. The infrastructure gets commoditized. It always gets commoditized. What doesn't get commoditized in the long run is the applications, and that's where C3 AI plays. The game that we're playing is to see if we can establish and maintain a market leadership position in enterprise AI. Hey, we have 90 applications in the market today, very successfully.

And so that's the investment thesis. And we might fail. We might end up being number two or three. Okay, there's your bear case. So it, you know, that's what the big picture looks like, Mike, and that's what we're after, and we're getting after it. We're having a lot of fun. I think we've been leaders in the industry for many years. Hey, I was talking about enterprise AI before anybody talked about AI. And we're now, you know, people talk about the AI Hype Cycle. I'm sure there's some truth to that. But we're in the AI implementation cycle at very large scale, in small organizations, in medium sized organizations, and in large organizations around the world.

Mike Cikos

Awesome. Well, with that, we'll close it out, Tom. But thank you so much for the time, and thank you to the audience here as well. Great. Thank you to everyone for joining us today.

Question-and-Answer Session

End of Q&A