

Satya Nadella on AI: @NoPriorsPodcast x Latent Space Crossover Special at Microsoft Build 2026

41 MIN · YOUTUBE · [HTTPS://WWW.YOUTUBE.COM/WATCH?V=CFNI2FORAc0](https://www.youtube.com/watch?v=CFNI2FORAc0)
<https://www.youtube.com/watch?v=cFNI2FORAc0>

SUMMARY

Satya Nadella discusses Microsoft's evolving AI strategy, emphasizing the importance of creating an ecosystem that allows companies to leverage AI effectively. He highlights the need for a robust training strategy for AI models, the significance of private evaluations for intellectual property, and the transformative potential of AI in various sectors, including coding, productivity, and education.

- Microsoft's AI strategy focuses on building an ecosystem for companies to participate as first-class participants in AI development.*
- The MAI models emphasize clean lineage and the ability for companies to create specialized AI applications.*
- There is a shift towards using AI to augment human tasks, particularly in coding and workflow automation.*
- Nadella stresses the importance of a multimodal harness for integrating models, data, and tools effectively.*
- The future of engineering roles may evolve towards generalists who manage AI agents and infrastructure, with fewer specialized roles.*
- Community impact and tangible benefits from AI are crucial for gaining public trust and acceptance.*
- The education sector may see significant changes as new models for learning and credentialing emerge, driven by AI advancements.*

[music]

[music]

>> Please welcome Swyx, Sara Goa, A Lad Gill, and Chairman and Chief Executive Officer of Microsoft, Satya Nadella.

>> [cheering]

>> Hello. Satya?

>> Uh I'm so excited to be here. Welcome to a crossover episode of No Priors and Latent Space with Satya Nadella. Um congratulations on an amazing build.

>> No, thank you so much and it's great to be with both of you. I listen to both of you or both the podcasts all the time. It's great to be on it.

>> Thank you so much.

>> Sara has been talking about um these amazing uh announcements from across the Microsoft estate all morning for I think 3 hours. What is the uh what's the most important reflection or takeaway you have?

>> I I'd say there are uh

perhaps the the biggest one for me is let's sort of conceptualize this more as an ecosystem play as opposed to a single model or even a single platform, right?

I mean

yeah, whatever I at least for me having grown up at Microsoft, having seen whatever four major platform shifts, uh I sort of fall into that uh

uh camp where a platform is defined by fundamentally its ability to create more value above the platform versus what's captured in the platform. And so if you you view what's happening right now, I think this morning's keynote was how can any company whether it's an AI-native company a traditional enterprise company, participate as a first-class participant where they can point to AI they created. Right? It's not that they don't use other people's AI. Of course, they will. But to me, what's the path? What's the recipe? How do I do it? What does the stack look like? What does the tooling look like? What is valuable? How do you do that? That's it. That's sort of our job to do.

Yeah.

>> Ecosystem strategy is uh very complicated, right? Because you end up building certain components, partnering

for certain components, supporting them.
You just announced this big suite of
models. Like, tell us a little bit about
the

uh training strategy for Microsoft.

>> Yeah. So

So, the thing that we wanted to do with
the MAI models was to build, and as
Mustafa talked about, first of all, a
great lineage, right? Starting with
pre-training

uh with very good data quality,
uh doing all the ablations, making sure
because in in some sense, it's become
even harder to build a clean lineage
model. Yes, because there's so much
stuff out there

uh that you truly need to ablate out to
be able to have a fantastic pre-trained
model. In fact, that's one of the
challenges of a lot of the open weight
models is they look great on one
benchmark or two, but they're not great
on practice. So, that's why, in fact,
even in the RFDs are
are pretty gone really excited about
these MAI models because how the heck
can a small 5B model hill climb?

Uh and it goes back a little bit to what
I think is ultimately the key thing to
do, which is try to pursue finding that
cognitive core. Uh so, to me, starting
with a clean lineage, then creating that
ability for companies to be able to use
this, right? Not just as a generalist,
but to create their own specialist by
building this hill climbing scaffold
around it, right? So, it's not just the
model, but you have a hill climb
scaffold around it, then
you will start building your RLE. You
will start collecting the traces. Most

importantly, you'll have private emails because we know all the emails out there are good, interesting, but they're not really that critical at this point because they all can be maxed. And so, the point is each company will have its own private email. And so, that end-to-end platform story around our models is sort of uh what I think is interesting. And then the other thing, Sarah, is since you brought that up is

I do feel there's a new frontier. Like people talk about the frontier and are you operating at the frontier?

Um interestingly enough, if you add a little temporality to it, you can use, let's say, in in in fact, that the Land O'Lakes demo we showed was pretty cool. We used whatever GPT-55, right? Then you collected a bunch of traces, and then you took a 5B reasoning model and achieved higher uh so, that is another aspect of what it means to appear you know, operate at the frontier.

>> Yeah. I I think uh I first of all have to congratulate you on basically building a frontier neural lab inside of Microsoft in 2 years. Um I'm wondering, you know, you have all this AI strategy that you're rolling out. I'm wondering, what do you know now that you wish you would tell yourself 2 years ago where or 2 or 3 years ago? 3 years for the Jensen partnership, 2 years for uh MAI.

>> Yeah, I mean, I think the the thing with that I reflect quite a bit, right? Which is sort of obviously I got into all this when I got excited by the the scaling laws paper and, you know, when, you know, even the OpenAI partnership came

about when those folks said, "Hey, we're going to really throw a lot of computer transformers." Uh and they've held, right? The thing that I always look back and say, "Wow, these things um do have capability that they're climbing up with I mean, this, you know, this crude way of saying it is intelligence is log of compute kind of works.

Now, what I think we underestimated perhaps is the real world complexity of deploying these so that they actually deliver the value in the real world, right? So the outcomes as measured by any benchmark is interesting, important, but the true eval is when people out there are able to do unique things that they only can value.

And it's very measurable.

Right? That I wish we had sort of even like had more in our consciousness, right? Which is as an industry because right now I think when people say, "Wow, I don't want a token max." It's an artifact of us not having thought ourselves as an industry that we are using tokens to create value every step of the way. So I think that's kind of what I wish we had gotten there, but I'm glad we are here.

>> What are some other use cases that you've seen that have created the most value for your customers? Because I know that people talk a lot about code and I think it's pretty clear that that's something that's having very large-scale impact. Are there other areas that you find in common that your customers are really benefiting?

>> Yeah, I think yeah, to your point, obviously coding is now got but it's

interesting by the way Elijah to even talk about the coding, right? Which is coding is worked so well that we now have to rebuild the IDE, right? I mean it's kind of nuts to see what we saw large is like, "Oh my god, I have these 100 agent sessions. I the cognitive load it transfers back to me as a human is so excessive that now I need a new UI. Uh oh, by the way, I like the the chat as the only artifact is also impossible. So that's why we need a canvas. So it's kind of interesting for all the things about where is software needed or where is UI needed? Uh you kind of need that even for code, right? In a fully agentic world. But that said, one of the things that we are starting to see, we started seeing with co-work, but even some of the work we showed with auto autopilot, right, on what you see with claws, is a good one, because if you sort of think about a lot of human capital is doing the glue work, right? If you now can augment that with tokens {slash} agents that are long-running, durable, right, then your ability to scale even what is still judgment and glue work gets amplified like coding does. So you can like I'm positive that 6 months from now we'll all be saying, "Oh, wow, like all through night the night there was a bunch of stuff that all these autopilots that I have working on my behalf with my delegated authority, so to speak, right? I can sort of given even my identity, did a bunch of work. Then of course I'll need my new ADE to say, "What did you do?" Like that. My Did I do this work and so on. So I think

that that's where compressing of workflows, completing of tasks, that's where I think a lot of the value gets created.

>> you raise a really interesting point, which is there's the actual agent is doing the code, and then there's a harness around it. And that's the environment, that's the context, that's everything you're setting up as a developer around actually a coding agent. What is the harness for the enterprise? Is there an equivalent concept for broader productivity work, or how do you think about that concept sort of generally?

>> So so some sense you kind of want the harness to define the models, the the data, and the tools, and so that you have a loop across those three. And so what we are trying to first of all make sure is each of our products that we build, right, whether it's GitHub Copilot or the security copilot the stuff we showed with M-Dash, or even the discovery for science, it doesn't matter. All of them are multimodal harnesses

um

tools access so that you can do this progressive uh disclosure of tools even so that they're token efficient. Uh and then you're feeding it with very rich context because that's sort of the other hard lesson we've learned in the last 2 years is oh my god the amount of work you need to do to prep the context layer uh such that your plan can execute in the most efficient way is where the magic is. So we have in our case we have the get up harness which essentially we're using across all our products.

It's available in foundry and we're open like you can use your llama harness whatever or you can use the um you know any open harness or any harness of yours and train with your tools and multiple models and your context. And so that's the pitch because right now a lot of dialogue is um hey if I train the harness plus tools and the model together you get evals. And what we're proving out is and the best example of that is what we did with M dash right because when it launched uh it found bugs or vulnerabilities that were not found by mythos.

Uh and so there is existence proof I would claim that you can have a multimodal harness uh that can in fact be more uh performant in the real world.

>> So premise behind the uh training at the independent frontier labs is really you know we're going to have these models and we'll have an API business and we'll support enterprise and startups but a first party product be it productivity or code or search drives the majority of revenue. That's a different value equation than you're describing. I think with the Microsoft ecosystem uh if if that's the case tell me if it's the case cuz obviously you have first party products and you have enablement products. Um what is the role of the developer like what's going to be hard and the set of skills and the value capture the developer has in that world?

>> Yeah, so I think that there's always going to be the case that someone who's super successful and as a platform builder can also have first-party products.

It was true with Windows, it was true uh with uh the the SaaS side and the cloud side as well with us and others and so on. But, the thing that is it should not be a limiter to other people achieving that same success, right? That I think is the core difference, which is the the network effects this time around around intelligence are such because they learn from data and not really lots of data. It's just the few samples that you have to see to understand what's novel about something. So, that's why the game becomes how to protect, so that's why I would say every company having private evals maybe the biggest IP. Right? I think about it. Like, what's that private eval that you can then use even a frontier model to hill climb on and not leak the traces maybe one of the biggest drivers uh of IP. Like, so in other words, another acid test is you have an eval that's private. You're using a a model A. Can you switch it to model B and you know, climb up? If you can, then you're in control. If you can't, you're not in control. And that's where even the harness becomes super important, right? So, therefore, having an open harness, letting all models come in, having your evals, your contacts, your tools help you hill climb, I think is the skills that an AI native startup needs, a SaaS company needs, or every enterprise needs.

>> Yeah, I think in a very real way, you are Microsoft historically as an operating systems company and then becoming a cloud company. Maybe like the third act is that you're harness or eval's

company. Whatever whatever the the sort of conglomerate of concepts that you want to put together. Um I I I think like enabling every company to have like frontier intelligence or what what I forget the the exact term that you used. Um is the is the mission, right? That is that is the platform promise that you build with us, you will get your intelligence for your data.

>> That's it. That to me that is the like if there was one tagline for this entire developer conference is can everybody operate at the frontier with their frontier intelligence, right? To me that is so important because otherwise I I don't know how you achieve stable equilibrium, right? Which is how do I then go and say, "Wow, my company is going to have a terminal value because I now know how to continuously compound on top of what's a platform that gets better, right?" So when like Windows obviously came out, Adobe built, Autodesk built, or even like take but Jensen said, "We built DX." And he built, you know, CUDA on top of it.

Right? I mean, I always say to Jensen, "God, I got the short end of that." But I wish we had recognized it. But nevertheless, but that idea that you can build a platform layer that someone else can then extend out and build their own intelligence layer in this case, I think is everything, right? Without it, why have a developer conference? I can just come and have you all sort of just worship at the altar of one model. But that's not a developer

conference.

>> Uh backstage we you had a discussion about what is IP or what is the the value in a company. It used to be the length of human experience at a company. And now it's this other thing, which is the eval's, the experience in sort of applying agents to the company. He I just want you to like flesh that out a little bit more cuz

>> Yeah, it's It's way to frame it, right? Because you have at the end of the day every company is going to have both the human capital that is still going to be super valuable because humans and their ability to find the gaps that exist at all times is going to be the way we all will create value, right? I mean so I'm definitely in the camp that this is going to be about expressing new forms of human agency and ambition even as token capital goes up, right? So let's say

any corporation has lots of tokens and lot of human capital. The question is how do you compound the two? So if you have a like if you take in teams, I have a bunch of agents doing work and a bunch of humans doing work and the traces between those

that is really important context of the how that enterprise is creating value.

Then that goes back to train not a generalist model, but to train the train the company veteran agent.

Right, that is super valuable again, right? Which is when a company goes says it should in fact go on to the balance sheet is how I think about it, right?

That's what in fact they may be like human capital was never possible to go put on a balance sheet

because you didn't know how to capture the tacit knowledge, whereas now I think you can with the agents that have learned through the through time, through all the traces.

So that's what at least we think will happen.

>> I think the SEC is going to have to have accounting standards for token expertise.

>> Uh

you're talking about the equilibrium state

and a stable equilibrium where companies have this compounding value and can see terminal value for themselves. Another challenge to you know, the considered equilibrium of okay, there are applications and workflows that are sort of common to a vertical or a horizontal and this was like the generation of SaaS companies and you know, Microsoft has lots of SaaS properties as well. And then there are things that are very specific to every enterprise that they're differentiated against.

Um I'm sure you have heard much and participated much of the debate about the end of software because all these workflows are are cheap to generate now.

Um do you think the equilibrium looks different between what agents get built in enterprises versus in their vendors in the future?

>> Yeah, so I think what's happening there is see we we had a particular way we captured

uh

I would say workflow

in apps, right? Because we built a a data model, right? We schematized some part of some business process.

>> Mhm.

>> We then built a bunch of business logic.

>> Yep.

>> And then we put a bunch of UI on top of it, right? So that's kind of what every SaaS company

>> configuration, right?

>> 20 20 years that was

>> Right.

>> And that was it. So interestingly enough, now you kind of get to re-litigate that vertical stacking, right? So I still think, for example, that data model that you built underneath every SaaS application is super good, right? It like why reinvent it? Like I I my general ledger better be a general ledger. I don't need new schema creation.

>> Yeah.

>> Uh in fact, that entity relationship uh is actually pretty good robust thing that I want to feed.

>> And you want it to be stable.

>> That's right.

>> Yeah.

>> Then same thing with business logic, right? If If you look at uh we have this product called Power BI, right? It is like dashboards galore people created. The beauty underneath that dashboard is a very rich semantic model, right? Someone took the pain to create a dashboard and do all the measures. And you want that that's business logic, right? I want that to be available to me. So I think the challenge of the SaaS business model is we packaged one way. We now have to learn how to unbundle these things and re-bundle in new ways and discover new business models, right? I mean if you look at it, te- what's

happening today with Microsoft 365 is a great example.

Right? We have this thing called Work IQ. In fact, like what we are realizing is, oh my god, like you know, if you look at In fact, there's a historical parallel to, right? We sold first Exchange and SharePoint and uh you know, before Teams, we had a thing called Link Server and what have you. And we thought, oh, that's all going to move to the cloud, but little did we realize that oh the number of people who will use servers in the cloud is 10x 100x, right? Because people were not buying servers, they were just buying a subscription.

The same thing is now happening with M365 because with Work IQ, we have exposed what is perhaps the most important database in a company that never got used as a database because it is only captive to our apps, right? It is all email operated on it, Teams operated on it, Word, Excel, PowerPoint, SharePoint. But now, like this is one of the coolest things I get to do with Work IQ. I go to a GitHub repo and I say, "Hey, I attended a bunch of design meetings last week related to this repo. Can you capture all that and tell me what changes I should make?" I mean, think about that. Right? It literally can go look at all those transcripts, come back with a plan to change a code base. Right? Previously, you could never have thought of using M365 for something like that. So, the value creation opportunity now in the agent world is in fact 10x more. But it does require us to have, for example, there's going to be usage around M365, right? Which is going

to be perhaps more than even the end users. And we have to even re-architect. Like in fact, like what I used to serve an inbox or a mailbox cannot be used to serve an agent.

Uh and so, that's sort of what we're doing.

>> I don't believe in like permanent business models for any of these domains, but in the near term, do you have a prediction between uh you know outcomes based pricing token based pricing enterprise bundles

>> Yeah the way I think about this is always we have like let's even take the per user pricing.

>> Mhm.

>> The per user pricing is really an artifact of someone creating a budget needing certainty.

Right because it's the most important thing like somebody wants a budget.

>> Mhm.

>> They need a per user and and per user is just a set of entitlements to usage.

Right that's kind of what it is and so the way is if the first bundling will be take some usage bundle it into per user stacks and you know then sell

subscriptions. So subscriptions I think are going to be there per user is going to be there. Then the next big thing will be consumption. So people will say

I want consumption. And it's also possible that people will say I don't even want to pay for any of the subscriptions or the consumptions outcome. But remember

most people love outcomes until they have an outcome. Because once you have an outcome it's like giving away

royalty. Right I mean like I've talked to customers who love you know outcome based pricing and I say I'm all in until they oh my god like what are you talking about you're sharing in my outcome no no no I want you to go back to per user pricing and I want you to consumption price. Right so I think that debate will go on

and all all the all of these business models have a particular time and a place versus one to rule them all. And if anything if you're a SAS vendor or you're a platform vendor having that flexibility and quite frankly we face this with GitHub. Right we just recently announced a per user pricing on GitHub. Because little you know

GitHub co-pilot was constructed at a per user level before we understood even the intensity of usage of agents. Right it was an interactive way for a developer to use code complete maybe task. It was not like oh I launched 10,000 you know, agents that are going on all day. Right? So, that is what the adjustment is about. So, now that we really want there will always be a per user, but there will have to be a consumption meter.

>> How do you think about the durability of SaaS more generally? One thing I've observed is in a lot of enterprises internally, there'll be teams that almost have agent euphoria. They're so excited about the explosion of things they can build that they're trying to rebuild a lot of applications or going to other SaaS vendors and saying we're not going to work with you anymore or we're considering an internal project. And it seems like in 6 to 9 months maybe

some of those people will come back and say actually we we can't rebuild everything. How do you think about what's durable in this world and what isn't?

>> I I think I think we have to go through one full budget cycle on this to really see the uh

uh the sort of the emergence of the equilibrium. Because at the end of the day

there's marginal cost to even generating the app. Right? So, in fact, there can be even a a simple way to say it like if you should always acquire something if the marginal cost of building and maintaining

uh something on your own is higher.

Uh right? That should be like it's a quantifiable right? A quantifiable thing. And the maintenance part is important. Right? Even like you got to remember like hey, you know, all the security stuff that now AI will find you better fix them too fast. Uh of course, there's a coding agent to help you with, but then that burns tokens. Right? So, whose responsibility is it? It's kind of like a a cycle that you've got to think through. And I think we have gone through the excitement that I can generate a lot of software. I think the next thing would be what software do I really want to generate? What software do I want to use from others? How do I compose these two into some agentic workflow that I have agency over? Right? Because I think there'll be very little tolerance for anybody who's inflexible uh at the vendor level.

Uh but at the same time, I think that anyone who has got that flexibility

shows up, delivers the value, will be back at again, right? We're selling software, but with just different business models, in fact.

>> Uh speaking about building software, um one of my favorite moments from I think a previous build, maybe 1 or 2 years ago, was they had a they they there was a section of you building your own software. I'm curious if you're building anything now.

>> Yeah, so I I think the you know, first of all, let's face it, right? Building software has made it possible for even the incompetence of a CEO of a company like ours, [laughter]

uh you can build. So, thank God. But that said, I I I do feel that you know, something like um GitHub Copilot to me, and especially the new sessions app or the new app, has just made it so much more possible for you to have agency over artifacts that you felt you couldn't touch before, right? So, to for me as a CEO even, to go to a codebase, uh to be able to learn about it, like I remember joining Microsoft long back, you know, first, and then you say, "Well, everybody had to go in and look at, you know, whatever, Cutler's Mallet or what have you, to learn how to do good C C++ code." Um

so, now that ability to be more full stack up and down is so good. But that doesn't mean every one of us should be doing the same thing. The question is, how do you then have the ability to inspect things, learn things, see things, um I think it's just so much more. And

so, to me, what I'm building a lot of is these long-running foundry agents, uh right? So, there's autopilots. So, the easiest thing is, to me, I think I just built one uh even last week, where the idea was, hey, can I have an agent that is continuously monitoring, essentially, my own chief of staff autopilot, right? We're going to have that obviously in Scout. That's what we showed, but it is so easy and trivial to build. I took it work IQ. I said take work IQ, go and build a foundry long-running agent, store all the memory in using Raven, right? Basically at my back end as a service. And lo and behold, it built it and not only built it, I could say publish to Teams and it published the damn thing to Teams. So, the ability to have you know some end-to-end project like this complete is just pretty miraculous.

>> How do you think that impacts the different types of engineering roles that exist in the future because right now I think there's you know a dozen different types of engineers that you can be from QA, front end, etc. You know, there's a big swath. I've heard some people argue that in four or five years we'll basically end up with four engineering roles. It'll be people who are managing agents. It'll be forward deployment engineers or FTEs. It'll be security engineers and then people working on large-scale infrastructure for a small number of services and then everything else just collapses into the agentic world.

>> Do you think that's a correct view of the world?

>> Yeah, I mean I think I think we'll have to experiment our way through it. But what you said is what there are some very at scale things. At LinkedIn, they did structurally change and it you know basically built up a new discipline called full stack builder, right? So, they went and said, "Hey, let's bring people from design and product management, front end engineering, all put them together, but also have an edge, right? It's not like the design person still doesn't have the design edge or the front end person doesn't have the front end edge, but you can give yourself bigger scope in role so that you're not confined to one role. And then equally infrastructure has become very critical, right? So, in other words, like I mean our R&D team, I mean one thing we've is even for the Excel team, for example, building the RLE in which a reward can be learned is actually one of the hardest sort of infrastructure problems. And so you kind of need even new talent, right? Distributed systems people even in what was considered an end-user app team because it's a different skill set. So yes, infrastructure science is the other one, obviously. So I think we'll see how these evolve, right? Where's the reals, I mean, always the world will have a bunch of specialists. You know, I think the generalist role is going to be the most exciting, right? Because the leverage of a generalist is where we're going to see the maximum returns, right? When when you said,

"Hey, I'm coding. I'm now a general like what I've basically translated knowledge work, right, which I did where I created a Word document or a spreadsheet or even and now I can build an app, right? It's in the same sentence, right? That idea that, "Oh wow, my generalist skills have gotten a higher leverage." I think is what we're going to see across the board.

>> Music to the ears of CEOs and VCs that are like a little dangerous and a lot of fun.

>> Golden age for idea people.

>> Idea people with a lot of agency. If you take that idea of personal agency and you just zoom it out to the organizational context, um

uh my partner Mike Vernal, who actually started his career at Microsoft, just wrote an essay where one of the big takeaways is it's an age where you can be much more ambitious and you need to be given the pace of the environment and how quickly actually users and companies are open to adopting new technologies. Um how do you think about I feel silly asking this of somebody running a, you know, trillion-dollar-plus company already, but how do you think about how Microsoft can be more ambitious now?

>> It's a great question. Um I think um I think the the thing in these type of transitions is to have a conceptual model of how work can change to go after outcomes that you could hardly imagine previously, right? In fact, Kevin Scott has this nice line,

right? Which is
um when you can make the impossible like
when you're making hard things easier
that's sort of one point of leverage
but true ambition is about making the
impossible possible.

So
now the thing that is missing a little
bit in all of our organizations is what
is that new conceptual model of what can
we build? What was impossible and what
can we build?

And I'll give you one example of this,
right? Which is I take great inspiration
from sort of the people who were
managing the Azure net network.
And they came to the This was from even
last year.

You know, we were scaling. You saw that
I I talked about sort of how we built in
the last 15 months more Azure capacity
than we built in the first 15 years. I
mean, it's crazy.

Right? It's pretty wild. And it's the
same team. So they saw that and they
said, "Bob, this just ain't going to
work if we don't reconceptualize our
work."

So they built Essentially, they said,
"Our job is not to do Azure networking.
Our job is to build the agentic system
does that does Azure networking." Right?
These are the folks managing the 500
plus fiber operators managing the van,
right? All over. And fiber operations
ultimately is a physical operation.
Things get cut. Things get you know,
have to be repaired. You know, we have
fancy words called devops and so on.
Basically, emails are coming in and you
got to go respond to them, take care of
it. So they built this agentic system.

They even have a character for it. It's called Miles and it sort of does all this stuff. Right? They started sort of screaming for more tokens and so on. And so they were saying, "Look, we don't need head count. We need tokens in order to be able to manage our operation."

That reconceptualization of what their work is, right? They They basically took their work and made it meta. Right? That meta work is now their new work.

>> Mhm.

>> Right? In the '80s, if somebody had come to us and said, "4 billion people are going to get up in the morning and start typing." My model would have been, "We need 4 billion typists." But we're not doing typing. We're doing knowledge work. So that to me, I think, is it, right? Which is whether it's Microsoft or whether it's any organization is to give ourselves permission to do new types of meta cognition, meta work using these new tools to change the outputs that matter, uh and then really make the impossible possible. So completing that dot or that that connective tissue across those, I think, is where a lot of the enterprise value will get created.

>> Can we talk about the data centers?

>> Yeah, please ask.

>> Oh, okay. Well,

we This leads nicely into the data center buildout. I always think just I'm just impressed with the sheer scale of the buildout from Microsoft but also everyone else that this is redefining what it means to be a hyperscaler. And I just feel like that that that is unprecedented scale on finances,

on the way you run the company, but also the communities that are that are impacted. Um

then just talk a bit more about what you're seeing on the ground. Like when you visit your

>> Yeah, I think there are two aspects of it. Obviously, the the buildout is extraordinary.

You know, nothing like this has happened and it's great to be a one of the participants in it. Uh but you brought up the other part, right? I think at this point it's clear that unless we as an industry uh

are very principled about ensuring that the benefits of all the stuff we're talking about are felt in real ways at the community level, right? Because this is not just a a campaign.

Right, it has to be real where people are saying, "Look, this is not changing the prices on energy for me." In fact, if anything, it's bringing down prices because long-term there's going to be a better grid. There's going to be more energy. Water consumption is in fact not sort of in fact water is being replenished, right? You got to really, you know, educate folks on truly what's happening the the close-loop systems we're building. We have to invest in the training, the jobs, the tax base. In fact, the least talked about stuff is the amount of jobs that get created during construction, after construction, what's the tax base that's there in the community. And and all this has to be real.

And and if that is the case, then we will have permission. If it is not, we

won't have permission. It's as simple as that, right? Which is we we I think we have to take it as an industry pretty seriously.

I think it's good for communities to be skeptical, ask the hard questions for us to do the hard work, earn that.

But at the end of the day, if this if we can really be the

I've always felt like in human history, if you use a lot of energy, but also create a lot of value for society, the story has been fantastic. If you don't do that, it's not been that great.

And this time around I'm a firm believer that ultimately, if you do have a token economy that drives productivity, that drives economic growth, that drives broad-spread, you know, participation, better health outcomes, then I think we will be in a great place. And that's at least what we all have to be focused on.

>> Yeah, it makes me think actually that with all these initiatives that you're doing, might be easier to see ROI in the communities first before in an enterprise.

>> Yeah.

>> [laughter]

>> I I mean, I think both sides. In fact, it comes back together. It has to be the people in the communities are going to be employed, are going to be participants

uh in the real economy, right? That's I think the question is like if we if the broad economy is doing well and the communities are doing well, the dots get connected. It's sort of the market forces are such that we will connect the

dots. And that I think is it. Like you ought to be able to see the evidence. You can't be about any one company. Uh but it has to be broad economic growth and broad you know, community permission.

>> Yeah. What's the biggest mistake you've made about currently or whatever you most updated your personal models on regarding societal impact of AI?

>> So, you're saying what's the the

>> What have you updated most on in terms of societal impact of AI?

>> I think the

um the

the most uh

critical thing is the first question we even started with, which is we need to tell

the story and make it real

that everybody has a real shot to

participate as a first-class participant

in this new economy. Right? That's kind

of I think in the next 12 months, 18

months, we need a way for people to say,

"Oh, wow, I get it." Right? There's

going to be tremendous capability,

tremendous amount of infrastructure,

but I can see what is going to happen

whether it's the benefits like health

outcomes or my ability to create a

startup or my ability to run my local

sort of uh store more efficiently, it's

just happening and I see that uh benefit

myself. Right? That to me

you know, earning that permission in a

path-dependent way, we can't wait. See,

the one thing you learn I've now learned

is I think the world is going to be way

skeptical of tech and tech companies

that say, "Trust us.

We've got it. The future is going to be glorious."

Uh you kind of have to deliver tangible benefits. Um

And but frankly, politicians winning elections

uh because they have advocated for that.

That will be at least my adjustment

because without it um

thinking that somehow because it's too

important this time around. It's too

much of the economy for it not to be the case.

>> So, one very simple framework I have for, you know, what are what is going to be the broad benefit of AI um

beyond the communities just working in technology uh our sort of wealth

creation is going to happen in a ton of

different companies, startups and large

companies. Then you have healthcare. Uh

you you had amazing demos today. There

are companies like Open Evidence. I

think that is happening. Um education

seems like another one that's an obvious

good where we haven't seen as much

impact as I'd expect. Do you have a

hypothesis on why that might be or if

it'll come?

>> Yeah, I mean I think this is where

again, how we think about education, how

you know, recently I met with uh the

founders of Alpha School and learned a

lot about what they were going and going

about it. It is fascinating to listen

uh to how to even rethink uh what is

education really look like because I

think it's actually very important. Uh

and I'm not saying anything

traditionally being done is less

important, right? I was even looking at

the

It's fascinating to see I forget to which Stanford class it was uh the the Asian guidelines for CS something. Uh because you still need people to learn. Uh like it is an interesting AI class that they were making sure people were learning how to apply softmax appropriately versus saying, "Hey, fix my training run." Uh so, I think learning concepts is important. It's going to be a critical. But, the way we create the incentives, what are the credentials, how we value those credentials, what is the employment opportunity for those credentials? So, I think that there is a complete change that has to happen uh given the way to get to information, way to educate yourself, way to continuously keep yourself updated has changed so much. So, I think interestingly enough maybe the next big startup and success story could be someone who builds a new university um or a new um pedagogy even of how to get someone to go through a curriculum and find economic opportunity uh that's highly valuable.

>> Well, that has felt uh perhaps impossible for a long time, but it's a great note to end on and something that might be possible. Yeah. Thank you, Sayan.

>> Thank you so much. Thank you. I appreciate it. Thank you all.

>> [music]

[music]