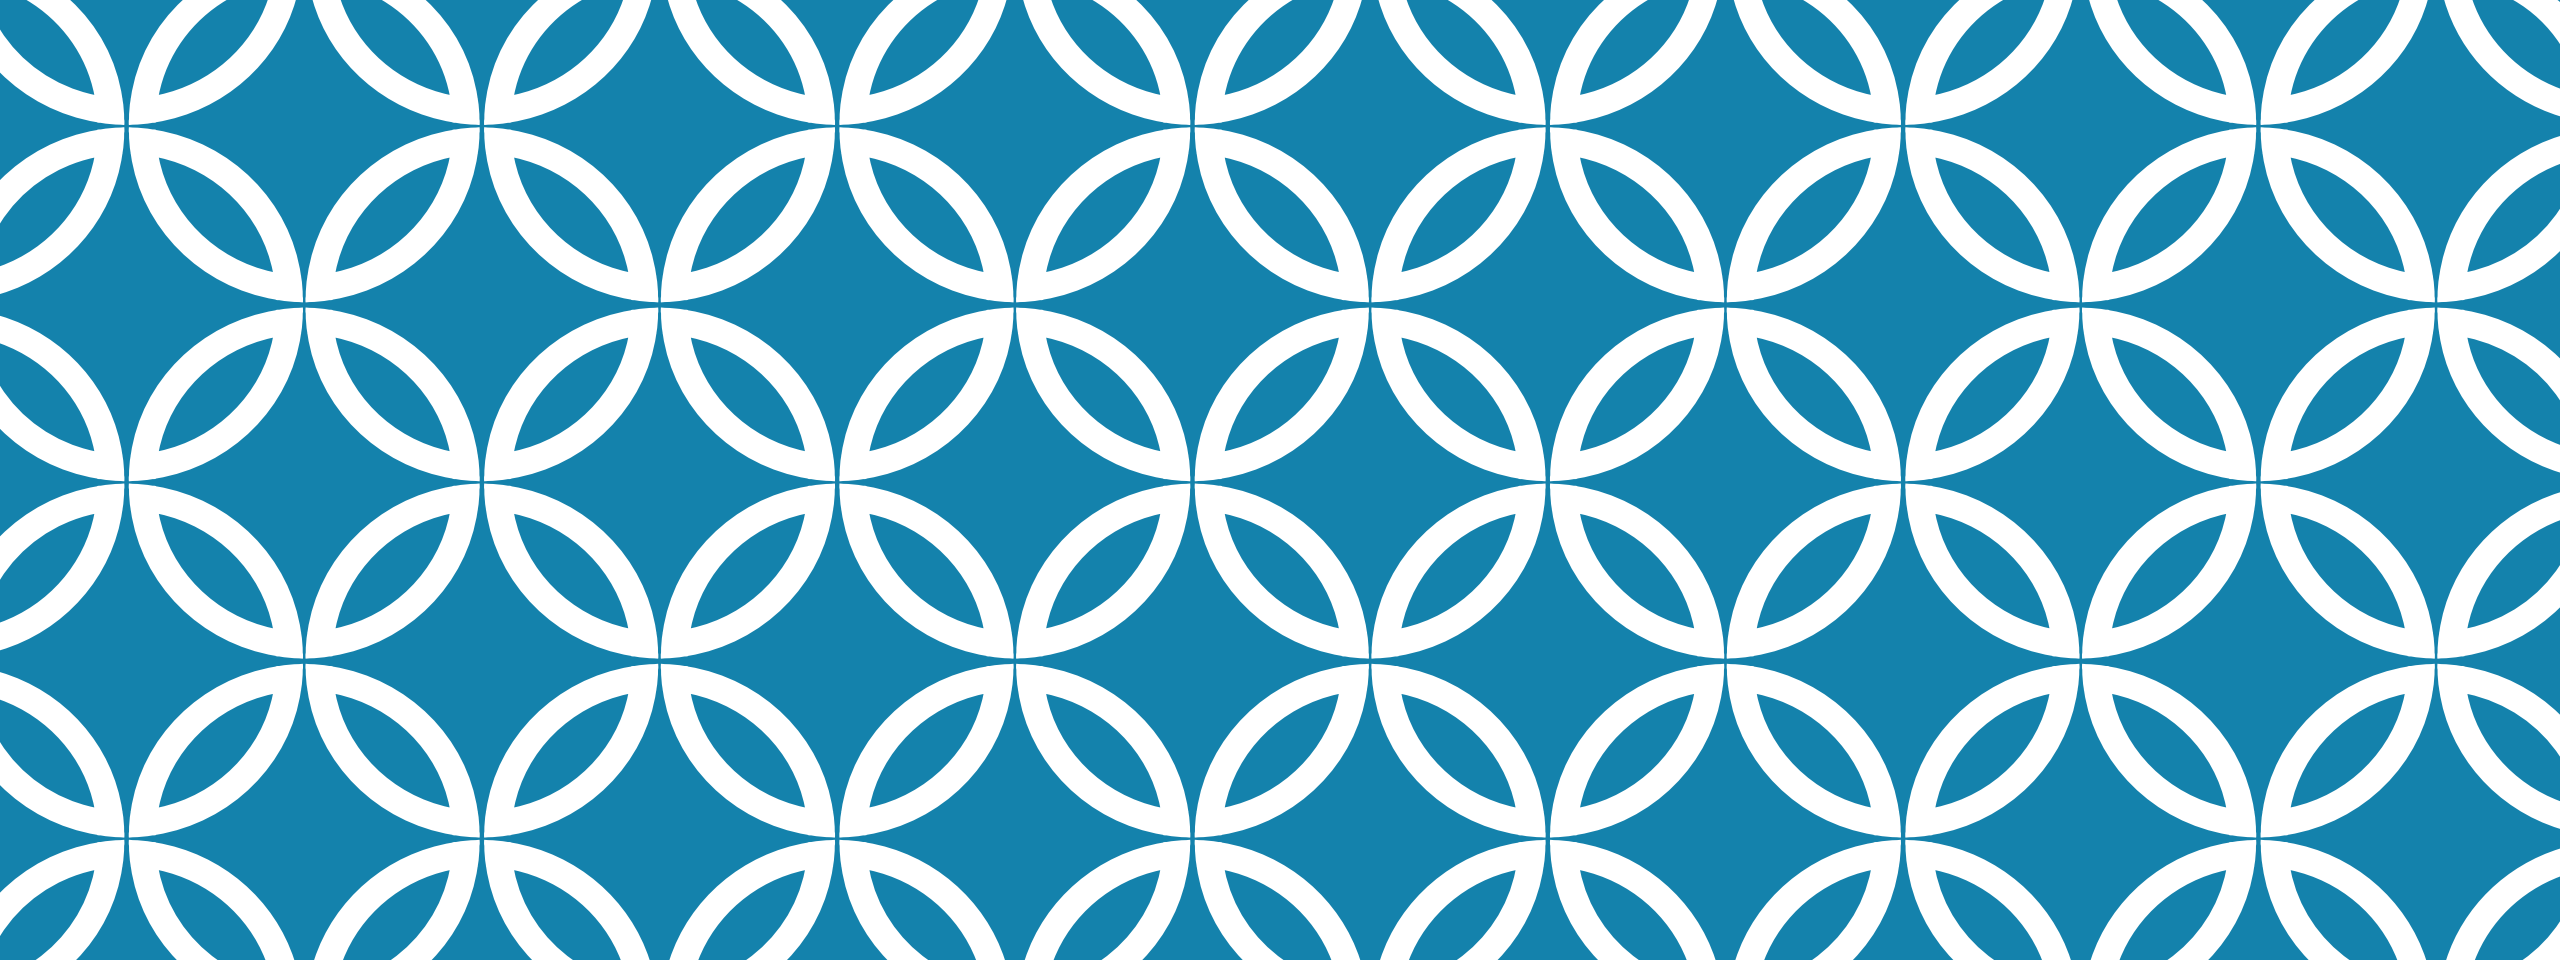

CS5412/LECTURE 28 THE FUTURE CLOUD

Ken Birman

CS5412 Fall 2022

OUR LAST LECTURE!

Today, we'll peer into a crystal ball (using lots of web search)

Look at business articles about the future of the cloud, drawing on white papers published in various sources (listed on the syllabus page and in the notes).

CORE QUESTION

We have better technology!

I dominate the market!

Who wins? And why?

THE EVOLUTION OF THE CLOUD AS A MARKET SHAPES THE CLOUD AS A TECHNOLOGY BASE

To understand where the cloud will go in five or ten years, we should try to understand how cloud use and demand will change (or not change).

Investments in new technology concepts and development align with cloud use patterns that scale poorly, and opportunities to increase efficiency.

In contrast, things we do well now and will continue to do more of in the future evolve towards greater cost-effectiveness, not new technology.

REMINDER: CURVE SHAPES

We've previously discussed how rapidly-expanding markets are often close to saturating just as they seem to be growing exponentially quickly.

We also discussed the "two peak" adoption curve from Crossing the Chasm

Keep these curves in mind because as we look to the future, both are relevant!

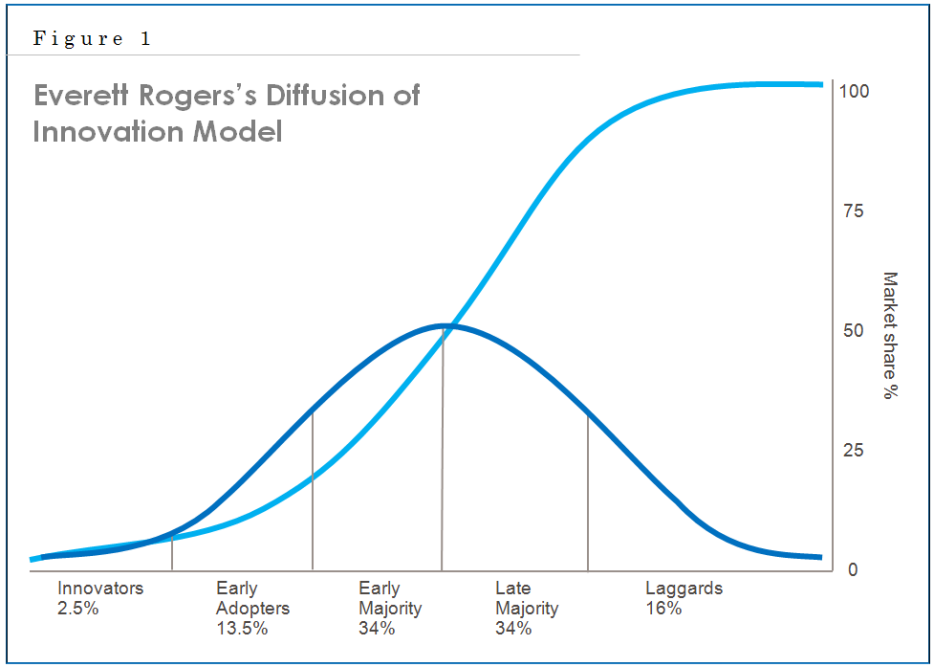

REMINDER: CURVE SHAPES

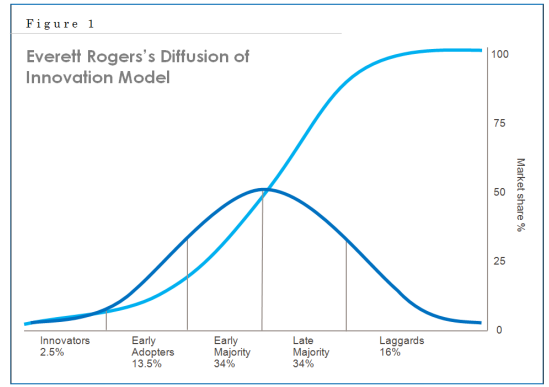

We've previously discussed how rapidly-expanding markets are often close to saturating just as they seem to be growing exponentially quickly.

We also discussed the "two peak" adoption curve from Crossing the Chasm

Keep these curves in mind because as we look to the future, both are relevant!

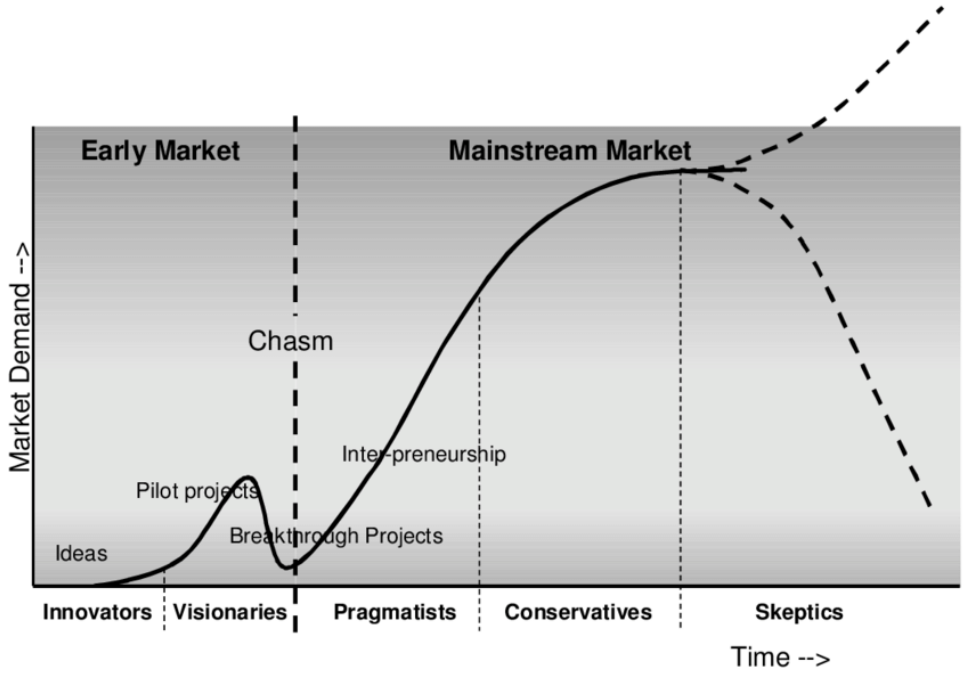

CAN WE REALLY PREDICT FUTURE GROWTH?

One idea is to split the question into two sub-questions

For things the cloud does now, what trends are revealed by market research focused on expenditures?

- Money is our metric of growth.
- This assumes that that what is happening now would more or less continue to happen for a while longer – "probably, no big shakeups".

Things the cloud isn't doing now are different: here we are speculating about new markets, perhaps based on unsatisfied demand.

HOW DO COMPANIES DO PREDICTIONS?

They basically do data mining, or pay someone to do it.

The usual process is to pose a bunch of questions (assumptions) about the market they are excited about. Lots of questions.

Then through interviews with leaders at the top potential clients, they try to learn about the broader trends without limiting themselves to only their favorite clients or their investors. This gives some sense of the picture.

EXAMPLE: PUBLIC REPORTING BY A BUSINESS ANALYSIS FIRM

The report is public because they are hoping to drum up business from readers.

They see a shift from PaaS to IaaS to SaaS.

PaaS: Cloud as a "platform" for tasks like direct sales or company web sites.

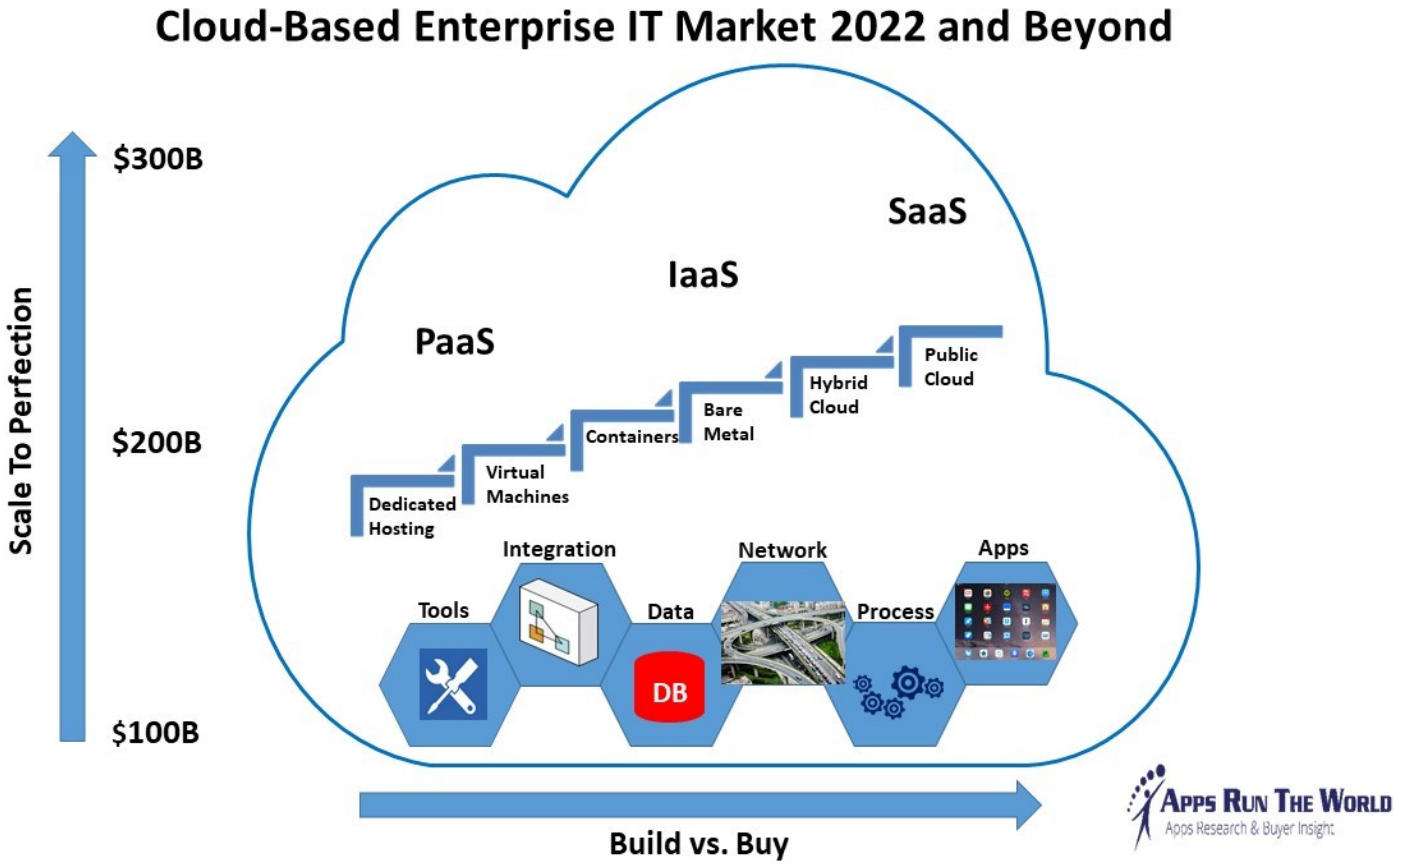

https://www.appsruntheworld.com/first-look-at-2017-2022-paas-iaas-saas-markets-as-14-top-cloud-computing-billionaires-expand-across-enterprise-it/

WAIT, WHAT'S WITH THE AAS THING?

Most cloud courses talk about XaaS on day one. In CS5412, we didn't really get to it at all! There are incredibly many "as a service" concepts tied to the cloud.

Examples: IaaS, SaaS, PaaS, plus some you need to write out in words (for example, SaaS means "system as a service" but "storage as a service" is a big deal too).

In general, "as a service" means "Available for rent on demand. A vendor offers this technology or capability to a large (usually public) clientel."

IF AAS WERE USED IN OTHER SETTINGS…

Uber/Lyft do “transportation as a service.”

AirBnB does “accommodations as a service.”

Cornell does “knowledge as a service.”

BACK TO THAT BUSINESS ANALYSIS…

This firm sees a shift from PaaS to IaaS to SaaS.

PaaS: Cloud as a "platform"

IaaS: "Rent bare metal", use it to create hybrid solutions.

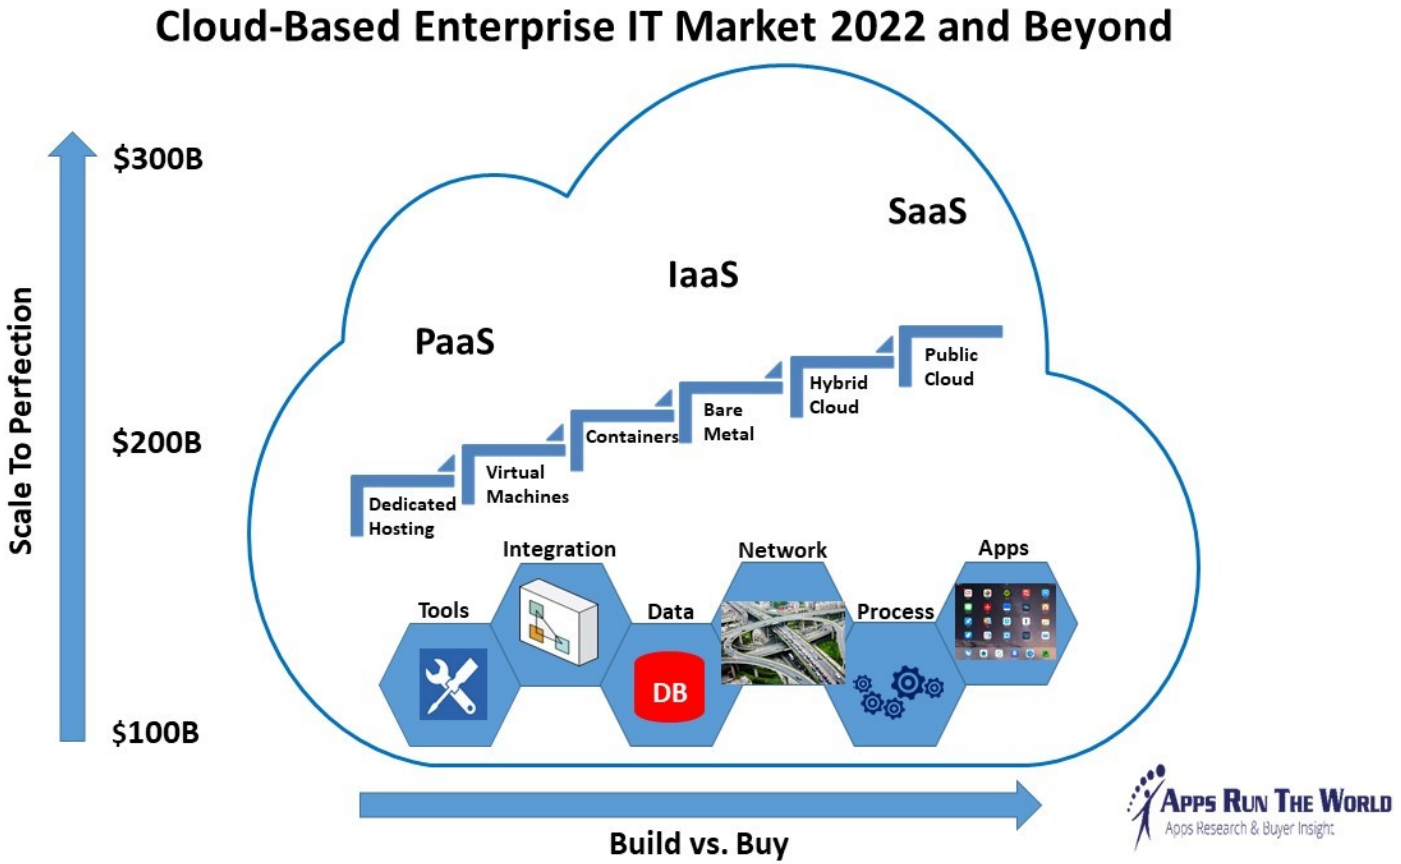

https://www.appsruntheworld.com/first-look-at-2017-2022-paas-iaas-saas-markets-as-14-top-cloud-computing-billionaires-expand-across-enterprise-it/

WHAT DO BUSINESS ANALYSTS SAY?

This firm sees a shift from PaaS to IaaS to SaaS.

PaaS: Cloud as a "platform"

IaaS: "Rent bare metal"

SaaS: Cloud hosts "solutions" like Office 365, Azure IoT

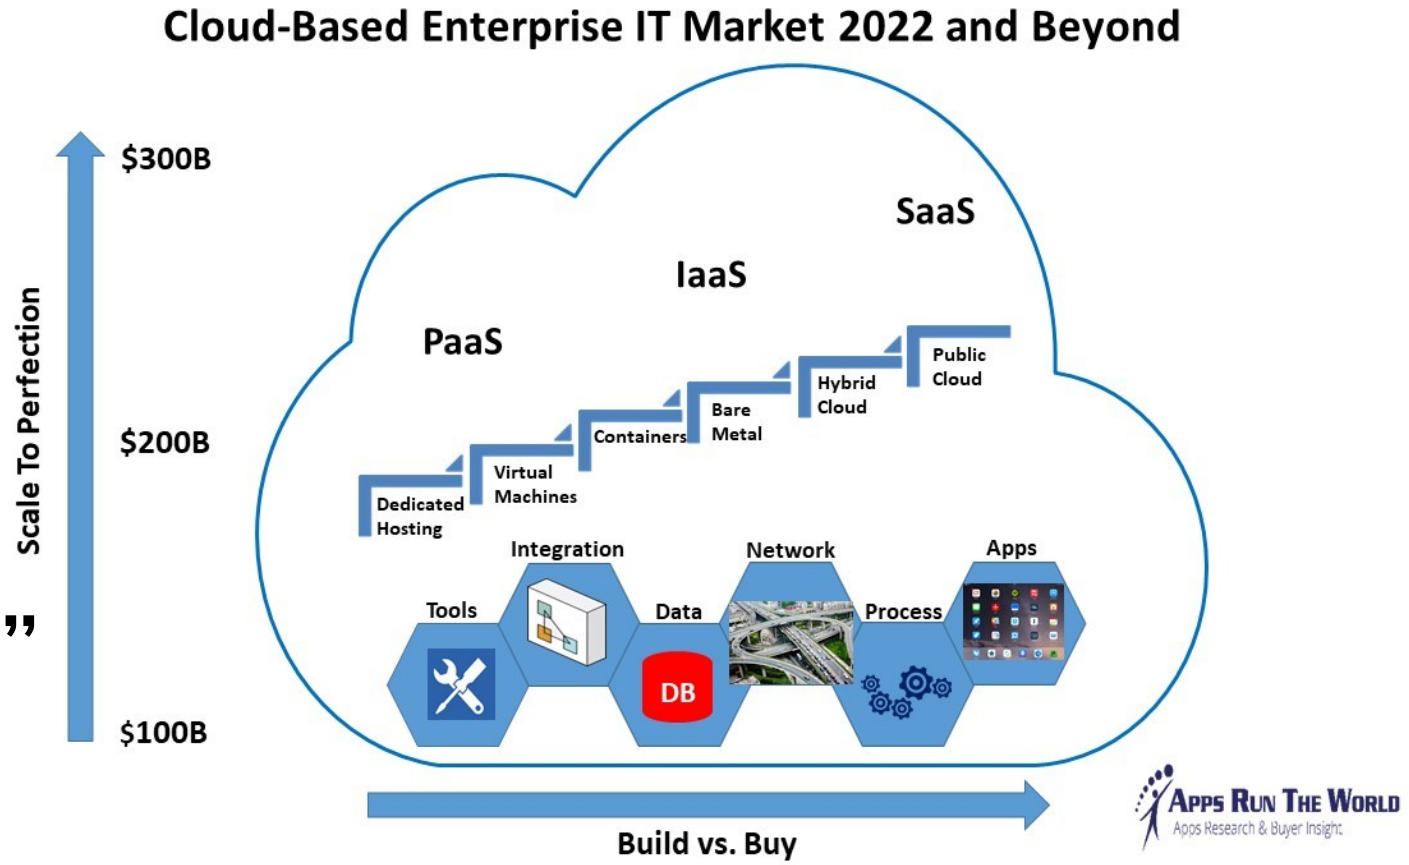

https://www.appsruntheworld.com/first-look-at-2017-2022-paas-iaas-saas-markets-as-14-top-cloud-computing-billionaires-expand-across-enterprise-it/

WHAT DO BUSINESS ANALYSTS SAY?

This firm sees a shift from PaaS to IaaS to SaaS.

PaaS: Cloud as a "platform"

IaaS: "Rent bare metal"

SaaS: Cloud hosts "solutions" (also known as "systems")

As we transition towards SaaS the cloud hosts Apps.

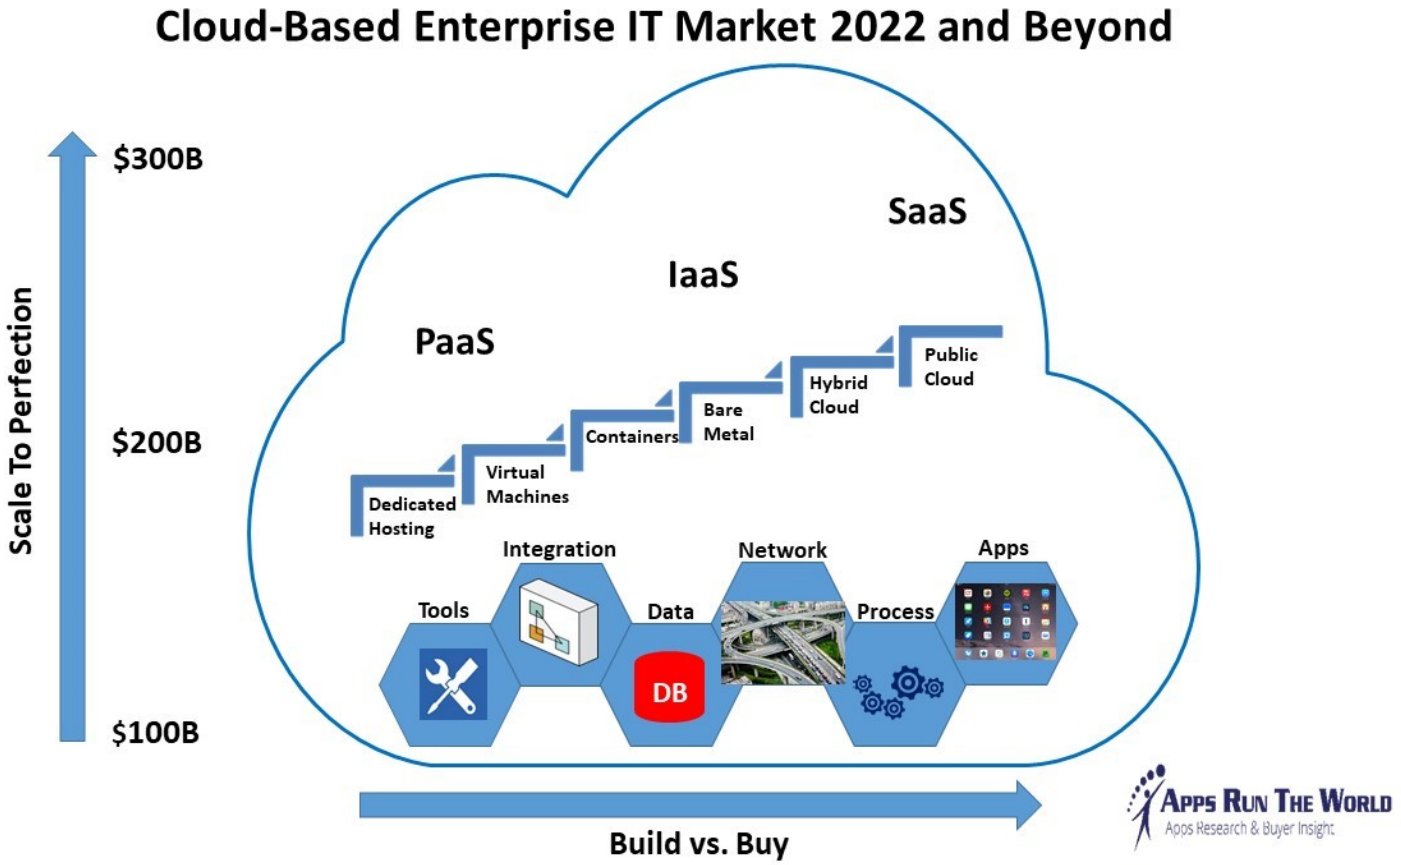

https://www.appsruntheworld.com/first-look-at-2017-2022-paas-iaas-saas-markets-as-14-top-cloud-computing-billionaires-expand-across-enterprise-it/

THEY TALK ABOUT AN "APP MODEL". BUT WHAT MAKES AN "APP"?

The term is a reference to our iPhones and iPads, or Androids: solutions to end-user problems that heavily reuse vendor-provided software tools and infrastructure.

Think about Microsoft Farmbeats: the 2019 version was "home built" and has an IaaS feel. But within a year or two, Farmbeats++ could easily shift towards being more like a collection of Digital Agriculture Apps that share a common Azure IoT Infrastructure. This author is predicting such a shift.

HOW DOES THIS IMPACT OUR TOPIC?

In CS5412 we mostly looked at components and how one uses them.

We also dug deeper to learn how they work, and how we could build versions on our own that would work differently and bring new functionality.

An App perspective suggests that "developers will favor standard components."

What technology gap does this expose?

… CASCADE PREMISE

As researchers, we asked ourselves this exact question

We concluded that the world needs a way to run ML solutions with better latency (lower delay), and high throughput, and high platform utilization.

… and also with strong consistency guarantees.

… and also back-compatible with existing ways of creating ML code.

SHOULD THE LACK OF PRIOR SOLUTIONS WITH THESE PROPERTIES WORRY US?

Yes and no!

No, as researchers. Cornell and other top schools exist to demonstrate the feasibility of options many people thought were inaccessible!

But yes, if we think about market uptake for this kind of work– about "impact". To impact the broad community, Cascade would need to be a product adopted by Microsoft, Amazon, Google, etc.

REVISED QUESTION TO ASK

Do those guys “need” what we researchers can offer?

… badly enough to adopt third-party solutions

… created by grad students, post-docs and a professor?

HOW IS THE MARKET EVOLVING?

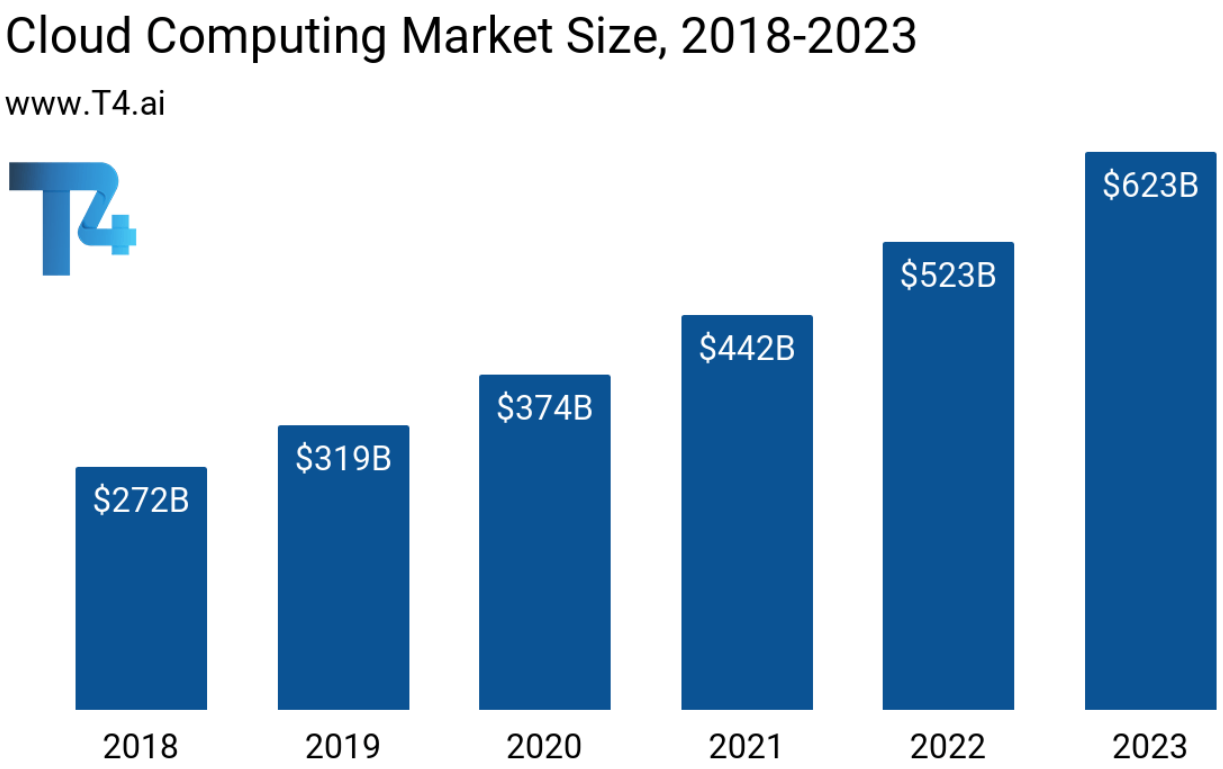

HOW IS THE MARKET EVOLVING?

In this analysis, from 2017, Microsoft broke through and for the first time ever Amazon was less dominant than Microsoft.

The reason is that this chart focuses on just PaaS, IaaS, and SaaS, plus looks at the subset of cloud revenue expected to "recur" annually

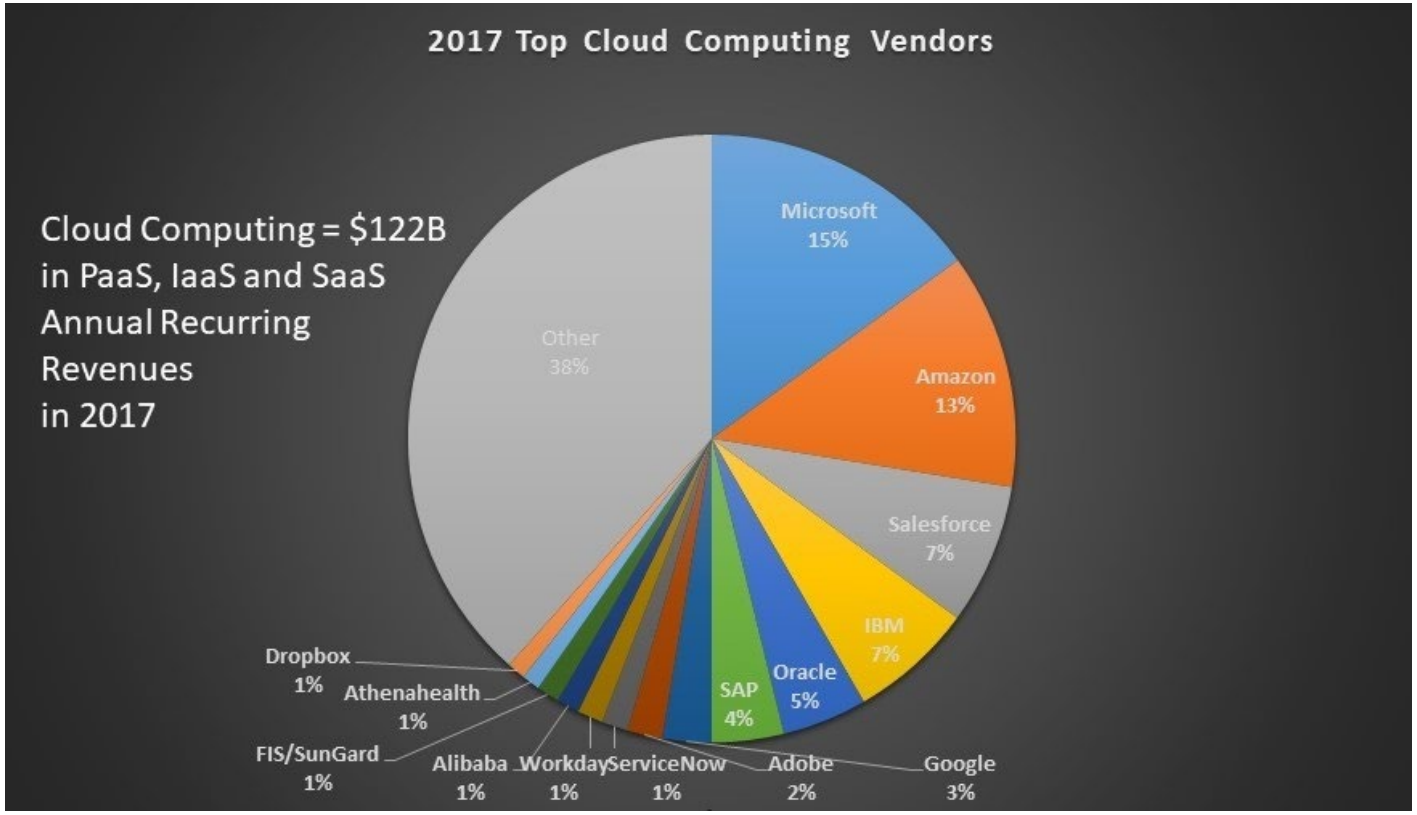

https://www.appsruntheworld.com/first-look-at-2017-2022-paas-iaas-saas-markets-as-14-top-cloud-computing-billionaires-expand-across-enterprise-it/

DIVING IN…

The article actually breaks down the revenue they used to build the chart, which is somewhat unusual for articles of this kind (many pull charts from thin air, and this is why I liked this one).

What it reveals is that the Microsoft acquisition of LinkedIn, combined with very strong growth for Office 365, has fueled a very strong growth in Microsoft market share, which was quite a bit smaller as recently as 2015

MARKET SHARE TREND ANALYSIS

This chart from IDC shows that there is a noticeable shift from traditional owned data centers towards public clouds (AWS, Azure)

With Office 365, one explanation is that companies are shifting their in-house business tasks into cloud SaaS

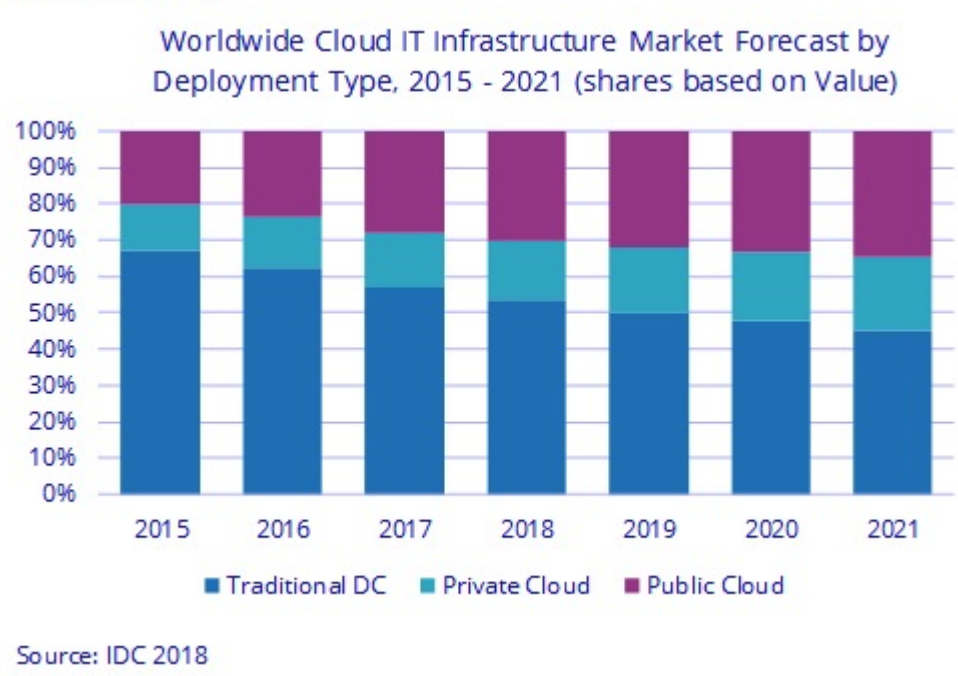

SLICED IN DIFFERENT WAYS, YOU CAN REACH VERY DIFFERENT CONCLUSIONS

This Q3 2018 data from awsinsider.com (a company with an obvious reason to favor AWS) shows AWS strongly ahead of Microsoft and every other vendor.

But notice that it doesn't include SaaS, and instead uses a category called "hosted private cloud". What is that?

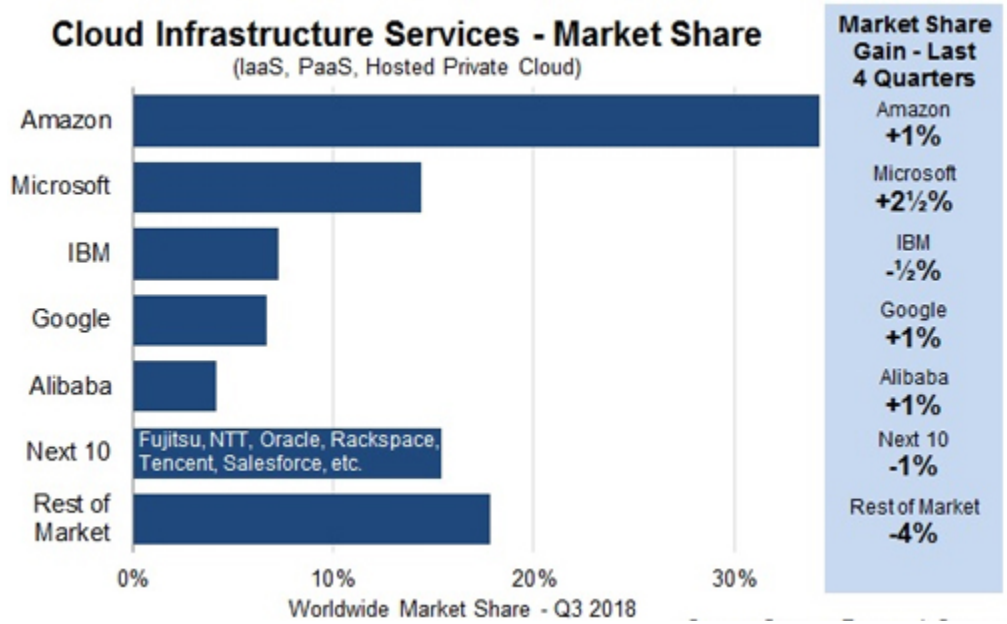

MOST "ENTERPRISES" ARE EXPLORING SHIFTS TOWARDS THE CLOUD

This is a market research report from zdnet.com

It shows that Azure and AWS have similar use patterns, but with more experimenting or planning to adopt Azure than AWS

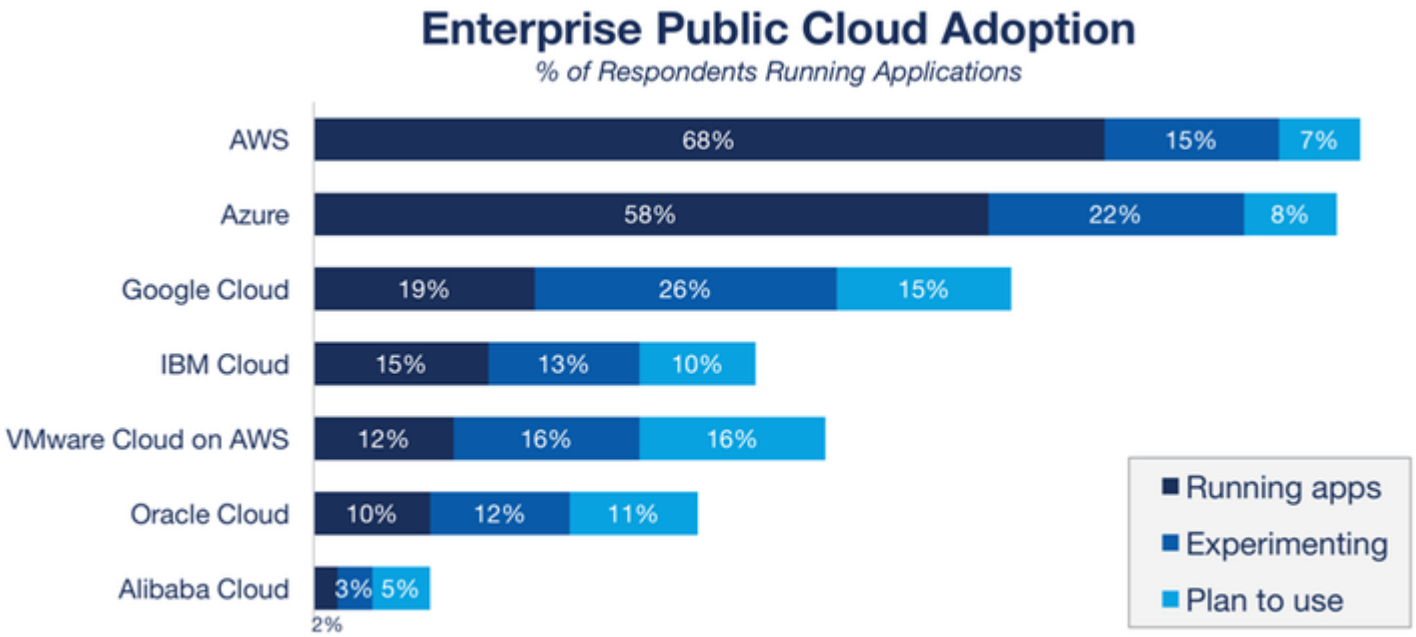

Source: RightScale 2018 State of the Cloud Report

AWS USE IS MORE ORIENTED TOWARDS IAAS

This is a market research report from zdnet.com

It shows that Azure and AWS have similar use patterns, but with more experimenting or planning to adopt Azure than AWS.

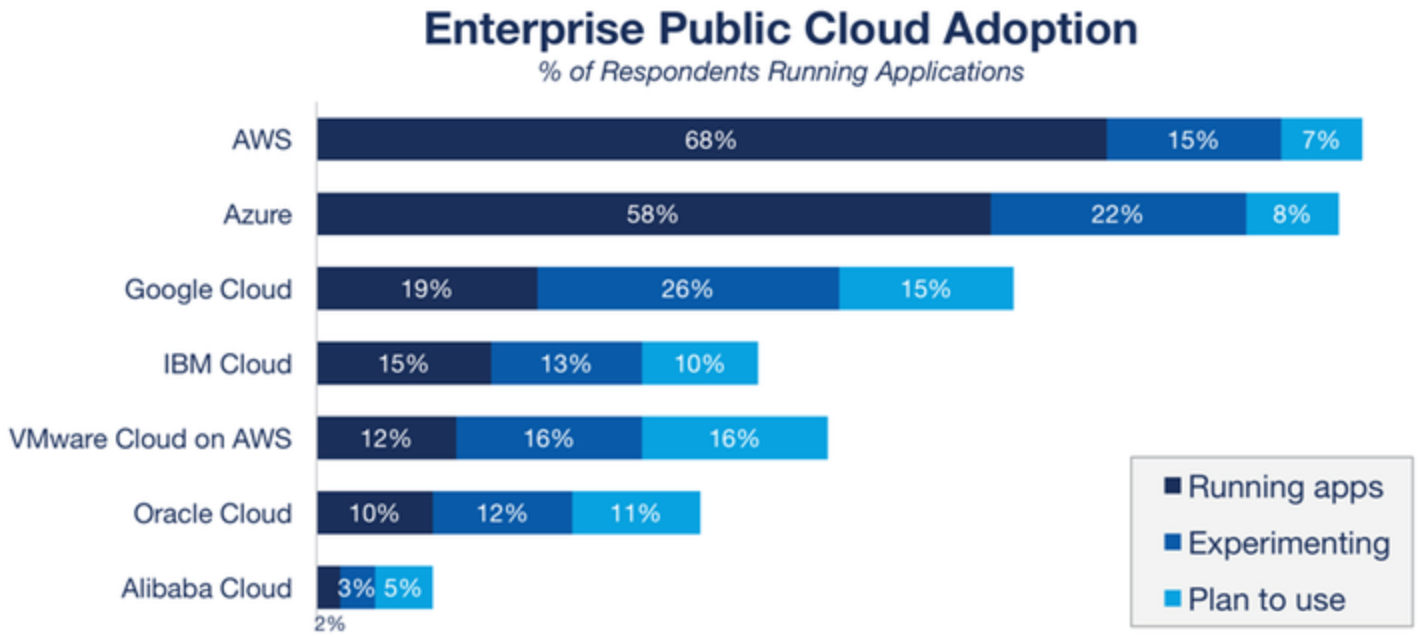

Source: RightScale 2018 State of the Cloud Report

MOST "ENTERPRISES" ARE EXPLORING SHIFTS TOWARDS THE CLOUD

If we look just at VMs we are focusing on IaaS rather than SaaS

Here AWS is the clear winner. And more companies are experimenting with AWS too.

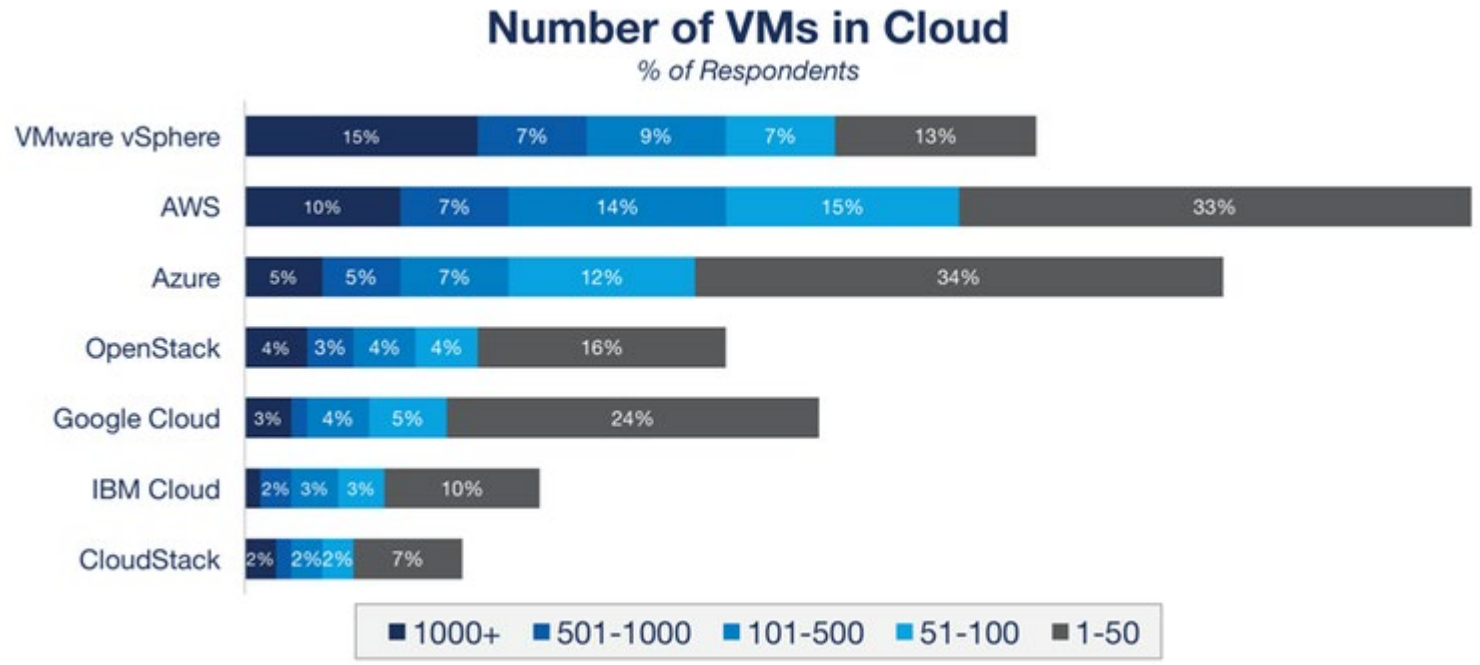

Source: RightScale 2018 State of the Cloud Report

AWS GOVERNMENT CLOUD

Hosted Private Cloud means "using AWS infrastructure but not a shared public cloud." The biggest such example is "government cloud"

This is a line of business in which Amazon builds and operates clouds specifically for the US government and military, securely.

That particular line of revenue greatly increases the AWS revenue for cloud as a whole, but in a distinct category.

WHAT ABOUT IOT AS A MARKET?

This chart from "statistica" looks at how many smart devices are out there world-wide

It includes everything from routers and smart TVs to home thermostats that do smart sector control and heat different rooms in different patterns.

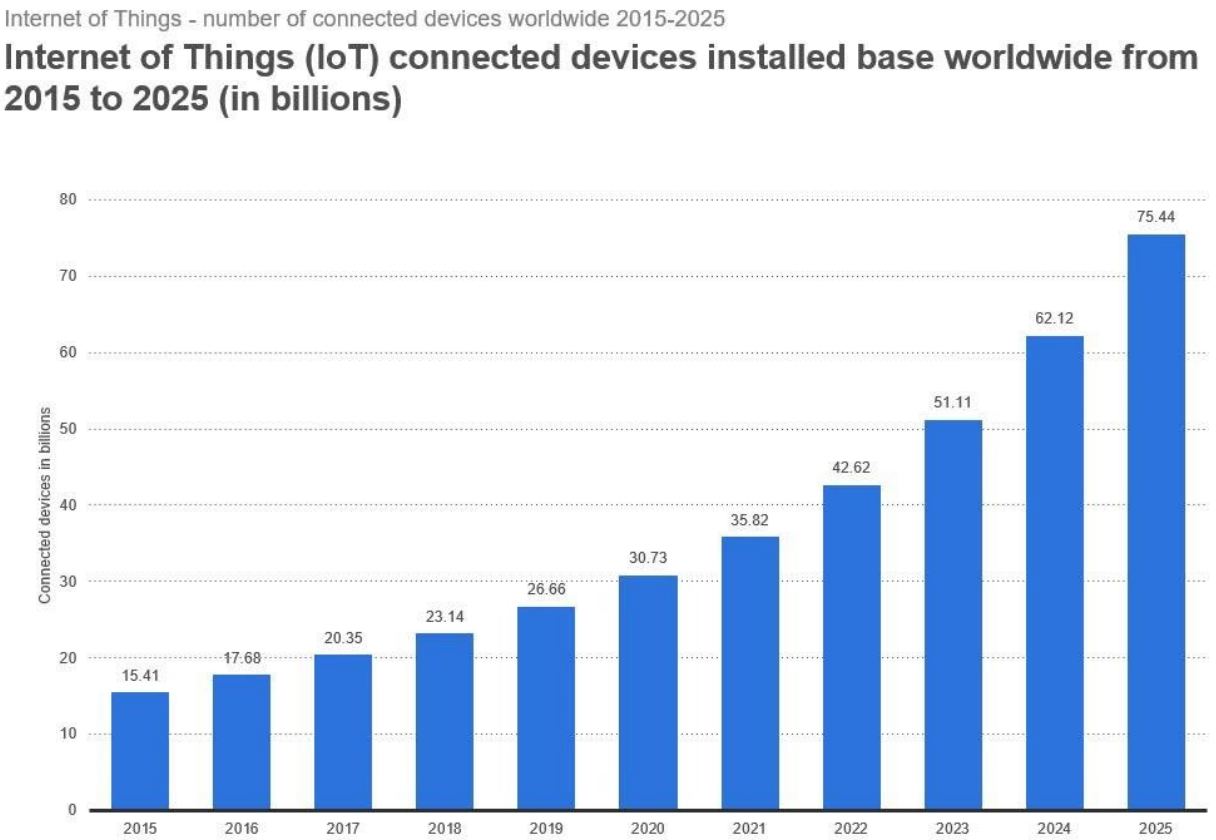

MULTI-USE EDGE CLOUD (MEC)

(Also called multi-access edge cloud and mobile edge cloud. IBM like the term “cloudlet” but somehow it never took off)

An emerging term for this market area. A MEC is a cluster of computers running on premise (“in the closet”) or in a 5G cellular control station near the antennas. The MEC offers a mix of cloud properties and “bare iron” properties, with very low latency.

Many things could run on a MEC that are hard to run on the main cloud…

MEC INFRASTRUCTURE MARKET PROJECTIONS

Exhibit 62: Chart on US - Market size and forecast 2021-2026 ($ million)

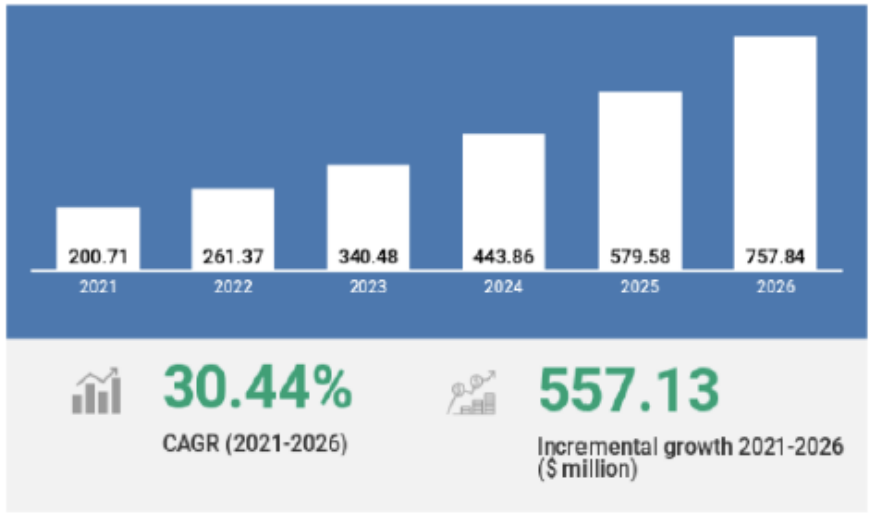

Exhibit 63: Data Table on US - Market size and forecast 2021-2026 ($ million)

Year	2021	2022	2023	2024	2025	2026
Market size ($ million)	200.71	261.37	340.48	443.86	579.58	757.84

Technavio: Mobile Edge Computing Market by Component and Geography - Forecast and Analysis 2022-2026

IOT MARKET: SMART GRID

If the power grid of the future will be smart, what will make it smart?

Part of the answer is "it will need to be, to handle distributed energy resources."

Then they point to various specifics that are mostly IoT and cloud, like ML/AI, mobility, cloud computing, data analytics, blockchain!

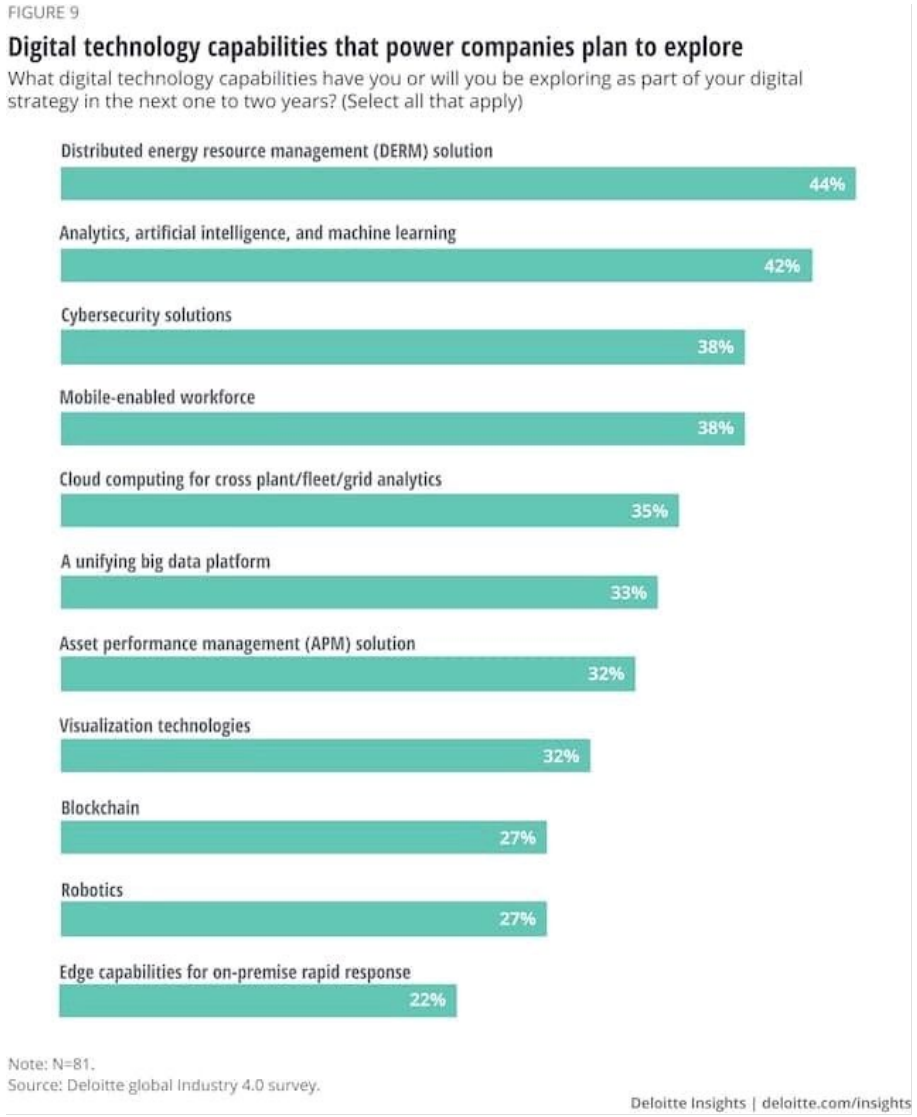

EVERY AREA HAS SIMILAR STORIES!

Smart grid has become somewhat mature and the numbers are more trustworthy

If we focused on augmented reality or smart farming, the projections would be less solid.

The edge has a lot of such stories, hence we can bet on the IoT/MEC concept even if we are uncertain about which sectors will be big winners

IoT DOLLAR VALUE ESTIMATES (COMPOSITE)

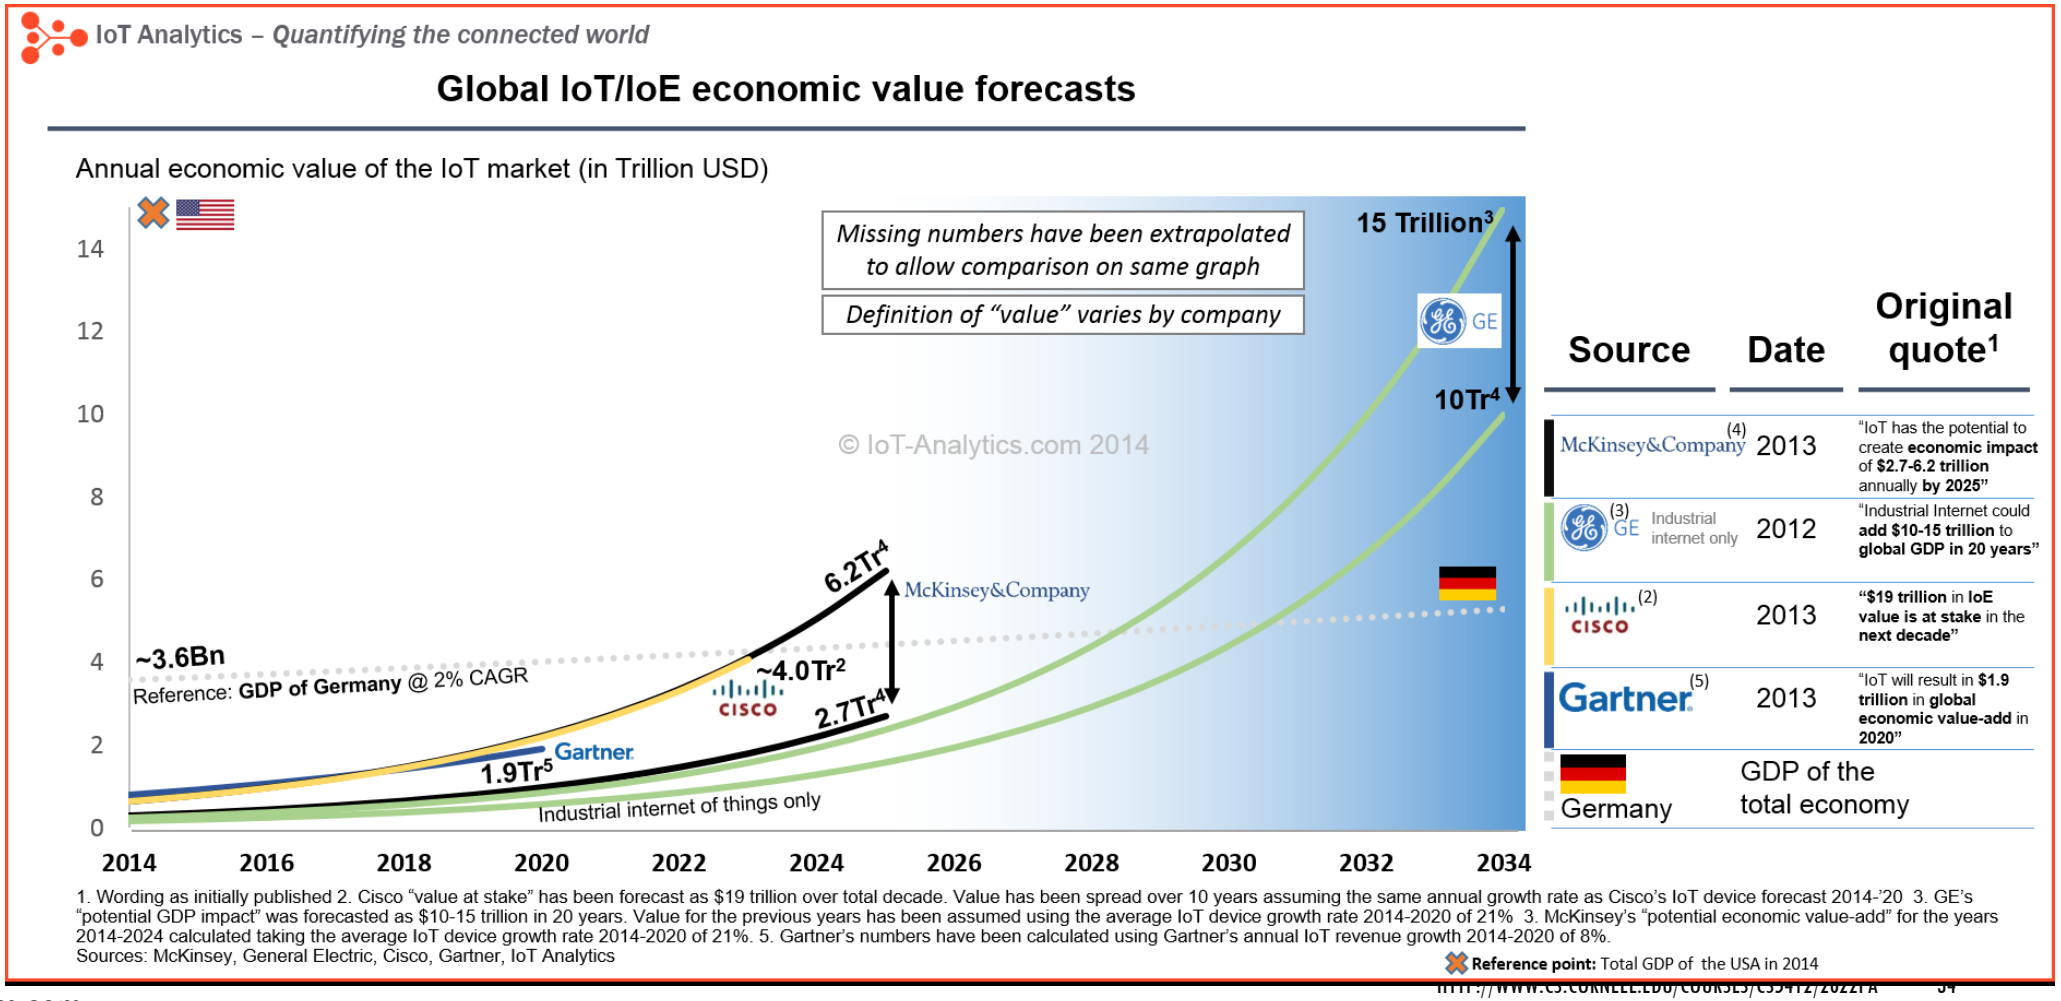

Source: IoT-Analytics.com

CLOUD SHARE OF IoT MARKET?

This estimate, by Machina Research Inc, suggests that the device dollar value will soon be much smaller than the platform and app value

So this supports the view that Apps could become a very big market for the cloud fairly quickly

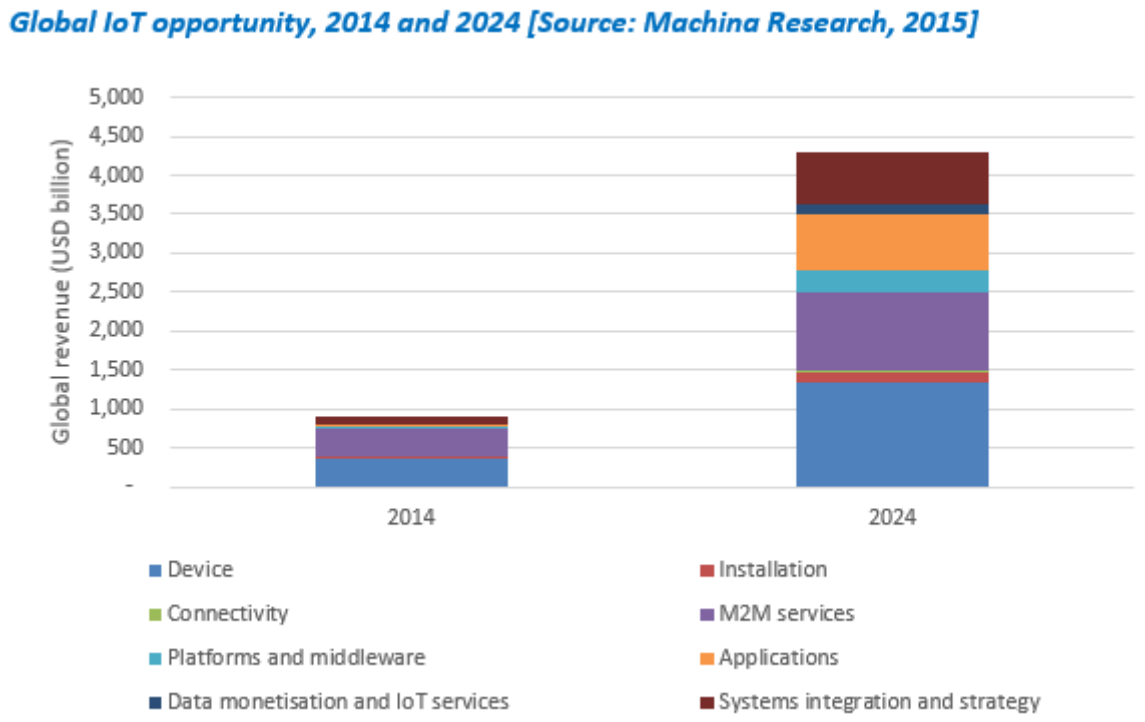

WHICH INDUSTRIES ARE DOMINANT?

Again, by Statistica.

Strong growth across the board but especially in “non” consumer electronic categories, like automation of retail stores and smart energy systems.

Smart agriculture isn’t included here.

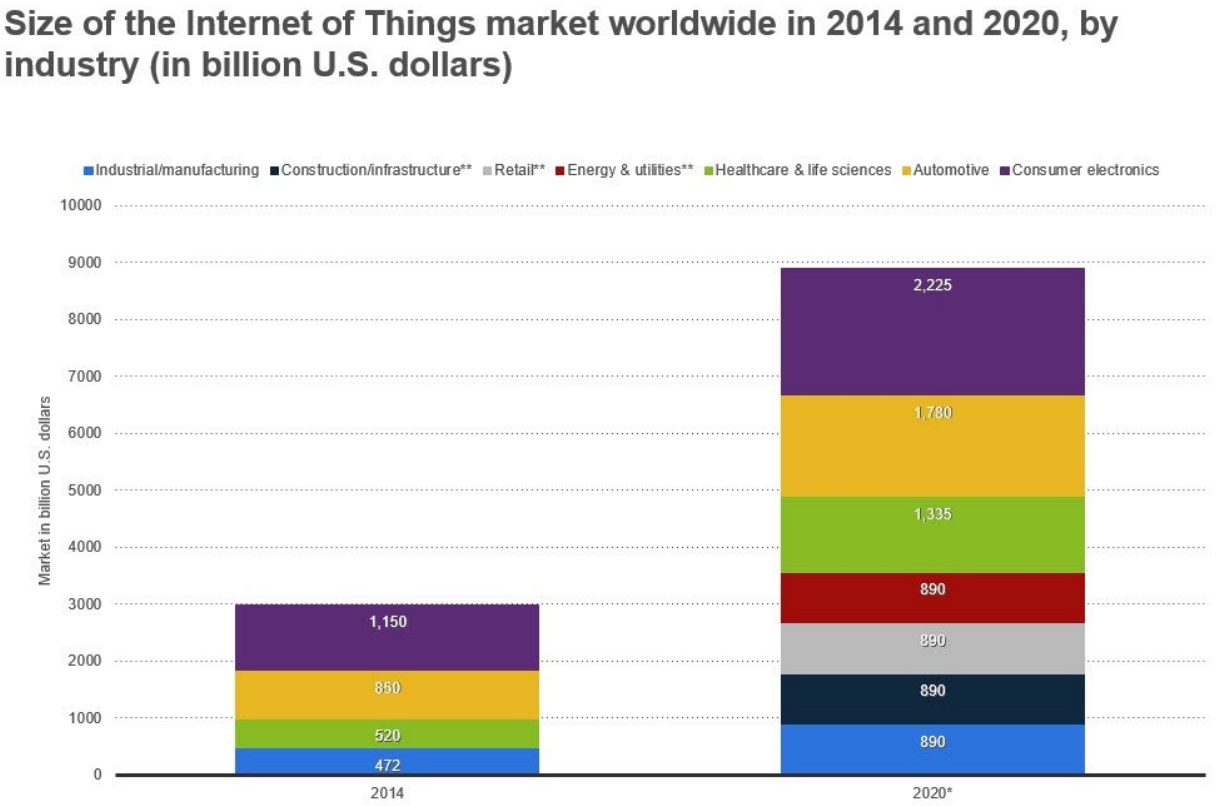

WHAT CAN WE CONCLUDE?

This data generally backs the view that the cloud will become a major player in markets related to IoT during the next five to ten years.

It also supports the view that integrated solutions (SaaS) are growing more rapidly than cloud as a "bare metal hosting VMs" but that there is still huge growth ahead in the IaaS sector.

And it suggests that the main *uses* of IoT are in managed enterprises, not random situations where sensors are just dumped into the environment.

SO HOW CAN WE LEVERAGE THIS INSIGHT?

Back to our focus on actual cloud infrastructure.

It seems to be telling us that the big need isn't going to center on cloud scalability for things the cloud is already doing perfectly well.

- On the other hand the sheer scale of the use causes it to be a very high value market, where products that improve efficiency could find uptake.
- So on this more mature side of the cloud, expect the emphasis to be on sqeezing more work out of the infrastructure using less money / energy.

MEANWHILE…

On the IoT side, we've spent a semester seeing that the technology base today is just now taking baby steps, but that the opportunity is real.

These predictions suggest that the market may look more like an App market (like Office 365) than like an IaaS market (VMs and hybrid cloud)

So companies that build Apps, or enable people to build Apps, will find strong demand for their solutions (if the solutions are good ones).

BIGGEST OPEN QUESTION?

To me, it centers on customized hand-built μ-services, for individual apps (in which case the need for IDEs and other developer tools is the biggest need)

… versus some kind of general purpose "ML as a service" solution, where a precreated vendor-supplied service somehow magically covers all the needed intelligent behaviors.

- Such a service could definitely leverage hardware accelerators, a big win
- But that same hardware could also be made accessible to custom μ-services through a suitable library, so this isn't a decisive advantage

HOW DOES IT WORK TODAY?

You need to build your μ-service (and it probably is a group of processes, so you need to manage membership, state, handle failures…). Derecho can help on those tasks, if you don't mind coding in C++.

You'll need to register it with Azure or AWS (using JSON files) and then build triggered functions that can talk to it (more JSON files). Some steps are not very intuitive and will require "research" (hours wasted!)

You need to improvise your own debugging and performance analysis tools, and are on your own for long-term "life cycle" aspects.

MY BET?

People like the cloud and want to preserve and expand the use of their existing technologies and solutions

But they also need the low delay of MEC hosting

So they will want a "mini-cloud" that looks like AWS or Azure, and has PaaS service options aimed at low latency, like Cascade. Check back in a few years to see if my bet is right! But to make this work….

ASPECTS A PROFESSIONAL SOLUTION WOULD NEED

Clearly, automated help creating those configuration and "task control" JSON files, integrated with AWS Cloud9 or Azure Visual Studio/VSCode.

For the APIs offered to the functions, automation of the API "declarations" and the logic to import them and call them from the function side.

Careful tuning (by the vendor) of the resulting paths. We want our functions running within 1ms or less of an event, and won't have time to launch the container at the last moment, or bind to the service. So those have to be done ahead of time, anticipating the need.

MORE PROFESSIONAL TOOLS WE WILL WANT

Launch service

Authenticate if needed

Register micro/service to accept RPCs

There should be an easy way to create functions able to call the service, using those RPC APIs

We need an efficient upload path for image objects

There will need to be tools for garbage collection (and tools to track space use)

… and tools for managing the collection of configuration parameter files and settings for an entire application

…. and lifecycle tools, for pushing patches and configuration changes in a clean way.

PLUS, SOME LARGER CHALLENGES

Code debugging support for issues missed in development and then arising at runtime

Performance monitoring, hotspot visualization and performance optimization (or even, performance debugging) tools

Ways to enable a trusted micro-service to make use of hardware accelerators like RDMA or FGPA (even if the end user might not be trusted to safely to so)

- Many accelerators save money and improve performance but are just not suitable for direct access by hordes of developers with limited skill sets.
- Some could destabilize the data center or crash nodes, and some might have security vulnerabilities.

CAN IT BE DONE?

I'm doomed! My boss wants me to lead the μ-service development group!

Actually, no problem!

This is not such a terribly long or frightening list.

The big vendors, like Microsoft and Amazon, can definitely build this sort of technology, and they will do an incredibly good job on it, too.

- And they will, if the market is as promising as it looks.
- Flow of revenue leads to investment in profitable businesses.

Amazing! My boss wants me to lead the μ-service development group!

CONCLUSIONS?

The cloud infrastructure is an expanding opportunity area!

- Using it to create new Apps will fuel many small and big companies
- Leveraging the power of the cloud to learn from examples (big data) is an incredible enabler to do things that were impossible with stand-alone code.

Cloud IoT is more of a speculative bet, but looks like a good one.

- Even if the cloud just securely manages IoT, this already brings value!
- Apps for the IoT edge: maybe a case where 1+1 is way more than 2? E.g., could secured IoT devices plus apps that leverage them take off?