IPCO

Technology Advisory Panel

PO Box 29105, London

SW1V 1ZU

2 November 2021

TAP Cloud briefing document

Introduction to Cloud – what actually is it?

1. Cloud is using computers and on-line services without worrying about buying and managing hardware computers yourself.

2. Facebook is a Cloud service. You can read messages and post new ones without using a big computer that you need to maintain and understand. When Facebook gets a new, prettier, user interface, you don't have to change anything – all the work is done by them, in the Cloud.

3. When you change your phone or tablet, you just need to tell it your main passwords, and you can access all your previous messages on it without a problem. They are stored in the Cloud.

4. Hasn't it always been this way? No. It used to be much more work for the user, and the computer owner, but it was also easier to understand where all the data was and what was processing it…

History of Cloud – what made it happen?

5. In the ancient times, computers were large and very expensive. Only governments, universities, and big businesses had them. People at home may have had some dedicated electronics for simple games on a TV screen, but they didn't have email, and computers didn't really talk to each other.

6. The basic PC/Mac arrived 40 years ago, and most businesses could then use them for accounts and inventory and word processing.

7. Email really got going about 30 years ago, as all the PCs started to get connected to each other over modems (that weird noise a Fax machine makes when it's trying to connect is a modem, which connects two computer-like things together over a phone line). Email would be sent, stored in, and collected from a few big computers designed specifically for that job and acting like electronic sorting offices. All email messages were eventually stored on the PC of the receiving person, so backing up your computer was critical, or you'd lose everything.

8. Even with the invention of the World Wide Web over 30 years ago, for those who had the bandwidth and were interested in the content available (mostly academics!), the important data still lived locally, on your own home or business computer.

9. About 20 years ago, the first UK broadband lines arrived, and the connections were both faster (maybe 3 Mbit/sec download!) and also "always-on". This meant you didn't have to pay per minute, like an expensive phone call. The bill was per month, which meant you might as well stay on there for longer…

10. Crucially, broadband also meant uploading data was much faster than ever before, and pretty soon, people were so confident of being able to view "the Internet" whenever they needed to, that they no longer had to keep all their data on their own computer, at home.

11. The original World Wide Web had really just allowed you to download a file from a big computer somewhere on the Internet to your own computer, so you could read and edit it, before maybe sending it back. It was like email for documents.

12. Search engines took this to the next level: if you asked for "books by Neil Smith", the server would run around looking on all the other computer index files, find where all Neil's books were described and sold, make a page with links to those books, and maybe a short snippet or summary, and then send you that page. You could click on the link to a book you liked and download (or even buy) it.

13. This search results page you downloaded was made just for you, and only when you asked for it. If Neil had uploaded a new book a few minutes before, it would have been included.

14. Making pages up according to user requests and actions is called Web Services and is what Cloud services are all about. Facebook is a Web Service. You click Home and it makes a page with different sections, and new people you might want to connect with. If you click one of those names (links), it makes up a new page with a focus on just that person.

15. So, what do we need to make this all work well? Good internet connection speeds (only really available to most people in the last 10 years or so), good user access equipment (screen, keyboard or touch surface, enough CPU power to draw the graphics out nice and smoothly), and the commercial models to make sure the people providing the services still get paid somehow.

16. What we don’t need at home any more is lots of data storage, and computers that we understand and can maintain. A smart-phone or tablet is fine for most people.

17. Nowadays, the computers that provide these great user services and experiences provide Web Services to each other, too. One company might do an amazing job of converting images from one format to another, and only want to do that. A company like Facebook, which wants to allow users to upload their own photos, and then display them in a consistently pretty way, might send user images that are posted to them in weird formats to this company, by posting them to a special web ‘site’, and getting the converted images back in a common format. Other specialist services might be text language translation.

Different flavours of Cloud – and examples of each

18. So, Cloud is lots of computers all talking to each other, and users, over common web services protocols, a bit like web pages? Yes.

19. Is everyone doing the same thing, then? No. There’s a sliding scale of how much involvement people and companies have over the design and control of the computers doing all the work. These each have snappy acronyms and/or abbreviations, plus commercial offerings:

On-prem

20. “On premises” means you have all your own computer servers, and data, held locally. You know where everything is, but you also have to manage and support (back-up, update the operating system, install anti-virus, change the disks when they break, etc.) everything. For people with sensitive data, this might be essential. It’s often the default for people who have ‘always done it this way’, too.

21. Sensitive medical databases might be an example of where security is so important (and the user access is so limited) that the overheads of having to manage the hardware locally are well worth it.

Co-location

22. Co-location is where an external company looks after some of the annoying bits of supporting big, noisy, computer server hardware for you. They will provide stable power, cooling, good networking connections to the rest of the world, and physical security (a big co-hosting location will have security staff and passes, and CCTV, so people can’t just wander in). You will still have to provide, install and remotely manage the servers themselves. You are just paying to have them out of your own premises, so you don’t have to secure and cool them, and give them good network connections 24/7.

23. This is often a first step “into the Cloud” for companies, as their server numbers get larger, and as the number of customers rises: they can’t rely on one office building housing their customer-facing web servers, so they move them to a purpose-built “co-lo” location.

Infrastructure as a Service (IaaS)

24. Once companies get used to the advantages of not having local servers, they start to push the envelope further. "Do I really have to maintain security patches for 20 servers, and do my own back-ups? I just want to provide the main web service." This is often the start point for start-up companies and SMEs.

25. Infrastructure as a Service is where a company provides remote access to what looks like a dedicated server, running the latest Windows or Linux OS, and does the boring admin tasks for you. If one of the hard drives dies, they will replace it without you noticing they've done it. All you do is develop and provide your web service for your customers. You still have to worry about managing user data and passwords for your staff, but the annoying hardware-dependent parts have been taken off your hands.

26. This stage starts to get really flexible, too: if your current 'server' isn't powerful enough to serve all the requests from your growing customer base, you can quickly add another server, and have your critical data copied over, so you can run two at once. No buying hardware or travelling to a co-lo site in a van: use your credit card and you have doubled your server capability in 2 minutes.

27. Most businesses without super-sensitive data get to this stage quickly after getting used to co-lo.

Platform as a Service (PaaS)

28. The complete journey (as far as hardware goes) is where a customer company just specifies what it wants: storage at different back-up levels and sizes, compute facilities with the ability to auto-expand as needed, user accreditation and certificate management, networking segments and firewalls, data and service resiliency by cloning everything across multiple regions and countries.

29. If it needs a database, then instead of asking for a server (IaaS) and then installing database software on it to make a working database, it asks for a database and the Cloud company does both the server and database software maintenance in the background. The customer just gets a reliable database service. Most of the difficult functions in a modern company can be provided in this way, including the networking to glue everything together.

30. Amazon Web Services (AWS) is the biggest IaaS/PaaS supplier, with Microsoft (Azure) next in line, and others way smaller.

31. Most modern start-up companies work solely in this way. They wouldn't dream of owning or supporting actual hardware servers, unless they needed something really bespoke. Most medium and large companies would like to operate this way, too, but find it harder to convert everything to this mode, as they have much more currently-deployed IT, and often different services (maybe from company acquisitions) running on different IT, with different controls.

32. The engineering people find this PaaS approach much easier. The legal people can find it quite tricky, as achieving resilience usually means having multiple copies of data in disparate places (so they don't all break at once). It might help resilience to keep a live copy of EU data on the far

side of the USA, and store a slower back-up copy in China, but the legal regimes might then mean “Lawful Access” is now available to US and Chinese government officials for the EU business. This was one of the drivers behind the development of the GDPR.

Micro-services

33. Is PaaS the end of this road, then? Not quite, no. Nowadays, some people are looking at the discrete service examples (e.g. the image transformation service, above) and asking why they’d need to care about whether it was running on a small or big server. Increasingly, big services are made up from smaller ‘lego brick’ services that are each easier to develop and test.

34. These collections of micro-services (“Lambda services” in AWS, and Azure Automation in Microsoft world) further distance the engineers from annoying real-world constraints of hardware and geographical location.

What are Private Cloud, Hybrid Cloud, and Secure Cloud?

35. Private Cloud is building your own complete Cloud infrastructure, as though you were going to sell it to the public, like Amazon and Microsoft, but keeping it on-prem, and using it only for your own data and services, without any insecure access to the real-world. Often achieved in the real world by contracting with specialised teams in the companies that provide large-scale Cloud services to the public to come into your building and manage your own infrastructure; this ensures the latest technologies can be moved into your Private Cloud as soon as possible after general release.

36. What’s the point of that? It allows companies or government agencies to maintain the security of on-prem solutions (no one external can access the data) whilst allowing their development and support engineers to learn and use the latest Cloud technology approaches. If a company wants to employ a skilled developer nowadays, she may well have spent the last 5 years working in a Cloud environment. A Private Cloud allows the company to bring her in with no friction. Private Cloud means the company or agency has to keep the Cloud infrastructure updated with the latest tech as it happens, just like Amazon and Microsoft have to, but with far fewer users over which to amortize their costs.

37. A Hybrid Cloud looks similar to a Private Cloud (to the engineers), but some of the less sensitive data stores and services are actually in the main (public) Cloud. Managing the interfaces is tricky, because an attacker could access the public-facing parts and try and work their way towards the sensitive data stores. Maybe back-ups are fully encrypted on-prem, and then resiliently stored on a public Cloud storage site.

38. Secure Cloud is where a Cloud Services company provides a Private Cloud for a customer at secure locations (like military bunkers or dedicated datacentres) and provides secure comms to and from those secure sites for the customer’s staff. It’s a commercial arrangement that allows the customer to have everything a real Public Cloud environment would offer (almost) without the security concerns of trying to manage the infrastructure or access, or to have the complete Private Cloud inside their own premises. This is similar to Private Cloud, but the ‘bunker’ might have 100 companies’ Cloud environments being managed in it.

39. Why wouldn't everyone just have a Secure Cloud? It's great for a clean start, but not for companies with large existing IT and data estates. The provider has to keep up with whatever Amazon and Microsoft are offering, so it's almost always one of those two suppliers that has to be used. A number of clever AI services (like language translation) only work on a real public cloud, because they fall back to using real people to do the work if the AI fails. These can't be replicated in a Secure or Private Cloud.

How people, companies, criminals, police, and security agencies exploit these services

40. OK. We've explained how people and companies might gain benefit from these Cloud services as we went along. In summary: developing and deploying new services is easier; costs and management are in a single place; they can even be more secure than on-prem and co-lo, as the hardware-based attacks and the 'lazy Operating System (OS) updater' risks are removed. The problems can be legal ones around where customers' data are being kept, and whether data is really deleted or not.

41. Criminals can use these services to avoid the more basic evidence-gathering approaches of (say) breaking down the door and seizing all the servers and desktops. If they can run their (illicit) businesses entirely in the Cloud, and protect their access devices (smart-phones, tablets, maybe laptops) with good biometric security, then seizing IT hardware becomes pointless. Child Sexual Abuse Imagery hoarders can easily store their images within some on-line storage service run from a country that doesn't have cooperative law enforcement relations with the UK.

Police

42. Police can use Cloud capabilities to improve their effectiveness in analysis, and their presentation of evidence. Using the Cloud securely isn't very easy for non-engineers, so mistakes can easily be made which are even worse than exposing the data on one desktop machine. Lots of data has been exposed by naïve users accidentally making private storage buckets world-readable (!) This will happen to organised criminals, too, as they start to use Cloud services like a legal business would.

Security Agencies

43. Many of the 'Big Data' techniques developed by government agencies and data-scraping companies like Google only really work in Cloud-like environments, so in that way, data-focused government agencies are already Cloud-natives. Security adversaries (e.g. nation-state actors and global threat actors) are themselves being forced to use these Cloud approaches and may well make many security mistakes as they learn, which investigators can exploit.

44. Despite the security instinct to keep all data within a security boundary, using Private Cloud will get increasingly difficult, as what Cloud can achieve is now driven mainly by commercial uses. Using Hybrid Cloud architectures skilfully and securely is probably the answer, as well as contributing much more strongly to the Open Source projects that are driving Cloud capabilities.

45. The United States Department of Defense (US DOD) is looking to a Secure Cloud, but the "only two vendors" problem, coupled with the huge budget available, is slowing the project down with legal challenges from whichever supplier is currently missing out in the process.

Frequently Asked Questions

Q1: why is it called "The Cloud"?

> A1: on normal network architecture diagrams, details are given about the computers and network elements in the interesting locations you are trying to describe, but the general Internet that connects all the sites together is usually draw as an amorphous cloud, to show "lots of computers connected together". This image of a cloud to represent the "Internet beyond your known infrastructure" has become common, so services offered "in the Cloud" means "somewhere in that Internet blob I don't know the details of exactly".

Q2: is it possible to define where you want your data to be?

> A2: within large legal boundaries, yes. You can say "only in the EU" or "only in Ireland", but you can't say "in Lambeth". If you want it that localised to a known place, then you have to use a "Co-lo" which robs you of most of the advantages of modern Cloud services.

Q3: why would I want to use a service that I don't understand fully for protected data?

> A3: you may want to use complex analytics on it (involving a Machine Learning task) and that will need many computers working together to complete a model (say 1,000 for a day), but then afterwards running the trained model only needs 1 computer. Cloud services allow you to do this sort of task easily and cheaply; buying in 1,000 'real' computers (on-prem) and only using them for a day wouldn't be economic.

Q4: does anyone really understand what's happening inside a Cloud environment?

> A4: the environments are well-designed and built in functional layers with lots of logging between each active component, so it's easy to understand as much as you'd like about what's happening in your Cloud. Some of the ways of working with Cloud use techniques that don't map directly to how you would do things in an on-prem solution (the networking for instance), so sometimes pre-Cloud engineers may take some time to properly understand a Cloud infrastructure and assess its risks.