

Date: 28th September 2018
Released by: The [Coed:Ethics](#) Community
Authors: Paul Johnston & Anne Currie
Title: The State of Data Center Energy Use in 2018
License: Creative Commons 2.0 (Attribution)

Executive Summary

This whitepaper is an overview of the state of energy use in data centers (DCs) worldwide in 2018 and the likely changes over the next 5 years. It's aimed at companies and individuals who purchase data center capacity as servers on prem, in co-location facilities or as Cloud resources. That's basically every enterprise on the planet.

As major global consumers of electricity, data centers play a key role in any move to restrict climate change and maintain the global security of energy supplies. The tech industry's electricity strategy is therefore of interest to all enterprises and citizens. In this paper, we outline the positions of the major cloud providers and on premise DCs.

Recommendations

We recommend the tech industry adopts a policy of secure, 100% sustainably-powered servers for data centers by 2024 (for example: solar, wind, hydro or nuclear).

Specifically, we suggest:

- Enterprises choose 100% sustainable locations for new Cloud instances (e.g. use Azure, Google Cloud or the sustainable AWS regions of Dublin, Frankfurt, Oregon or Canada) and transition existing VMs there as soon as possible.
- Enterprises inform their data center providers they want secure, sustainably-powered compute resources by 2024.
- Individual technologists state their commitment by signing the [Sustainable Servers](#) petition.

Cloud League Table

Cloud	Rating	Sustainable Servers?
Google	B+	100% with offsets
Azure	B	100% with offsets and energy certificates
AWS	C	100% with offsets in 5 regions, elsewhere unknown with estimates in <30-40% range
Oracle	C-	100% with offsets in a few regions <30% overall
IBM	C-	~50% overall

Alibaba	D-	Unknown but China a major market, and not known what energy is purchased
---------	----	--------------------------------------------------------------------------

Introduction

Cloud computing is a vital part of the digital revolution changing our world. The Cloud offers on-demand services via particularly massive data centers.

All DCs require power to operate - currently 2% of the world's electricity. This puts data centers at the current energy consumption levels of the aviation industry and projected to exceed aviation by a factor of 4-5 by 2024.

Electricity is a commodity that most of us in the developed world take for granted. We rarely consider where it has come from, whether it will always be there or how much impact it has. Cloud providers take responsibility for infrastructure and the requisite energy generation away from enterprises - that is their job. However, are enterprises holding themselves and their Clouds sufficiently to account for the climate and energy security impact of the data center capacity they purchase?

Caveats

We have done our best to investigate and get the facts straight in this report but it is very hard to get hold of this data. If we have missed some fact that changes the position of any of the cloud vendors, please let us know and we'll be happy to update this living document.

Electricity, The Grid and Data Centers

How do data centers consume electricity and provide services?

Demand

Data centers provide their customers with Service Level Agreements or SLAs that guarantee operation. A loss of power, cooling or connectivity - all of which rely on electricity - would break the SLA. To avoid that, data centers use two or more forms of electricity supply and failover mechanisms.

The DC's primary supply will usually be from a "grid", e.g. the National Grid in the UK, which transmits electricity from generating stations to demand centers via distribution lines.

The purpose of a grid is to abstract away the details of electricity supply. When you use grid current you can't choose the source - you get whatever energy mix is immediately available wherever you are. You can see the mix of generation right now in the UK at <http://gridwatch.co.uk/>.

When you buy a renewable energy tariff as a consumer, you are simply requiring your energy supplier to:

- purchase sustainable energy from some generation company equivalent to your usage
- feed it into a grid somewhere.

So, if you pay a renewable energy supplier and plug into the grid then when you flick the switch you still get the same mix of energy as everybody else but you will have kindly offset your usage with more renewable power generation elsewhere.

No grid is infallible and there will occasionally be periods where it can't supply electricity (e.g. due to transmission lines failing). For those periods data centers employ backup supplies to provide power until the grid is restored. That backup might be local diesel generators or large battery storage. The backup ensures the DC meets its SLAs.

Data center power use is not constant or consistent. Usage changes depending on demand. An increase in the number of servers in the data center or a spike in requests could increase power requirements, network, or cooling needs.

Supply

A generation company creates electricity and provides it to the grid through a Power Purchase Agreement (PPA). For example, company A builds a wind or solar farm and the grid agrees to buy the resulting energy at a specified value.

You could provide electricity directly to your data center from solar panels on the roof, for example. This direct electricity usage is the most sustainable way to power anything because transmitting energy across a grid is lossy. However, it is unlikely that a roof or local area will have enough space or capacity to provide all the energy required for your DC. Alas, you also can't guarantee either sun or wind conditions for an SLA so you would still generally need a grid connection.

Offsetting Carbon Emissions

So, a **Carbon Offset** is a reduction in emissions of carbon dioxide or other greenhouse gases in one location so that emissions of those gases somewhere else are effectively cancelled out.

If a company buys from and sells electricity to the same or a different grid network, then they can be said to offset their usage. For example, company B with a data center that draws electricity from a grid may build a solar farm 1000 miles away to generate the same amount of energy and sell it to another grid. If the company didn't want to build their own solar farm they could buy the power from an existing one. That's all offsetting. Most of the cloud providers' renewables claims are managed via large scale offsetting rather than private wire solar, wind or other direct forms of sustainable energy.

Offsetting is great, but it raises an important issue. If data center electricity demand grows then the DC provider needs to purchase more renewable energy to offset the increased usage. If use grows fast then building renewable capacity quickly enough to cope may be a problem, in which case the only readily available source will be fossil fuels.

Offsetting works well in a stable scenario but where data center usage is increasing then it becomes more of an issue. If DC capacity will grow the best strategy is to get ahead of it by creating new energy sources.

The Green Cloud

Most servers are operated either "on premise"; in a server co-location facility; or via the Cloud. On premise or in a co-lo, companies have access to physical server hardware that needs to be powered. In the Cloud, companies have access to virtual machines and electricity use is hidden.

For an enterprise, there are significant advantages to the Cloud:

- no more managing machines, networks, and power
- hardware on demand, reducing capital risk
- access to specialized hardware (e.g. FPGAs)
- high quality additional services (e.g. secure backups, distributed databases).

As a result, research by Cisco suggests [95% of data center traffic will be cloud traffic by 2021](#). Even if that is optimistic, we can plausibly predict that "most" data center traffic may be Cloud within 5 years.

The other potentially good reason for moving operations to the Cloud is efficiency. According to AWS, Cloud providers achieve ~65% efficient use of resources (including electricity). In comparison, Gartner suggests on premise servers are 10-15% efficient on average. This suggests the Cloud should be a more environmentally-friendly hosting approach in the long run. However, there is a problem.

Cloud flexibility makes it very easy to increase usage. Indeed, research from [Huawei estimates that IoT, Machine Learning and AI may increase DC traffic five-fold by 2025](#). In our opinion, this is a conservative estimate. Without the traditional help from Moore's law It is currently unclear how the Cloud will cope with a 5x increase in demand without both relying on fossil fuel energy and putting the security of other electricity supplies at risk.

Options

Where companies own their own servers (on prem) they can buy offset electricity from renewable providers. The energy cost is usually similar to non-renewables and as the current is still drawn from the grid there is no increased supply risk.

We strongly recommend that enterprises who operate their own DCs switch their energy provider to a sustainable, offset one. We advise co-lo customers to contact their DC and ask about their energy mix and plans for moving to sustainable sources.

Cloud consumers have may fewer options for sustainability, but still a lot more than they realise. In the next section we will review the current renewable status, commitments and customer options for the major Cloud providers.

Gartner Top 6

in 2018, Gartner listed the six top cloud providers in their [Cloud/laaS report](#).

- AWS
- Azure
- Google Cloud
- Alibaba Cloud
- Oracle
- IBM

We focus on these below.

Amazon Web Services (AWS)

AWS, the largest Cloud provider, has a [long term goal to go 100% renewable](#) for its global infrastructure. That is a great target. We would love to see a date on it.

"Carbon Neutral" regions

The good news is AWS already has [5 carbon neutral regions](#):

- US West (Oregon)
- European Union (Frankfurt)
- European Union (Ireland)
- AWS Govcloud (United States)
- Montreal (Canada)

AWS customers can already run applications in these regions and by default they will be carbon neutral. To re-iterate, simply moving applications or environments to these AWS regions will automatically offset your servers' carbon emissions.

What "carbon neutral" means here is not discussed by AWS, but our assumption is that the electricity used is offset by renewables into the same grid.

Our view on AWS

AWS are clearly planning and executing an electricity strategy. Newer regions are usually more sustainable than older ones. Their Chief Evangelist [Jeff Barr has written about why cloud is better for the environment](#). AWS has also been [investing in renewable capacity](#). In addition, [they have been participating in groups aiming to advance energy policies](#). It

appears that AWS are hedging long term electricity bets and if they are so should every DC operator.

AWS are making good progress on sustainability. However, they do not release much public data. The AWS console also lacks UI nudges that would help their customers make more informed choices, such as flagging sustainable regions with a green label.

We recommend that AWS customers move their instances to carbon neutral regions where possible and site all new instances in those regions. They can also inform their AWS account managers that they want more renewable hosting (AWS are driven by customer demand).

Microsoft Azure

“Across the tech sector we need to recognize that data centers will rank by the middle of the next decade among the large users of electrical power on the planet.” Brad Smith, President Microsoft

Like AWS, Microsoft make much of the efficiency of large, optimised Cloud data centers over private ones ([93% saving over on-premises](#)) A full report is available [here](#).

Microsoft are keen to remind us "that consolidating distributed electricity demand from on-premises DCs into the cloud unlocks the potential for [large-scale purchases](#) of green power". That can spur the substantial funding required to bring major renewable projects online. Microsoft further state their goal is to "achieve 60% direct purchase of wind, solar, and hydropower early in the next decade".

Microsoft have a [strong commitment to carbon neutrality](#) and every business unit charges a fee on carbon that goes towards this end. The organisation also buys from green energy sources and purchases renewable energy certificates to reduce carbon emissions. This has led to Microsoft as a whole, and therefore Azure, being 100 percent carbon neutral since 2014.

Microsoft's strategy to purchase renewable energy certificates where they cannot use direct sources (because the data center is in an urban area, for example) is a decent short term way to ensure carbon neutrality.

Note that Microsoft are also working towards improving their efficiency through data center design and are part of the Open Compute Project which exists to further this goal e.g. [storage and security](#).

Our View on Azure

Microsoft have achieved a great target for renewable energy use for Azure - 100% carbon neutral. Their use of renewable energy certificates is good, but further investment in renewables would be even better (they plan to achieve 60% this way).

Azure customers are sustainable by default in that all the energy used to host their compute is either directly provided or offset. However, it is not entirely clear which regions are directly powered by total or partial renewable energy and which are offset with renewables or energy certificates.

Google

"Google will match 100% of the energy consumed by our global operations in 2017 with renewable energy and maintain our commitment to carbon neutrality." [Renewable energy at Google Cloud](#).

Google clearly states they offset their customers' digital footprint with clean energy and reduce the impact on the environment. In fact, [Alphabet is the biggest corporate buyer of renewable energy on the planet](#).

This should be applauded and celebrated. It is a big step in the right direction and shows that a successful technology provider can be a leader in renewable and sustainable energy. For further details read Google's [blog with stats](#).

"We say that we matched our energy usage because it's not yet possible to "power" a company of our scale by 100 percent renewable energy. It's true that for every kilowatt-hour of energy we consume, we add a matching kilowatt-hour of renewable energy to a power grid somewhere. But that renewable energy may be produced in a different place, or at a different time, from where we're running our DCs and offices. What's important to us is that we are adding new clean energy sources to the electrical system, and that we're buying that renewable energy in the same amount as what we're consuming, globally and on an annual basis."

Our View on Google

Google are doing a great job here. One point to note is that while Google is saying that your servers are sustainable (and they are) it is not saying that your servers are not causing new carbon emissions. They are merely offset. However, Google and Azure are the only major providers that have met a 100% sustainable energy goal, making them by far the industry leaders on sustainable servers.

Alibaba Cloud

Data for the sustainability of the Alibaba Cloud is, unfortunately, hard to come by. Lots of cooling and efficient "green" data center technologies are mentioned in [this article](#). There are also several mentions online of Alibaba Cloud attempting to be "green". However, not many articles specifically provide the key information, which is its purchase of renewable energy or offsetting. e.g.

<https://www.alibabacloud.com/press-room/alibaba-cloud-to-help-empower-green-olympic-winter?spm=a2c5t.10695662.1996646101.searchclickresult.38af1923S7z2m0>.

Many of Alibaba's data centers have been built with innovative cooling technologies, and there appears to be research into interesting cooling options such as placing data centers in

favourable locations. However, there is evidence that their efficiency efforts are not focused on the key problem of the electricity mix in the Chinese grid that powers a large proportion of the data centers that Alibaba Cloud uses ([Alibaba leading China's push for cleaner DCs](#)).

Our view on Alibaba Cloud

While there is a global expansion of Alibaba's cloud services there are no public statements regarding offsetting of carbon emissions similar to those given by other providers. Nor is there any commitment to renewable sources. We therefore have to assume that using Alibaba Cloud probably has a negative impact on carbon emissions.

Oracle

Oracle proudly displays its credentials as a "Green" company, being ranked #17 in the world and #10 in the US in the [Newsweek Green Rankings 2016](#).

It also talks about its data center efficiencies, claiming to operate some of the [most efficient data centers in the industry](#).

In terms of renewable energy, they are aiming to "increase" the amount of renewable energy they use with some regions (notably the United Kingdom) using 100% renewable energy. Our assumption is that Oracle offset the energy rather than use directly supplied renewable energy.

In the [Oracle Corporate Citizenship Report](#) they state that 29% of their total electricity use came from renewable sources.

"As more sources of renewable energy become available and as open, competitive markets evolve, we plan to increase the amount of renewable energy we procure. In certain locations, including the United Kingdom, our data centers are using 100 percent renewable energy. *Moreover, renewable energy is one of our key considerations in data center site selection, including co-locations*".

Our View on Oracle

Oracle have an established goal to reach 33% renewable energy usage by 2020. This seems modest in comparison to the goals set by other providers, but it's not insignificant. Oracle customers have some information that would allow them to choose to host in renewable regions, for example the UK.

Oracle do include their total Greenhouse Gas Emissions in the above report and appear to be committed to reducing them. This is a good sign.

IBM

IBM have long standing environmental and energy policies dating back to the 1970s and have a strong history of energy conservation and climate protection programs. More about their [Climate Position](#) can be found on their website.

IBM has been increasing its [renewable energy procurement](#) over the years, and publishes information on the amount of energy it buys across all IBM-managed locations. For example, in 2017 they purchased 779,000 MWh of renewable electricity (22.9% of IBM-managed locations).

According to their [purchases of renewables for data centers](#) "nearly 50% of electricity procured for IBMs managed data center operations came from renewable generation sources". However, in co-location data centers only 27.4% of total electricity consumption came from renewable-generation sources. Fortunately, IBM places most of its data centers in locations managed by IBM.

IBM provides good statistics on its position and demonstrates commitment to increasing usage of renewables for data center operations. It doesn't identify clearly whether the sources are offset, which is likely.

Our View on IBM

IBM has a [good position](#) on Climate Change and a history of taking action in multiple areas, including employee commuting, and seeing the problem as societal, economic, and governmental.

At just below 50% renewable, their data centers are more sustainable than some other providers, notably Alibaba, Oracle and AWS (outside of specific sustainable regions). However, customers are unable to place a workload on IBM's systems that could be 100% renewable as they can with several other providers. That is an area that could be improved.

The Current State of Cloud

Across the six providers above, some have shown a strong and clear commitment to offsetting their data center energy with renewables and others have made more limited efforts in this area.

It is difficult to guess the exact amount of renewable energy that is used to power the Cloud, as that would require usage information from the vendors themselves. However, we estimate approximately 25%-40% of current cloud infrastructure energy usage is being offset by renewable energy sources.

The good news is sustainable server usage is likely to be significantly better in the Cloud than for on premise DCs. However, as stated above, [Cisco](#) and [Huawei's](#) research gives us cause for concern on the current rate of progress in renewable energy generation. If their projections are accurate, data center usage will increase five-fold by 2025 and we'll see a significant increase in energy requirements for what is likely to remain a handful of Cloud providers. We will need a massive increase in electricity provision worldwide to both satisfy data centers and maintain energy security for the rest of us.

This makes it even more important to get robust data center renewable power targets. As it stands, by 2025 assuming Moore's Law remains somewhat offline we'll have to build at least 10x the current renewable capacity used in data centers to make them sustainable. Without such a target, it is likely that the tech industry will become one of the largest polluters on the planet.

Offsetting Is Not The Final Answer

Note that offsetting is a very good solution for the current energy management scenario we have. However, given that energy usage is going to rise due to the growth of our and other industries, it is important to understand that 100% offsetting for tech will only slow the global increases in carbon emissions. It will not reverse them.

We believe that the technology industry must also be proactive in the development of new power sources for data centers, starting now. Relying on offsetting in the medium to long term will be not be enough. Sustainable servers by 2024 is only the starting point. We will require new DC power solutions over the next few years to meet the scale challenges of our energy needs. Our growth will continue upwards after 2024.

Sustainable Servers

Tech is a huge user of electricity but most of our industry doesn't have a strategy for making that use sustainable.

Data centers already require [~2% of the world's electricity](#). All of tech (including devices) uses around 12%. According to The Economist, Bitcoin mining alone takes an ["eye-watering amount of power"](#) (equivalent to the energy use of a small country). In one recent report, the power requirements from new AWS DCs in multiple US states may have significantly increased the [cost of electricity to poorer citizens](#). Perhaps it is no surprise AWS keep their DC locations secret.

How can we stop hurting and start helping? Surprisingly easily. All we need to do is pay attention and make a minor effort.

Our industry relies on electricity. We are major global consumers of it. We cannot ignore the impact this has on our own and wider energy security and the climate. Fortunately, the tech sector is incredibly innovative and we're forward-thinking people; we have an opportunity to support a global revolution in the energy industry and ensure that our growing energy requirements are met by new, more sustainable generation capacity.

The tech industry could lead the way by committing to run all our applications on Sustainable Servers - i.e. servers running on renewable, nuclear or offset power - by 2024.

Consumer Action

The good news is that 100% Sustainable Servers are achievable today.

- For Google Cloud users, you've already done it. Congratulations and thank you.
- For Azure users, you've mostly done it. Congratulations. Do ask Microsoft to speed up their switch from carbon energy certificates to renewables.
- For AWS users, you can do it by transitioning your instances/resources to sustainable public regions (currently Dublin, Frankfurt, Canada and Oregon).
- For Oracle Cloud, you can move your servers to the sustainable UK region, for example.
- If you operate your own servers on prem you can choose a renewable energy supplier.
- For everyone else, tell your Cloud or co-lo provider that Sustainable Servers - your Data Centre powered or offset by renewable energy - by 2024 is what you want.

Energy Security is Not Just About The Environment

A significant increase in carbon pollution is a big problem for the environment of course, but growing energy needs also require a growth in infrastructure to maintain cheap and secure electricity supplies for everyone.

It is not possible to foresee exactly what the political climate will be like or the price of electricity in 5 years time but using [UK government statistics](#) it's reasonable to infer that oil and gas prices will go up by around 10-20% and that is likely to increase costs to both consumers and data centers.

A renewables installation is a 20-30 year investment and can provide a good ROI over that period. It also delivers more stable energy prices. It makes sense for Cloud Providers to hedge their bets on energy costs by fixing a long term price on their electricity. We suspect that is the main motivation behind the moves to renewables described above. We expect more companies to explore renewables projects as the cost of electricity starts to rise, especially if governments begin to implement measures to combat carbon emissions via regulation.

Sign The Pledge

Sustainable Servers run on 100% sustainable energy either offset or via direct power. The tech industry could lead the world in energy investment. That would be great for our companies and for society by giving us a cleaner and more secure world.

We want as many industry individuals as possible to sign our simple statement: "Sustainable Servers by 2024"

[Sign The Pledge](#)

We will champion your views to the Cloud providers. Please help.

This paper is available under Creative Commons CC 2.0 Attribution license. Please link to this paper.

Addendum - AWS 2019 (added 10/Nov/19)

Since we wrote this paper in 2018, there has been a sign of progress from AWS. They have set a target for [80% sustainable servers by 2024 and 100% by 2030](#).

This is a major step forward, although it doesn't meet our target of 100% by 2024 (a goal Google has already achieved).

In a recent talk at the MapCamp conference in London, Adrian Cockcroft, AWS's VP Cloud Architecture Strategy, stated the position clearly: if we want sustainable servers faster we need to ask for them or, better yet, move our workloads to existing green regions.



Authors

<p>Anne Currie Anne is a tech industry veteran engineer, startup founder, writer and speaker. Based in London, she is the Chief Technology Strategist for Container Solutions. Twitter: @anne_e_currie</p>	<p>Paul Johnston Paul is a multiple startup CTO, interim CTO and a Director of a Community Energy company. He is also a former Senior Developer Advocate for Serverless at AWS. He has many years experience in cloud, technology and energy companies. He is based in Milton Keynes in the UK. Twitter: @PaulDJohnston</p>
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Other useful information sources

- [Guardian report on IoT and carbon emissions with lots of useful links to studies](#)
 - Proposed data center in Ireland... Needs 8% of the grid's electricity to run and 144 diesel generators as backup.
- [Greenpeace Click Clean!](#)
 - Upside: pressuring companies
 - Downside: not pressuring "the internet" or private DCs where it is a much bigger issue
- [Renewable Energy Adoption and Data Centers](#)
- [Microsoft 100% notice](#)