

Saving Energy in Europe by Using Amazon Web Services



**Black
& White**

October 2021

Commissioned by



451 Research

S&P Global
Market Intelligence

©Copyright 2021 S&P Global Market Intelligence. All Rights Reserved.

About this paper

A Black & White paper is a study based on primary research survey data that assesses the market dynamics of a key enterprise technology segment through the lens of the “on the ground” experience and opinions of real practitioners — what they are doing, and why they are doing it.

About the Author



Kelly Morgan

Research Director, Datacenter Infrastructure & Managed Services

Kelly joined 451 Research, a part of S&P Global Market Intelligence, in April 2011 as an analyst covering the economics and finances of the datacenter and hosting industries. In her current role, she leads the analyst team that tracks datacenters, interconnection and content delivery worldwide. Coverage includes trends, providers, services, market size, supply/demand, M&A and technology.

Before joining 451 Research, Kelly spent nearly 10 years in private equity, focusing on investments in telecommunications, IT and datacenters. Prior to that, she worked for several years at the OECD in Paris, managing the organization's budget process. She has a BA from Wesleyan University and an MA from Tufts University.

Executive Summary

This study estimates the energy savings and associated carbon-emission reductions of moving business applications from on-premises enterprise data centres in the European Union (EU) to cloud options such as Amazon Web Services (AWS). Building on modelling from a similar study of US firms in 2019 and of APAC firms in early 2021,¹ we estimate that running business applications on AWS, rather than on-premises enterprise data centres in Europe, could reduce associated energy usage by nearly 80% and carbon emissions by up to 96% for many businesses when AWS purchases 100% of its energy from renewable sources.


These energy savings result from the higher energy efficiency of both the servers and data centre facilities of cloud providers such as AWS. This is because cloud providers use server systems (powerful computers that offer their resources over a network) with great attention to power optimisation, integrating the very latest components. These servers run at higher utilisation levels, leveraging the cloud provider's ability to share and dynamically allocate resources among multiple customers. At the facility level, designs that use less energy for both cooling and power distribution boost efficiency for cloud providers in their owned and leased data centre sites. All of this translates into considerably less energy used to perform the same unit of work – such as processing financial transactions, running business operations, executing online orders, enabling government services or serving webpages – than would be required at a typical enterprise or government facility. This reduction in energy consumption and resulting carbon footprint is more pronounced if cloud providers are able to use low-carbon renewable energy sources. AWS is on a path to 100% renewable use by 2025.

There are organisational differences between the cloud and the technology traditionally used by enterprises and agencies that also impact energy efficiency. The cloud business model is focused on delivering IT services at scale: Maximising energy efficiency (and thereby reducing power costs) can have a direct impact on the bottom line. Thus, cloud providers such as AWS have an incentive to make the entire technical organisation work together, from design to operations and from servers to data centres, to improve energy efficiency. This includes operating servers at a higher utilisation rate and designing facilities to reduce energy and water consumption. By contrast, enterprise infrastructure tends to be fragmented across disciplines (server, storage, networks, facilities management, real estate, etc.), with efficiency initiatives that are narrower in focus and harder to scale. Data centre energy costs and the associated carbon emissions are not top priorities for senior leadership at many organisations because digital infrastructure is not a core business.

Our survey and modelling show significant potential for energy savings from a move to the cloud in the EU. We estimate that AWS is *up to five times more efficient than the average EU enterprise and that moving a megawatt of typical workloads from an EU organisation's data centres to the AWS cloud could reduce carbon emissions by up to 1,079 metric tonnes of CO_{2eq} per year*. This compares to our 2019 study of the US, where AWS was estimated to be 3.6 times more energy-efficient since surveyed EU enterprises were found to be less efficient than those surveyed in the US two years ago. For example, while we found that the average power usage effectiveness (PUE) for US enterprises in 2019 was 1.66, the average PUE for enterprises in the EU in 2021 was 2.0, offering even greater opportunity to improve energy efficiency and reduce carbon emissions than in the US.

1. The Carbon Reduction Opportunity of Moving to Amazon Web Services, 451 Research, October 2019 and The Carbon Reduction Opportunity of Moving to the Cloud for APAC, S&P Global Market Intelligence, July 2021

Figure 1: Carbon Reduction Potential of Cloud Infrastructure Compared to Surveyed EU Enterprises



Source: 451 Research/S&P Global Market Intelligence

Grid emissions vary greatly across the European Union, but the European electrical network is largely pooled (the Continental Synchronous Area is the largest grid in the world), with future interconnection upgrades aiming to further connect the European grid. According to the European Environment Agency (EEA), electricity generated in the pooled European network emits 255g of CO₂-equivalent per kilowatt-hour on average, due to a continued presence of fossil fuels in the mix. As such, improving energy efficiency up to 80% by moving business workloads from on-premises infrastructure to cloud services such as AWS could dramatically reduce carbon emissions. For a 1MW (about 1,000-square-metre) enterprise data centre at an assumed 30% electrical utilisation in the EU, switching all applications to cloud services could reduce emissions by about 1,079 metric tonnes of CO_{2eq} per year. This is the equivalent of removing over 500 cars from the roads or offsetting the annual electricity-related emissions for over 50 households across the EU (average).

Bringing additional renewable energy sources online could further reduce carbon emissions. Large cloud providers such as AWS have robust renewable energy procurement programs that are tied to decarbonisation goals. Through corporate power purchase agreements (PPAs) and other efforts, cloud providers such as AWS enable new renewable energy generation on the grids where they operate and reduce their carbon emissions. Amazon's purchase of renewables has added renewable power-generation capacity to the grid, including wind farms and solar projects in Ireland, the Netherlands, Sweden, Finland, Germany, Spain and the UK, that will provide gigawatts of renewable energy. The carbon-emission reductions from shifting workloads to the cloud are estimated at 1,079 metric tonnes of CO_{2eq} per year per megawatt, but when a cloud provider is powered by 100% renewable energy, as AWS has set out to achieve, those carbon-emission reductions could be as much as 1,292 metric tonnes of CO_{2eq}.

Introduction

More organisations are increasing their focus on the impact they have on the environment. Roughly 65% of nearly 4,000 companies around the world in a recent survey² now have public goals to reduce carbon emissions, up from 50% in 2017. The European Union (EU) recently approved an updated Climate Law, which requires cutting emissions 55% by 2030 against a 1990 baseline. Through its forthcoming Corporate Sustainability Reporting Directive, the EU will require many organisations to publish reports on their environmental impact. The UK is on a similar path, having proposed Mandatory Climate-related Financial Disclosures, which will require companies to report sustainability data.³ Sustainable investment funds are on the rise as well, with Europe leading the way, accounting for 70% of global sustainable investing.⁴

As companies seek to increase energy efficiency and reduce emissions, IT infrastructure is one potential avenue for improvement. IT equipment and the data centres housing it vary in energy consumption. Seemingly small improvements can have a large impact on overall energy use; however, for most organisations, running data centres and IT equipment is not a core competency, and they do not have the experience or expertise to maximise energy efficiency. To determine the room for improvement, we modelled the energy use and energy efficiency of IT equipment and data centres for a variety of organisations and compared them with the energy efficiency of cloud services,⁵ using Amazon Web Services (AWS) as an example. According to our analysis, moving workloads to the cloud could dramatically reduce the carbon footprint of most organisations' IT operations.

To understand the energy and environmental considerations of enterprise businesses across the European Union, 451 Research, a unit of S&P Global Market Intelligence, surveyed senior stakeholders at over 300 companies from a broad set of industries across France, Germany, Ireland, Spain and Sweden. The companies surveyed have annual revenues of \$10m-1bn, and they range from well-known enterprises to businesses that face IT challenges similar to those of larger firms, but with smaller budgets and less IT expertise.

Utilising these survey results, other 451 Research and third-party industry data on cloud service operations, we devised and populated an energy-efficiency model that offers a 'chip to grid' view. At the core of this model is a measure of the amount of work performed for the total energy used by the data centre infrastructure, covering both data centre facility-level power usage effectiveness (PUE) and server-level energy efficiency. As we elaborate below, analysis of the model results helps us better understand just how much organisations across the EU can reduce the energy consumption and related carbon footprint of their workloads by moving them to the cloud from on-premises data centre infrastructure.

2. The KPMG Survey of Sustainability Reporting 2020

3. <https://www.gov.uk/government/consultations/mandatory-climate-related-financial-disclosures-by-publicly-quoted-companies-large-private-companies-and-llps>


4. ESG funds defy havoc to ratchet huge inflows, FT, February 6, 2021

5. Throughout this report, we use the word 'cloud' to refer to what is commonly known as hyperscale compute and data storage services. We define this as providing IT resources that can scale in a matter of minutes, practically without limit.

The Energy Efficiency Model

Building on the methodology employed in the 2019 US study, 451 Research began with a survey of 308 enterprises across France, Germany, Ireland, Spain and Sweden from a variety of industries.

Figure 2: What Industry Do You Work in?



Source: 451 Research/S&P Global Market Intelligence (n=308)

We asked EU survey respondents a series of questions about their server infrastructure and operations. We used these responses in our model to establish profiles of enterprise server infrastructure in the EU, focusing on two main measures of energy efficiency: the age of servers in the installed server base (as servers have become more efficient over time) and their utilisation rates, which also impact energy efficiency.

Our survey results show that, on average, the EU companies surveyed keep their servers for slightly more than four years before upgrading, somewhat shorter than the nearly four and a half years reported by US enterprises in 2019. In the European sample, larger businesses upgrade more quickly than smaller companies, which is different than what we found in our US and Asia-Pacific (APAC) studies, in which we found no meaningful correlation of company size with server replacement. Among EU member states, Spanish and Irish firms were more likely to hold onto their systems a few months longer on average, whereas surveyed organisations in Germany and Sweden replaced their servers faster.


We used survey responses to estimate the proportion of various server technology generations in EU enterprise server stock. This is influenced by average age of the hardware, how quickly new technology is adopted and the change in server numbers over time.

Adopting the latest server technology has a material impact on energy use because chip technology and overall system design have become increasingly efficient. In our sample, EU companies on average start buying new server technology nine months from its general availability, which is longer than what we saw in either the US or APAC. This lag contributes to the efficiency gap between enterprises and cloud service providers such as AWS, which are known to access the latest server technology as much as a year ahead of its availability to the market at large. This is because cloud providers are willing to design, test and fund new chip and server technology.

The 451 Research model then estimated server utilisation levels by analysing adoption of virtualisation⁶ and workload consolidation practices.⁷ Virtualisation is common among EU enterprises, with 95% in our sample using this technique, even if only on a few servers. However, the average rate of virtualisation of workloads overall was 37% among EU respondents, compared to 48% of respondents in the 2019 US survey. There was little variation across EU countries but a higher virtualisation rate at larger firms.

Adopting virtualisation is a key step in boosting energy efficiency, but making the most of virtualisation via fewer and better-used systems is just as important. To understand how well servers are utilised, we asked about workload consolidation practices. Over a quarter of EU respondents said they are aggressively consolidating, targeting high levels of consolidation to reduce the number of servers required and to improve overall energy efficiency; a similar number proclaimed themselves pragmatic, or pushing for consolidation but careful not to impact application responsiveness. However, a large group claimed to use virtualisation for other purposes, such as simplified management of instances and higher availability (faster recovery times) rather than for reducing server count.

Figure 3: Workload Consolidation Practices



Source: 451 Research/S&P Global Market Intelligence (n=308)

These responses offer guidance on server utilisation. The model calculates that based on the 2021 sample, average server utilisation across surveyed EU enterprises is slightly above 15%. This compares to slightly below 15% in APAC and an estimated 18% in the 2019 US sample.

This level of utilisation is not surprising given that for a typical company, utilising business servers more than this is quite difficult, even if the systems are sized for particular workloads, because little processing occurs outside of business hours. By contrast, hyperscale providers utilise cloud servers at much higher levels. Our research suggests that cloud providers typically aim for much higher utilisation for shared (multi-tenant) hardware in a bid to strike the right balance between efficiency, cost savings and application responsiveness.

6. Server virtualisation is a collection of techniques that decouple running software from hardware via layers of abstraction. This creates flexibility in where the software runs and what hardware resource it accesses.

7. Note that our model performs detailed calculations for each enterprise to estimate their respective infrastructure efficiency. Other factors such as the resource sizing of servers (i.e., processors, memory, storage drives), as well as the broader storage and network infrastructure, also affect overall efficiency, but we opted not to include these because they are more nuanced and specific to workloads.

This is possible due to load balancing across hyperscale infrastructure – for example, by running low-priority computational jobs at times when there is spare capacity left by other, higher-priority customer applications. Thus, for example, the same servers that run video calls for businesses during the day could provide movies or videogame content at night. Cloud utilisation may edge even higher in the coming years. Based on vendor disclosures, 451 Research expects upcoming servers to have peak efficiency points near or beyond 80% utilisation. This, combined with the adoption of new modular software development techniques such as microservices and serverless computing (running many snippets of code scaling to meet demand as opposed to large monolithic software), will promote further improvements in cloud utilisation and, therefore, energy efficiency.

Another element influencing enterprise server efficiency is the change in the number of servers over time. On average, EU businesses in our survey have seen their server numbers grow by 22% over the past three years (in enterprise on-premises data centres). Server growth means that the installed base of servers includes a larger number of servers using more recent technology.

However, cloud providers such as AWS have the edge here as well. Owing to rapid uptake of cloud services (leading to double-digit growth for cloud providers), the cloud server population includes a larger portion of more recent server generations than the typical organisation's server population. This effect is further amplified by cloud providers' investments in new technology, such as AWS developing the Graviton processor series, and early purchases of the latest technologies. For example, 451 Research believes that the Intel server processors currently utilised in cloud infrastructure are heavily weighted toward 14nm processors (such as Skylake, Cascade Lake/Cooper Lake), based on sales patterns reported by Intel in its financial filings. This is coupled with a rapid uptake of more recent chips from both Intel (Ice Lake) and AMD (Rome/Milan) that deliver major leaps in energy efficiency and performance, particularly at high utilisation.

451 Research finds the root of these differences in a structural factor: organisational setup. In the cloud business model, inefficiency has a material impact on the bottom line. Thus, cloud providers' technical organisations take a more unified view of infrastructure, with IT and data centre infrastructure teams working together to optimise design and operations. In addition, cloud providers' scale enables the funding of custom engineering efforts and improvements.

A consistent theme from our studies around the world is that survey respondents in general do not have comprehensive IT sustainability programs that would raise all aspects of enterprise data centre operations in line with best practices. Only a handful of surveyed companies reported higher than average performance across multiple metrics. This means even the best-performing enterprises as a cohort fall well short of cloud efficiencies.

On balance, we consider our model conservative in estimating the energy-efficiency delta between enterprise and cloud servers due to a number of factors that favour enterprises – for example, we did not include the emphasis enterprises typically place on price and right-sizing servers for specific applications instead of overall server efficiency.

The Cloud Data Centre Facility Effect

Data centre cooling systems – such as chillers, pumps, air handlers and air-conditioning units that keep the temperature and humidity under control – are major energy users and contribute to inefficiencies at enterprise data centres throughout the world. While industry bodies such as the globally recognised American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) recommend widening temperature bands in the interest of energy conservation,⁸ most enterprises still aim to keep server inlet temperatures low (typically under 22°C) in the belief that this is necessary to keep IT component failures low (even though industry data indicates the concern is misplaced). This practice is highly energy-intensive because most enterprise data centres rely on mechanical refrigeration (compressors) to cool the circulated air. The climate across much of the EU is temperate; compared to some other regions of the world (such as Southeast Asia), it allows for the use of more energy-efficient cooling methods, but survey results indicate that these are not being utilised by EU enterprise data centres.

Overall, our survey paints an unfavourable picture of EU enterprise data centre energy consumption. The PUE on average across the more than 300 European firms surveyed was 2.01, well above the 2019 US reading of 1.66 and only slightly better than the average of 2.03 in APAC, where climates are typically hotter and more humid. This means that, on average, the energy European on-premises data centres use for cooling, power distribution and ancillary functions (lights, security systems, office space) is equal to the energy used by IT systems, chiefly servers.

This stands in contrast to the efficiency of cloud data centres, which are known for optimising energy efficiency across their operations. A significant factor in cloud data centre efficiency is a more relaxed cooling regime in which temperature and humidity move within wider bands. For example, AWS states that it has optimised its evaporative cooling systems to minimise water usage (when water is required during hotter temperatures), and AWS data centres in Ireland and Sweden use no water for 95% of the year. This results in far less energy used in cooling a cloud data centre in return for a relatively small amount of water. As a result, hyperscale cloud data centre sites across much of the EU can deliver an annualised PUE below 1.4 in most cases, and sometimes as low as 1.1.

Most enterprises would find it difficult to implement these practices in their data centres. Concerns about risk to IT hardware – such as increased component failures – are still common, and application downtime can prove very costly to the business. New hotspots (locations where temperature exceeds the limits set) may develop where cold air delivery to IT is insufficient. But even if that's not the case, without optimisation of airflow in the facility, and changing server settings with higher operating temperatures in mind, enterprise data centres' air handlers cannot attain more efficient performance levels. Furthermore, concerns about a potential increase in IT hardware component failure tend to deter enterprise technical leads from promoting such efficiency initiatives.

We also note that cooling, although a major energy user, is not the only factor that contributes to facility-related energy overhead enterprises. New UPS systems are considered to be 92-95% efficient (losing only 5-8% of power to operate the UPS), while older (5-10 years or more) models can be less than 90% efficient.⁹ The biggest hurdle for enterprises to overcome this is a combination of the high capital cost of electrical upgrades and the risks (perceived or real) to live operations.

8. Thermal Guidelines for Data Processing Environments, Fourth Edition, The American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2015

9. High Efficiency UPS Operating Modes, a presentation by Chuck Heller, Liebert Corporation, May 2011, and Implementing Energy Efficient Data Centers, Schneider Electric White Paper 114 Rev 1

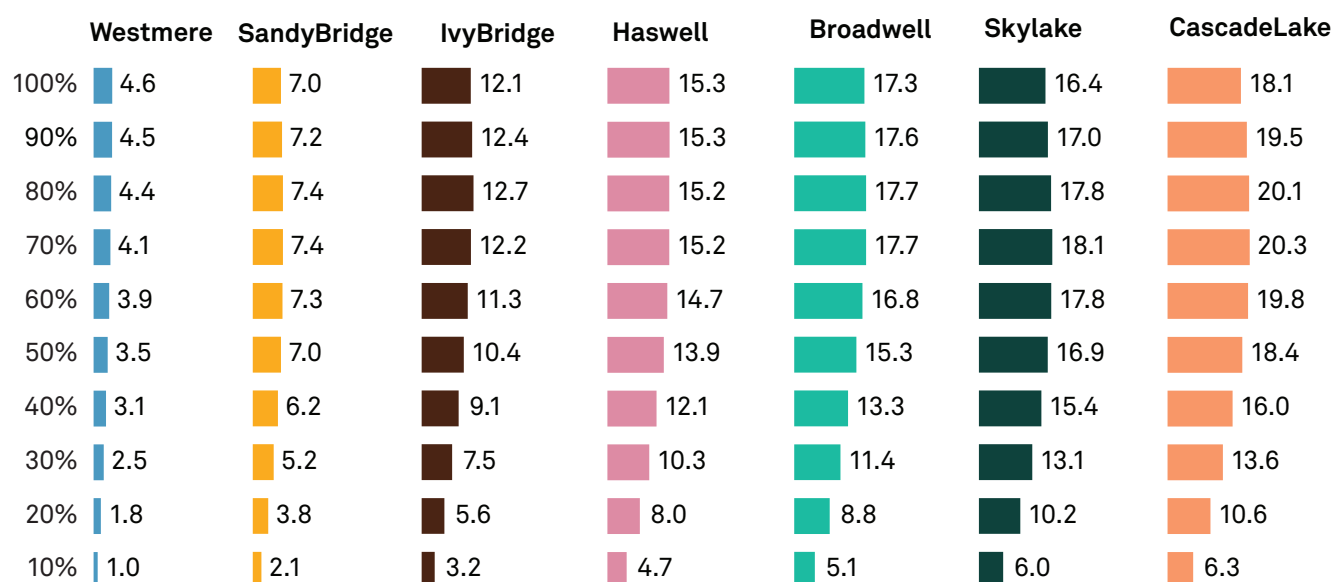
Specifics of the Model

Using the results of the survey, 451 Research calculated the relative operational efficiency at the largest data centres of the enterprises surveyed. To establish a baseline, we used the power-efficiency database of industry body Standard Performance Evaluation Council (SPEC). The benchmark, called SPECpower_ssj2008, simulates a business application to measure the power consumed per second by operations on various types of servers.

We used Intel-based servers from 2010 at 10% load (which is a typical load for non-virtualised servers) to create a reference point. In our model, the average energy efficiency of a two-processor system based on the 2010 Intel Xeon server processor at 10% load is '1.' Using the survey data, we modelled the profile and utilisation of typical enterprise servers to estimate their energy efficiency. The modelling included the adoption rate of new server technology, how often servers were replaced and how many new servers were added to the installed base over time. We also used survey responses on virtualisation and workload consolidation to estimate server utilisation rates.

We utilised this modelling because server hardware is not equally efficient across its load curve. Figure 4 below shows how efficiency can vary both by server generation (from left to right) and by utilisation rate (from bottom to top). Our model suggests that the rapid growth, scale and utilisation of cloud infrastructure such as that of AWS has greatly increased server efficiency compared to most enterprise infrastructure and will continue to do so in coming years.

Figure 4: Relative Efficiencies (the Processing Done Per Watt) for Different Types of Chipset at Different Utilisation Levels Compared to the Westmere Chip at 10% Utilisation



Source: SPEC.org, compiled by 451 Research

Energy Savings from Moving Enterprises Workloads to Cloud Infrastructure in the European Union

According to our calculations, surveyed EU enterprises scored, on average, 6.9 in server efficiency, meaning that their server infrastructure is over six times more efficient than non-virtualised servers at the start of the past decade (in 2010). This is due to advances in server technology – for example, faster yet lower-power semiconductors, in addition to virtualisation. However, the average hides a wide variation. The best-performing enterprises have achieved gains of over 9x in efficiency, while the worst are below 5x.

Although top-performing EU enterprises are more efficient than their US or APAC peers, the average European efficiency is below that seen in some other regions. We estimate that most EU enterprises have consolidated fewer workloads and virtualised less than US enterprises, for example, while also taking longer to adopt the latest server technology. When compared to the results of our 2019 US study, we estimate the IT energy efficiency gap between enterprises in the US and EU to be about 25%, with the advantage going to the US.

Using the same model, 451 Research estimates that hyperscale cloud servers, as demonstrated by AWS, are roughly three times more energy-efficient than those of average EU enterprises, with a score of over 20 compared to 6.9 in server efficiency. This is mainly because cloud servers are much more highly utilised and a larger percentage of them are new, using technology that is more energy-efficient. In addition, cloud providers such as AWS have more energy-efficient data centres, as detailed previously, with better power usage effectiveness than typical on-premises enterprise facilities. When including facility-level efficiency, *AWS cloud infrastructure is five times more energy-efficient than typical EU enterprise infrastructure*. We estimate that moving from on-premises enterprise hardware and facilities to the cloud could reduce average enterprise energy consumption by over 80%. In addition, cloud providers, including AWS, have enabled significant volumes of new renewable energy projects to power their operations and, therefore, have a much lower carbon impact than most enterprise facilities.

Enterprise Efficiency Approach

European enterprises are mindful of the need to improve efficiency and the expectation of stepping up their enterprise data centre sustainability efforts. Yet a gulf remains between good intentions and effective action. While there is clear government pressure for businesses to strive for sustainability, and European respondents typically have PUE targets for their data centres, they find it difficult to make the investments in technology upgrades that deliver such gains. For example, a head of data centres at a midsize Swedish retailer observed, “[at the corporate procurement department,] they are not willing to look at the energy-efficiency factor or the ROI; their main focus is always on one thing, [...] if they can make do without this spending, then they do not proceed further on this.”


In addition, sustainability is typically gained as part of a cost-reduction effort, rather than being a specific goal. A technical lead for data centre IT operations at a French telecommunications provider said, “There are no particular goals for energy reduction or energy savings, but [rather] for cost saving; that’s what happens [at] most of the companies, and it somehow ends up saving energy as well.”

Based on this and other surveys, we believe most enterprises do not see data centre infrastructure sustainability as a strategic priority. Thus, the technical organisations running enterprise data centres and IT infrastructure have efficiency programs that, for all the benefits they deliver, are limited in scope and must work within the confines of existing data centre and IT infrastructure.

Outlook and Conclusions – Directions in Energy Efficiency and Energy Policy for Carbon Reduction

Based on our survey of over 300 businesses across France, Germany, Ireland, Spain and Sweden, as well as our estimates of server populations across European enterprises, moving business applications to the cloud could reduce IT and facility energy use by 80% on average in the EU compared to operating in an enterprise data centre (see Figure 5).

Figure 5: Energy Reduction Potential for a 1MW Enterprise Data Centre Load Moved to a Cloud Provider such as AWS




This could reduce carbon emissions substantially due to the use of fossil fuels such as coal and natural gas to generate electricity in many EU countries. Although the fuel used to generate power in the EU varies greatly from country to country, the highly interconnected energy networks create what can be thought of as one grid. According to the European Environmental Agency, the European electrical grid's overall emissions were 255g of CO_{2eq} per kilowatt-hour.

Reducing the energy required for IT by 80% on average in the EU would mean that for a 1MW (about 1,000-square-metre) enterprise data centre at an assumed 30% electrical utilisation, switching all applications to cloud services could reduce emissions by up to 897 metric tonnes of CO_{2eq} per year from IT equipment and up to 182 metric tonnes of CO_{2eq} per year from facilities for a total of 1,079 metric tonnes of CO_{2eq} per year on average. This is equivalent to removing more than 500 cars from the roads or offsetting the annual electricity related emissions for more than 50 households across the EU on average. When we include both the carbon intensity of consumed electricity and renewable energy purchases, which reduce associated carbon emissions, when AWS reaches 100% renewable energy, it will perform the same task with a 96% lower carbon footprint on average in the EU.

We estimate that reductions are achieved by shifting to cloud. This is because the overall efficiency of cloud infrastructure makes it easier and more economical to run a given workload using low-carbon energy when the energy required is a fraction of that needed for typical enterprise infrastructure. This reduction becomes even greater when factoring in cloud providers' potential use of renewable and low-carbon energy sources. AWS, for example, has committed to using 100% renewable energy across its global operations, and is on the path to achieve this by 2025. This is not only because cloud providers have the organisational resources, long-term view and large-scale demand for electricity to encourage new renewable energy generation projects. It is also because organisations such as AWS are making efforts to reduce emissions by prioritising the use of low-carbon energy and even signing power purchasing agreements, despite the complexities involved.

Migrating enterprise workloads in the EU to an AWS cloud using 100% renewable energy could save an additional 214 metric tonnes of CO_{2eq} per 1MW at 30% utilisation. This is on top of the sizeable CO_{2eq} reduction due to the improved efficiency of servers and facilities. If thousands of enterprise data centres were taken offline across the EU, this would significantly reduce greenhouse gas emissions, equalling the footprint of millions of households.

Country Profile: France




French survey respondents had the most energy-efficient data centres out of the EU countries surveyed, with an average PUE of 1.98. This was offset somewhat by a relatively long server life span of 51 months on average (compared to 46 months in Germany) and an estimated server utilisation rate of 15.4%.

Due to what we calculate to be relatively old servers with low utilisation, typical French enterprises could potentially expect to see energy savings of 67% by moving workloads to cloud, with an additional energy savings of 11% from cloud data centre efficiency, for a total of 78% energy savings. France has among the lowest greenhouse gas emissions from electricity production, thanks to its nuclear power plants and renewables, so the carbon emissions saved by improving energy efficiency is relatively low – at 192 metric tonnes of CO_{2eq} per megawatt annually from servers and another 30 metric tonnes of CO_{2eq} per megawatt annually from more efficient data centre facilities. When a cloud provider is powered by 100% renewable energy for its infrastructure in France, that could reduce emissions for enterprises that do not have 100% renewable energy by another 11 metric tonnes of CO_{2eq} per megawatt annually, for a total of 233 metric tonnes.

France has over 3,800 businesses with more than 250 employees. If 25% of these firms put 1MW of IT load into the cloud (a cloud migration project of moderate size), this could save as much as 226,450 metric tonnes of CO_{2eq} per year, equivalent to roughly a year's worth of emissions from nearly 22,000 French households' electricity use.¹⁰ As the country's ageing nuclear plants are taken offline over the next decade, this energy savings could become a key contributor in keeping the country's overall carbon output low.

10. Sources: Eurostat: "Number of enterprises in the non-financial business economy by size class of employment online data code: TIN00145 last update: 18/03/2021", IEA Atlas of Energy (data from 2018)

Country Profile: Germany




In Germany, survey respondents reported the shortest server life span among countries surveyed, at 46 months on average, which yields better energy efficiency. This was slightly offset by having only 39% of servers virtualised, which was below the number in Spain and France, and an estimated server utilisation rate of 15.6%. However, German survey takers had the second most efficient data centre facilities, after France, with an average PUE of 1.99. We calculate a potential energy savings of 65% on average from server infrastructure and another 13.5% on average from data centre infrastructure, for a total of 78.5% energy savings when a typical German enterprise moves workloads to the cloud.

Although Germany has been promoting wind and solar installations, the country still has the highest quantity of fossil fuels in its energy production mix of the countries we surveyed. Thus, the potential carbon emissions saved by moving workloads to the cloud could be as much as 1,159 metric tonnes of CO_{2eq} per megawatt annually from servers and another 242 metric tonnes of CO_{2eq} per megawatt annually from more efficient data centre facilities. When a cloud provider is powered by 100% renewable energy for its infrastructure in Germany, that could further reduce emissions for enterprises by 333 metric tonnes of CO_{2eq} per megawatt annually, for a total of 1,734 metric tonnes. Germany has roughly 11,300 businesses with more than 250 employees. If 25% of these firms put 1MW of IT load into the cloud (a cloud-migration project of moderate size), this could save as much as 4,903,600 metric tonnes of CO_{2eq} per year, equivalent to roughly a year's worth of emissions from more than 292,000 German households' electricity use.¹¹

11. Sources: Eurostat: "Number of enterprises in the non-financial business economy by size class of employment online data code: TIN00145 last update: 18/03/2021", IEA Atlas of Energy (data from 2018)

Country Profile: Ireland




Irish respondents to the survey had an average server life span of 52 months, and 35% of servers were virtualised on average, with an estimated utilisation rate of 14.6%. This was right in the middle of the group of countries studied; the same held true with data centre efficiency, at an average PUE of 2.06. Irish companies could potentially save 67.5% of energy on average from server infrastructure and another 15.6% of energy on average from data centre infrastructure by moving workloads to the cloud, for a total of 83.1% energy savings.

Ireland had the second-highest quantity of fossil fuels in its energy production mix, after Germany. This means that the potential carbon emissions saved by moving workloads to the cloud could be as much as 1,135 metric tonnes of CO_{2eq} per megawatt annually from servers and another 263 metric tonnes of CO_{2eq} per megawatt annually from more efficient data centre facilities. When a cloud provider is powered by 100% renewable energy for its infrastructure in Ireland, that will further reduce emissions for enterprises by 230 metric tonnes of CO_{2eq} per megawatt annually, for a total of 1,628 metric tonnes. Ireland has roughly 600 businesses with more than 250 employees. If 25% of these firms put 1MW of IT load into the cloud (a cloud-migration project of moderate size), this could save as much as 241,000 metric tonnes of CO_{2eq} per year, equivalent to roughly a year’s worth of emissions from around 12,700 Irish households’ electricity use.¹²

12. Sources: Eurostat: “Number of enterprises in the non-financial business economy by size class of employment online data code: TIN00145 last update: 18/03/2021”, IEA Atlas of Energy (data from 2018)

Country Profile: Spain




In Spain, survey respondents reported the longest server life span among countries surveyed, at 54 months on average. This was slightly offset by having the highest rate of server virtualisation at 44% of servers virtualised and the highest estimated server utilisation rate at 16.1%, which lead to improved energy efficiency. However, Spanish survey takers had the least efficient data centre facilities, most likely due to Spain having the hottest climate of all the countries surveyed, with a PUE of 2.13. We calculate a potential energy savings of 66% from server infrastructure on average and another 13.5% on average from data centre infrastructure for typical Spanish enterprises that move workloads to the cloud, for a total of 79.5% energy savings.

Spain has been working to add renewables to its energy production mix and currently has nuclear generation, which helps reduce the carbon emissions from generation. For an enterprise, the potential carbon emissions saved by moving workloads to the cloud could average 772 metric tonnes of CO_{2eq} per megawatt annually from servers and another 146 metric tonnes of CO_{2eq} per megawatt annually from more efficient data centre facilities. When a cloud provider is powered by 100% renewable energy for its infrastructure in the country, that will further reduce emissions for enterprises by 187 metric tonnes of CO_{2eq} per megawatt annually, for a total of 1,105 metric tonnes. Spain has roughly 3,500 businesses with more than 250 employees.¹³ If 25% of these firms put 1MW of IT load into the cloud (a cloud-migration project of moderate size), this could save as much as 955,160 metric tonnes of CO_{2eq} per year, equivalent to roughly a year's worth of emissions from over 72,000 Spanish households' electricity use. This type of reduction in energy use could become more important as the country looks to phase out its nuclear power generation.

13. Sources: Eurostat: "Number of enterprises in the non-financial business economy by size class of employment online data code: TIN00145 last update: 18/03/2021", IEA Atlas of Energy (data from 2018)

Country Profile: Sweden



In Sweden, respondents to the survey had the second-lowest server life span, behind only Germany, with an average of 49 months. Sweden ranked third in server virtualisation percentage, at 39.8%, and enterprises had the second-best estimated utilisation rate of 14.8%. Given the climate, it was somewhat surprising that Swedish data centres were not as efficient as those in Germany and France, with an average PUE of 2.1. This could be because the data centres are smaller, or perhaps there is less incentive to make data centres efficient because power in Sweden is relatively inexpensive and is generated with very low emissions. Swedish companies could potentially save 67.5% of energy on average from server infrastructure and another 15.4% of energy on average from data centre infrastructure by moving workloads to the cloud, for a total of 79.9% energy savings.

Sweden generates nearly all of its energy from low-carbon sources so, as in France, the potential carbon savings from moving workloads to cloud is relatively low, averaging only 44 metric tonnes of CO_{2eq} per megawatt annually from servers and another 10 metric tonnes of CO_{2eq} per megawatt annually from more efficient data centre facilities, for a total of 54 metric tonnes. Sweden has roughly 1,400 businesses with more than 250 employees.¹⁴ If 25% of these firms put 1MW of IT load into the cloud (a cloud-migration project of moderate size), this could save roughly 22,570 metric tonnes of CO_{2eq} per year, equivalent to roughly a year's worth of emissions from over from 4,400 Swedish households' electricity use, assuming those households and data centres are not using 100% carbon-free energy already.

14. Sources: Eurostat: "Number of enterprises in the non-financial business economy by size class of employment online data code: TIN00145 last update: 18/03/2021", IEA Atlas of Energy (data from 2018)

Appendix

Methodology Details

The focus of this model is to capture the carbon impact of key design and operational features of enterprise data centres compared to the hyperscale cloud, in particular AWS, and to understand how key components impact the overall efficiency picture. The core of the model focuses on Scope 2 emissions, or utility grid electricity. The model does not include Scope 1 emissions or direct emissions from site operations, such as cooling system refrigerants or diesel engine generator emissions, nor does it look at Scope 3 emissions or those embodied in buildings and IT products, such as concrete.

We decided to focus on Scope 2 emissions because enterprises can control some of the key factors that influence this type of energy consumption, and these have a substantial impact on data centre energy efficiency and carbon footprint. In our view, Scope 1 emissions are important but do not currently reflect the potential operational efficiency of particular data centres because this is typically a very small part of the carbon footprint from data centres. For example, virtually all data centres currently use diesel generators to provide backup power if the grid fails, but this is not a common occurrence. This is evolving, however, and we will study the potential for improvements in this area in the months and years to come. Future versions of this model may include Scope 3 emissions. However, currently we do not expect these to meaningfully alter the conclusion of the analysis because we believe they are not considerably different between cloud and enterprise data centres.

EXTENDED CARBON MODEL w/ embedded emissions

Embodied facility	+	Embodied IT	+	451 Research CLOUD ENERGY CARBON MODEL				
				Grid - offset by renewable purchases	X	PUE	X	IT efficiency

The carbon emissions model analyses three major areas: IT equipment energy use; data centre efficiency; and the carbon intensity of the grid. The objective of the model is to show the potential difference between enterprise and cloud operations.

The output of the model is a ratio that shows the relative energy and carbon-efficiency difference between enterprise IT and cloud, as represented by AWS. The model incorporates EU survey data, data from S&P Global Market Intelligence on EU data centre and energy markets, third-party industry sources, and data from AWS.

The survey of EU enterprises with \$10m-1bn in annual revenue sought to understand characteristics that influence efficiency metrics, including policies and attitudes toward consolidation, adoption of new server technology, and server replacement cycles. We believe that such an approach creates a more robust picture and provides better context than asking exclusively for technical specifications and operational metrics, many of which may not be available to survey respondents in required detail.

Grid carbon intensity – Carbon emissions per kWh energy; S&P Global Market Intelligence and third-party data.

PUE – Power usage effectiveness shows the facility energy overhead as a ratio of the IT load.

Server hardware power efficiency – The inherent design power efficiency of the server that is calculated using *server distribution by age, server utilisation and power efficiency* data from the Standard Performance Evaluation Council's database `specpower_ssj2008`.

- Server age distribution: proxy for server technology generation that largely defines the server's efficiency potential. To gain this distribution, we asked for average life span, speed of adoption of new tech (to account for additional lag compared to the cloud) and capacity change (skew of distribution). Q10: *How long is the typical life span of the servers running in your largest data centre?*
- Server utilisation: Instead of asking for server utilisation, which we find is not practical to obtain, we asked about the maturity of IT operations by gauging virtualisation levels, any projects in motion that aim to increase virtualisation levels over time, and aggressiveness of consolidation. 451 Research based assumptions on how these responses translate into utilisation numbers on third-party industry data.
- Power efficiency data from SPEC: SPEC maintains a database on server power efficiency per a test suite that simulates a complex business logic and benchmarks performance against power use across the load curve. Using this data, the model can assess the relative power efficiency of servers based on their technology generation (age) and utilisation.

While server makers aggressively fine-tune hardware and software specifically for the specpower_ssj2008 benchmark to attain the best possible result in ways that arguably do not represent a typical deployment case, we are relying on averages across multiple submissions and use the data to calculate efficiency improvements with newer server generations and with better utilisation. We believe these are representative of real-world behaviour of hardware and software in a generic enterprise IT environment.

While SPEC data is in 10% increments, we needed finer granularity of 1% for our calculations because we modelled IT operational efficiency of the surveyed US enterprises. We did that by using linear interpolation between measurements as an approximation to an implied efficiency curve.

- Based on virtualisation and consolidation levels, we calculated composite average utilisation of each server generation for each enterprise, then weighted such efficiency readings by distribution of server generations per enterprise. We tested this against a more detailed hourly workload simulation (e.g., internet traffic profile during a day) where a more complex calculation using hourly utilisation and energy-efficiency readings would generate the efficiency reading, but the total difference from using a simple average utilisation and associated energy efficiency reading was typically 1% or less.

To estimate the equivalent carbon emissions of average households that could be saved by businesses switching to cloud, we started by calculating the emissions per household using the IEA Atlas of Energy (data from 2018) estimates of CO_{2eq} emissions per person in each country, multiplied by the average number of people per household in each country (2020 Eurostat data). We used the Eurostat: "Number of enterprises in the non-financial business economy by size class of employment (online data code: TIN00145 last update: 18/03/2021)" report to conservatively estimate the number of businesses by country and took 25% of that number, multiplying it by the estimated amount of CO_{2eq} metric tonnes saved. The amount of CO_{2eq} saved is based on our model's outputs for moving 1MW to cloud in each country.

CONTACTS

The Americas

+1 877 863 1306

market.intelligence@spglobal.com

Europe, Middle East & Africa

+44 20 7176 1234

market.intelligence@spglobal.com

Asia-Pacific

+852 2533 3565

market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

Copyright © 2021 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not endorse companies, technologies, products, services, or solutions.

S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.