

451 Research

S&P Global

Market Intelligence **Black & White**

The Carbon Reduction Opportunity of Moving to the Cloud for APAC

COMMISSIONED BY



JULY 2021

©COPYRIGHT 2021 S&P GLOBAL
MARKET INTELLIGENCE.

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.



451 Research

S&P Global
Market Intelligence

Table of Contents

- Introduction 4
- Executive Summary 5
 - Figure 1: Carbon Reduction Potential of Cloud Infrastructure Compared with Surveyed APAC Enterprises and Public Sector Organisations 6
- Key Survey Findings 7
 - Figure 2: What Industry Do You Work in? 7
 - Enterprise and Public Sector Data Centre Infrastructure in the Asia-Pacific Region – Survey Overview 8
 - The Cloud Data Centre Server Effect – Structural Efficiency. 8
 - Figure 3: Workload Consolidation Practices by Country 10
 - The Cloud Data Centre Facility Effect – APAC Enterprises and Public Sector Feel the Heat 11
- The 451 Research Cloud Energy Efficiency Model 13
 - Figure 4: Efficiency Curves of Intel Server Generations per SPECpower_ssj2008 Database 14
- Energy Savings from Moving Enterprise and Public Sector Workloads to Cloud Infrastructure in APAC 15
- Outlook and Conclusions – Directions in Energy Efficiency and Energy Policy for Carbon Reduction 16
 - Figure 5: APAC Cloud Carbon Reduction Potential 16
 - Figure 6: APAC Cloud Carbon Reduction Potential at 100% Renewable Corporate PPA Mix 17
- Appendix: Country Profiles – Some Differences, but Similar Conclusions 19
 - Country Profile: Japan 19
 - Country Profile: South Korea 20
 - Country Profile: Singapore 21
 - Country Profile: Australia 22
 - Country Profile: India 23
- Methodology 24

Introduction

Cloud infrastructure¹ has become a key element of enterprise and public sector IT, providing flexible resources and enabling workloads to scale as needed, without having to buy and install new hardware.² However, there is an additional potential benefit that many enterprises and public sector organisations are not aware of: moving workloads to the cloud can also dramatically reduce their energy use.

IT efficiency and sustainability are becoming increasingly critical considerations and priorities for businesses and governments around the world. KPMG noted in a global survey that 80% of companies it works with publish sustainability information, and a majority have set targets to lower carbon emissions. However, the current targets may not be enough.³ A United Nations report from mid-2020 found that only about 40% of enterprises it surveyed are convinced that their targets are ambitious enough to meet the UN's Sustainable Development Goals by 2030.⁴

A major source of energy use and carbon emissions for enterprises and public sector organisations is self-owned and operated IT equipment and data centres. Yet improving their efficiency can be difficult and expensive. Many would like to boost efficiency and lower emissions, but do not have the resources to make dramatic changes to their infrastructure that would optimise efficiency. This challenge can be compounded by changing weather patterns and risks related to climate change, with enterprises and public sector organisations sometimes having to choose whether to invest in disaster-proofing their data centres or invest in improving efficiency. Putting workloads in the cloud assists with both of these challenges.

451 Research, part of S&P Global Market Intelligence, set out to better understand the energy used by enterprise and public sector IT infrastructure across the Asia-Pacific (APAC) region, and to compare it with that of typical cloud infrastructure. This process began with a survey of senior stakeholders at over 500 organisations from a variety of industries in Japan, South Korea, Singapore, Australia, and India—including over 100 public sector organisations and enterprises with annual revenue between USD \$10m and \$1bn in each country.

These survey results, along with analysis and third-party industry data, were used to model a 'chip to grid' view of efficiency that estimated the energy used by typical workloads in on-premises (self-owned and operated) servers and data centres on average. This was compared with an estimate of energy used by hyperscale⁵ cloud facilities in order to understand whether enterprises and public sector organisations across the APAC region could reduce energy consumption — and their related carbon footprint — by moving workloads from on-premises data centre infrastructure to the cloud.

-
1. Throughout this report, we use the word 'cloud' to refer to cloud computing, the on-demand delivery of compute power, database, storage, applications, and other IT resources through a cloud services platform via the Internet with pay-as-you-go pricing.
 2. Voice of the Enterprise: Cloud, Hosting and Managed Services, Workloads and Key Projects - Quarterly Advisory Report, 451 Research, June 28, 2019
 3. The time has come: The KPMG Survey of Sustainability Report 2020 KPMG IMPACT, December 2020.
 4. United Nations Global Compact, "Uniting Business in the Decade of Action", July 2020.
 5. "Hyperscale" as defined by the United States Data Center Energy Usage Report (LBNL-1005775) refers to both data centre scale ("warehouse-sized") and scalability (a large number of servers networked together with scalable server architectures and virtual networking).

BLACK & WHITE | THE CARBON REDUCTION OPPORTUNITY OF MOVING TO THE CLOUD FOR APAC

451 Research

S&P Global
Market Intelligence

Executive Summary

Moving APAC enterprise workloads to the cloud from on-premises data centres can reduce energy consumption and associated carbon emissions by nearly 80%.

This study estimates the potential reduction in energy use and carbon emissions that could be achieved by moving business applications from on-premises enterprise and public sector data centres to hyperscale cloud services in APAC. Building on modelling from a similar study performed in the U.S. in 2019,⁶ the findings of this APAC study demonstrate an average of nearly 80% energy savings from running business applications in the cloud rather than on on-premises infrastructure.

Hyperscale cloud infrastructure, at both the server and facility level, is significantly more efficient than that of most enterprises and public sector organisations, which translates into considerably less energy used to perform the same unit of work. On the server side, hyperscale cloud operators design server systems with great attention paid to power optimisation, using the very latest components. These cloud operators then run their servers to higher levels of utilisation, leveraging the ability to share and dynamically allocate resources to serve customers' workloads on the cloud. Hyperscale cloud providers also design their data centre sites to be highly efficient, using less energy for both cooling and power distribution. This drive for efficiency begins with the cloud business model, which delivers IT services at scale and which incentivises cloud providers to increase efficiency, from design to operation.

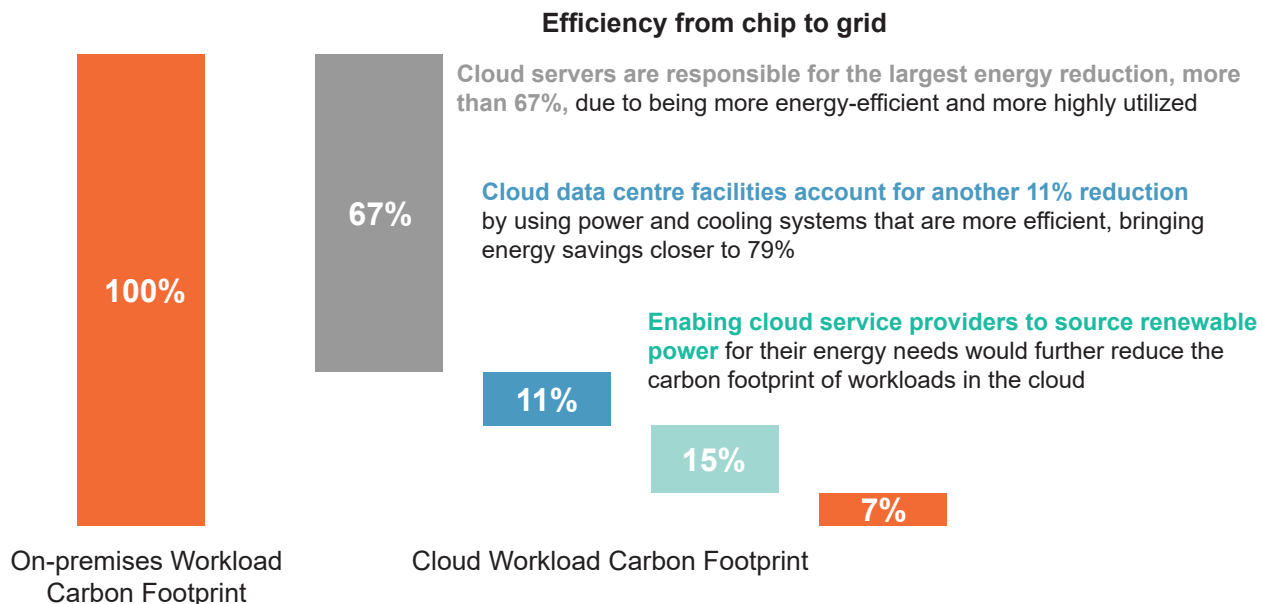
Hyperscale cloud data centres perform at an estimated five times the energy efficiency of the average APAC enterprise surveyed.

As a result of these efficiency advantages (and despite variation in enterprise and public sector data centre infrastructure conditions across the markets that we surveyed), the results of our APAC survey and modelling consistently show significant energy savings potential from a move to the cloud in APAC. We estimate that hyperscale cloud data centres perform at nearly five times the energy efficiency of the average of surveyed APAC enterprises and public sector organisations.

6. The Carbon Reduction Opportunity of Moving to Amazon Web Services, 451 Research, October 2019

Figure 1: Carbon Reduction Potential of Cloud Infrastructure Compared With Surveyed APAC Enterprises and Public Sector Organisations

Source: 451 Research/S&P Global Market Intelligence



What really sets the cloud gains apart in APAC markets, however, are the reductions in associated greenhouse gas (GHG) emissions. While grid emissions differ greatly from country to country, the APAC average of 576.5g CO₂-equivalent per kilowatt-hour is already well above that of the U.S. (417.3g in 2019). Consequently, shifting enterprise and public sector workloads to cloud infrastructure and reducing the energy footprint per workload yields even greater reduction in carbon emissions in APAC than in the U.S.

Combined with server- and facility-level efficiency gains, access to 100% renewable power for hyperscale cloud infrastructure can yield up to 93% reduction in carbon emissions.

Enabling cloud service providers to source 100% renewable power for their operations – as the major hyperscale cloud operators have pledged to do – would reduce the carbon footprint of cloud workloads even further. The current lack of accessible and affordable renewable energy options in much of APAC (such as corporate renewable power purchase agreements, or PPAs), leaves a significant amount of carbon reduction potential on the table. Even though each country’s grid is expected to use more renewable energy over time, boosting the proportion of renewable energy used by enterprises and public sector organisations, hyperscale cloud providers have more aggressive renewable energy goals, and the scale of hyperscale investments in renewables means that each new project adds tens of megawatts of new renewable power to the grid. We calculate that enabling cloud operators to source 100% renewable energy would reduce related emissions by another 15% (see Figure 1), which, when combined with the server and facility efficiency gains already discussed, could reduce IT infrastructure emissions by up to 93% across APAC.


BLACK & WHITE | THE CARBON REDUCTION OPPORTUNITY OF MOVING TO THE CLOUD FOR APAC

Key Survey Findings

The 2021 survey of 515 APAC enterprises and public sector organisations across Japan, South Korea, Singapore, Australia, and India shows that in APAC (as in the U.S.), when it comes to enterprise and public sector data centre sustainability efforts, there is often a gulf between intentions and effective action. All businesses and governments interviewed in our 2021 APAC study track energy consumption and consider efficiency when purchasing IT and facilities equipment. Some also monitor key metrics such as water use, free-cooling hours and server utilisation, and/or hire external consultants to benchmark data centre sustainability against peers. Only a handful of organisations have formed dedicated teams to promote best practices or have put in place more comprehensive programs.⁷

Figure 2: What Industry Do You Work in?

Source: 451 Research/S&P Global Market Intelligence (n=515)



Although many APAC survey respondents reported that energy efficiency and sustainability were high on their agenda, their answers revealed goals for infrastructure that generally lacked definition and clarity. Few reported having incentives or sufficient resources in place to tackle these technically complex problems. As one IT director at a manufacturing company in Singapore observed, while the company does indeed want to reduce its carbon footprint and is working in that direction (e.g., through better equipment procurement and optimising operations), at an organisational level, “There are no goals.”

7. These include airflow optimisation, decommissioning server systems running but not doing any useful work, catching unwanted deviations from expected readings, and potentially adding some on-site renewable energy capacity to cover for ancillary needs (e.g., adding solar panels to run lights in common space).

BLACK & WHITE | THE CARBON REDUCTION OPPORTUNITY OF MOVING TO THE CLOUD FOR APAC

Typically, enterprise and public sector efforts to improve efficiency tend to be limited to working within the confines of existing data centre and IT infrastructure, rather than fundamentally overhauling operational infrastructure and procedures. As one data centre manager at an Indian transportation company put it, “Budget constraints keep companies from procuring the best thing possible, considering how expensive, difficult, and complex restructuring the infrastructure is.”

Enterprise and Public Sector Data Centre Infrastructure in the Asia-Pacific Region – Survey Overview

The Cloud Data Centre Server Effect – Structural Efficiency

The survey began by asking APAC respondents a series of questions about their server practices. These were used to create profiles of enterprise and public sector servers across APAC that focused on two main measures of server efficiency: the innate efficiency of server hardware and server utilisation. Innate efficiency is largely determined by chip technology and system design, while the utilisation rate, or load factor, of a server determines how close it is to its optimal energy efficiency on average.

To understand innate efficiency, our survey first asked about the APAC enterprise and public sector server population and factors such as hardware age, speed of adoption, and change in server numbers over time. To understand server utilisation, we mapped out virtualisation⁸ patterns and workload consolidation practices.⁹ We consider our model to be conservative in estimating the energy efficiency delta between enterprise/public sector and cloud servers.

Our results show that, on average, APAC respondents keep their servers for less than four years before upgrading – a considerably shorter time than the nearly four-and-a-half years reported by U.S. enterprises in 2019. This pattern is not strongly correlated with company size or industry vertical, but reveals some regional differences. According to the survey, Japanese organisations, for example, are more likely to hold onto their server equipment for longer (slightly over four years), whereas Australian organisations refresh their server inventory at a significantly faster pace of three years and three months on average.

-
8. Server virtualisation decouples software from hardware by using layers of abstraction to create flexibility around where software runs and what hardware it uses.
9. Note that our model performs detailed calculations for each organisation to estimate its infrastructure efficiency. Other factors impacting servers such as processors and memory, as well as broader network and storage infrastructure, also affect overall efficiency, but we opted not to include these for they are more nuanced and specific to workloads.

BLACK & WHITE | THE CARBON REDUCTION OPPORTUNITY OF MOVING TO THE CLOUD FOR APAC

451 Research

S&P Global
Market Intelligence

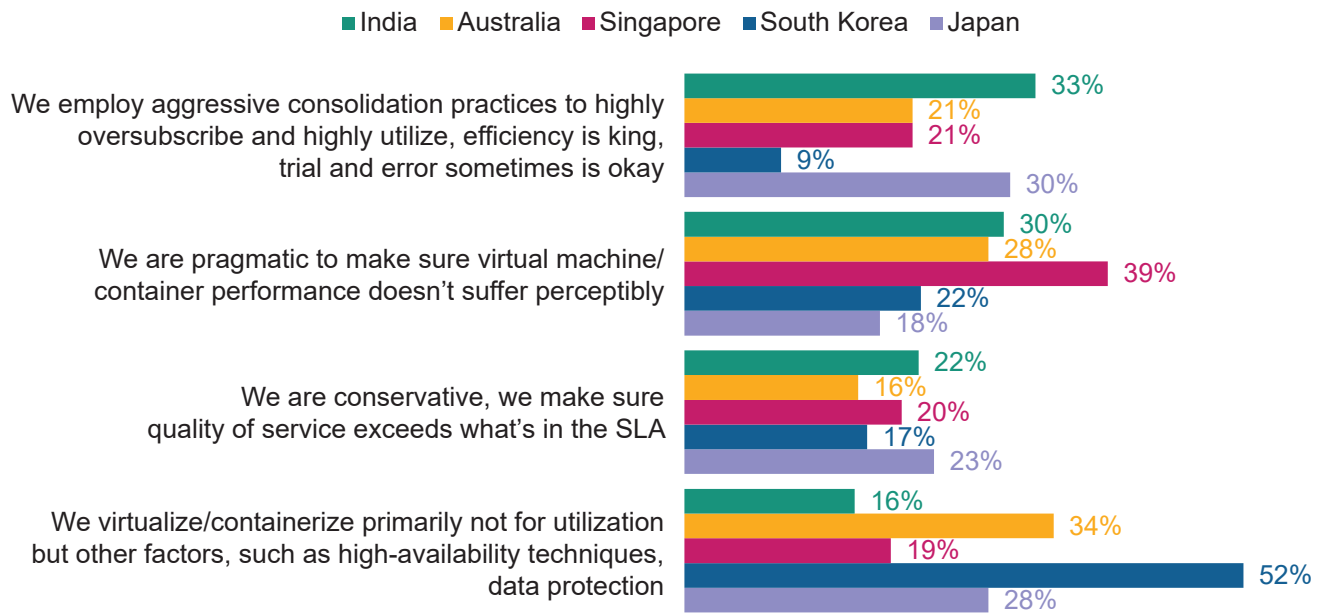
In addition to the server lifespan question, we asked how quickly new server technology is adopted. Unsurprisingly, larger organisations in the survey tended to be earlier adopters of new server technology than smaller ones. We also saw regional variations, with Japanese organisations taking nine months on average before buying the latest available server technology, compared with the six months average time frame of South Korean organisations. Adopting newer, more efficient servers at a faster pace has a material impact on energy efficiency. But even the fastest rate of adoption among APAC respondents could not keep up with that of hyperscale cloud service providers. Cloud operators have early access to the latest server technology – as much as a year ahead of general availability to enterprises or public sector organisations. This, combined with the relatively late enterprise/public sector adoption of new technology once it is available, adds up to a significant server-level efficiency gap between enterprises/public sector organisations and cloud providers over time.

Our survey showed that some level of virtualisation is common among APAC enterprises and public sector organisations. All survey respondents reported using this technique, even if only on a few servers. However, the average rate of virtualisation in APAC was 36%, with variation across countries ranging from 33-39%. This relative lack of virtualisation typically means servers are not as fully utilised and are therefore less efficient, contributing to the energy efficiency gap between on-premises data centres and the cloud.

To better understand server utilisation rates, we included questions about workload consolidation at APAC enterprises and public sector organisations (see Figure 3). APAC respondents reported lower consolidation of workloads per server than their U.S. peers, with South Korean organisations being the most cautious about allocating large numbers of workloads to a server (perhaps due to concern this would compromise application performance), even if this meant operating the servers less efficiently.

Figure 3: Workload Consolidation Practices by Country

Source: 451 Research/S&P Global Market Intelligence (n=515)



Server utilisation in APAC enterprises is estimated to be just under 15%, compared with well above 50% in hyperscale cloud servers.

These responses give us guidance on average server utilisation in enterprises and public sector organisations, which in the 2021 APAC sample is estimated at just under 15% (compared with an estimated 18% in the 2019 U.S. sample). These levels of utilisation are not surprising given that, for a typical organisation, utilising their servers much more than this is difficult since little processing is required outside of primary business hours. By contrast, our research shows that hyperscale operators utilise cloud servers at far higher levels, typically aiming well above 50% utilisation to find the right balance between efficiency and application performance.

Hyperscale cloud operators achieve higher utilisation rates by dynamically allocating server capacity among many customer workloads – for example, by running some computational jobs at times when other customers’ applications no longer need the server capacity. Cloud utilisation will likely climb even higher in the coming years. Based on vendor interviews, 451 Research/S&P Global Market Intelligence expects upcoming servers powered by Intel and AMD processors to have their peak efficiency points near or beyond 80% utilisation, which, combined with the adoption of new modular software development techniques such as microservices and serverless computing, will promote higher utilisation in order to achieve optimal efficiency.

The final element to understanding enterprise server infrastructures is the change in the number of servers over time. Most APAC respondents in our survey have seen their server numbers grow in the past three years, at least in their core data centres. Understanding this dynamic is important, because growth means the distribution of servers is skewed toward newer generations. Notably, survey respondents on average expected continued growth for at least the next three years. South Korean and Indian organisations expect the largest expansion (by over 20% in the next three years), while those in Singapore and Australia foresee a more moderate expansion of 15%-16%.

BLACK & WHITE | THE CARBON REDUCTION OPPORTUNITY OF MOVING TO THE CLOUD FOR APAC

Although adding servers will boost the percentage of new servers in an organisation's total server stock, cloud infrastructure still retains an edge. Owing to the very rapid growth in cloud services, along with early adoption of the latest technology and rapid server refresh cycles, we estimate that hyperscale cloud providers still will have a larger percentage of their overall server stock made up of the latest server technology than most enterprises or public sector organisations will.

The Cloud Data Centre Facility Effect – APAC Enterprises and Public Sector Feel the Heat

There are also efficiency differences when it comes to data centre facilities, particularly around cooling. The challenge begins with temperature. Major industry bodies, such as the globally recognised American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE), recommend widening the temperature range of air used to cool servers, in the interest of energy conservation.¹⁰ However, most enterprise and public sector organisations still aim to keep server air inlet temperatures low (typically under 22°C), due to worries that hot spots could develop where cold air delivery is insufficient, and concerns that higher temperatures will lead to IT component failures. In addition, when the IT inlet temperature rises, server fans may run more, offsetting some of the savings gained by running the building's cooling system less.¹¹

Most on-premises data centres also rely on mechanical refrigeration (compressors) to cool the equipment. By contrast, cloud data centres primarily rely on evaporative cooling rather than compressors, which uses less energy. This results in significantly less energy spent on cooling the cloud data centre than the typical on-premises data centre.

Cooling is not the only factor that contributes to facility-related energy overhead. Older power distribution systems perform poorly at low loads (e.g., less than 30% utilisation). Many APAC enterprises and public sector organisations have power distribution systems 5 to 10 years old or older and/or not highly utilised, so they are likely to be less efficient.

Adding it all up, APAC survey respondents use as much energy on cooling and power distribution as on their IT infrastructure itself, leading to an average power usage effectiveness (PUE)¹² measure of 2 in the survey sample, worse than the 1.66 average PUE measure observed in the 2019 U.S. survey. This is not surprising, given the high average heat and humidity in surveyed APAC locations, yet it demonstrates the difficulties APAC enterprises and public sector organisations face in making major improvements to the efficiency of their data centres. Hurdles include the capital cost of upgrades and the potential risk of downtime during such a major project. Application downtime can prove very costly to business operations. This may explain why APAC survey respondents generally lack comprehensive programs to optimise data centre operations, with only a handful of surveyed organisations reporting data that consistently outperforms their peers across multiple metrics.

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

11. The Unexpected Impact of Raising Data Center Temperatures, White Paper 221, Schneider Electric, 2016

12. PUE, the ratio of total data centre power and IT power

By contrast, cloud operators design and operate their servers with the expectation of higher temperatures, implementing the latest advancements in computational fluid dynamics to minimise the adverse impact of hotter and more humid climates in APAC. Our research shows that even in the more challenging climates of many APAC countries, hyperscale cloud data centres (a blend of owned and leased data centre sites) manage PUE measures of 1.2 or better on average – representing savings of greater than a third of total energy consumption.¹³

Finally, another difference between enterprises and public sector organisations and cloud operators stems from the business model of cloud. The cloud model delivers IT services at scale. To reduce operating costs, cloud providers have an incentive to make the entire technical organisation work in unison to optimise engineering choices and increase efficiency, from design through to operations. This is further strengthened by economies of scale and the ability to fund custom engineering efforts and constant improvements.

By contrast, enterprises and public sector organisations tend to manage IT infrastructure in a more fragmented way, with efficiency initiatives that are narrower in focus and harder to scale. The top IT priorities for APAC survey respondents are to keep up with changing business requirements and to guarantee the security, availability, and reliability of critical applications.

Even if efficiency improvements are possible (e.g., via system suppliers and integrators, if not in-house), for many organisations these remain possibilities rather than priorities. Because digital infrastructure is not their core business focus, enterprises and public sector organisations typically have not designed their IT systems to prioritise energy and carbon reduction. This is not surprising as, for most on-premises IT operations, the electricity bill is negligible compared with the cost of software licences and the IT payroll, while the cost of a system-wide upgrade of server and data centre infrastructure to maximise efficiency can be quite high. APAC survey respondents also show a general lack of comprehensive programs to optimise data centre operations in line with best practices, with only a handful of surveyed organisations reporting data that consistently outperforms their peers across multiple metrics.

13. While we note that PUE is an imperfect measure as there is no enforced standard, we nonetheless use PUE in our survey as the best available proxy for energy efficiency for data centre facilities. We also note that in some cases, efficiency gains behind the meter (e.g., efficiency gains due to a shrinking IT footprint, or the use of higher density equipment) may adversely impact PUE, but statistically speaking, these observations should not substantially change our analysis.

The 451 Research Cloud Energy Efficiency Model

Using the results of the survey described above, 451 Research/S&P Global Market Intelligence calculated the relative operational efficiency at the largest data centres owned by the enterprises and public sector organisations surveyed. To establish a baseline, we looked to the server power database of the Standard Performance Evaluation Council (SPEC). We used the SPECpower_ssj2008 benchmark¹⁴ to simulate a business application and estimate the power consumed by different servers at various utilisation rates.

To create a baseline reference point, we took results from Intel-based servers from 2010 at 10% load. In our model, the average energy efficiency of a 2-processor system based on the 2010 Intel Xeon server processor at 10% load (typical of non-virtualised servers) is '1.' Using the survey data, we then approximated the profile and utilisation of on-premises server infrastructure to estimate energy efficiency. We modelled the server infrastructure makeup of enterprises and public sector organisations by taking into account lifespan, technology adoption speed, and server infrastructure growth dynamic. To estimate utilisation for each cohort at the surveyed organizations, we used survey responses on virtualisation levels and workload consolidation practices.


Figure 4 illustrates how different server generations (left to right) and different utilisation rates (bottom to top) impact efficiency. Both are key to understanding why server hardware is not equally efficient across its load curve.¹⁵

14. https://www.spec.org/power_ssj2008/

15. Our modelling suggests that the growth of cloud services will push the efficiency of hyperscale cloud infrastructure ahead at a faster rate than efficiency gains in the enterprise in coming years. The proliferation of more hyperscale data centres in APAC using the latest server processors that favour workload sharing will deliver a jump in efficiency.

Figure 4: Efficiency Curves of Intel Server Generations per SPECpower_ssj2008 Database

Source: SPEC.org, compiled by 451 Research



BLACK & WHITE | THE CARBON REDUCTION OPPORTUNITY OF MOVING TO THE CLOUD FOR APAC

451 Research

S&P Global
Market Intelligence

Energy Savings from Moving Enterprise and Public Sector Workloads to Cloud Infrastructure in APAC

Cloud servers are therefore three times more efficient than the average APAC enterprise servers.

Our APAC survey results indicate that, thanks to advances in server technology (mainly due to faster and lower-power semiconductors), APAC enterprise and public sector server infrastructure has become much more efficient over time. On average, APAC respondents scored 6.8 in server efficiency, meaning that their server infrastructure is nearly seven times more efficient than non-virtualised servers at the start of the past decade. We estimate there is a factor-of-four difference between the best and worst organisations in terms of server infrastructure efficiency, with the best achieving 11-12 times efficiency gains, and the worst roughly 3.2 times.¹⁶


By comparison, 451 Research/S&P Global Market Intelligence estimates that hyperscale cloud server energy efficiency is around 21 times more efficient than the older non-virtualised servers, and therefore *cloud servers are three times more efficient than the average APAC on-premises servers*. This is due to the much higher utilisation of cloud servers, and infrastructure that is heavily weighted towards newer, more energy-efficient server technology that has been optimised for maximum efficiency.

Moving workloads from on-premises data centres to the cloud could reduce workload energy consumption by 79% on average for APAC enterprises and public sector organisations

On top of server-level advantages, hyperscale cloud operators also have more efficient data centres with much better PUE than typical on-premises facilities. This comes from the use of more efficient cooling methods, optimised airflow, and newer, more efficient electrical infrastructure that reduces power distribution losses. When combined with server-level efficiency, we estimate that *cloud infrastructure is five times more energy efficient on average than that of most APAC enterprises and public sector organisations*. Thus, moving a workload from an on-premises server to the cloud would produce a 79% reduction in workload energy consumption on average across APAC – even greater than the average 72% reduction estimated for U.S. enterprises in our 2019 study. This does not include the further reduction that could be achieved if cloud operators were able to source 100% renewable energy.

16. Our survey results also show that the APAC average on enterprise adoption of virtualisation and lower workload consolidation ratios was considerably lower than those reported in the 2019 U.S. study – a gap of about 16% between the two regional enterprise averages.

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




Our survey of 515 enterprises and public sector organisations across Japan, Singapore, South Korea, Australia, and India found that, on average, moving applications to the cloud could compress the energy footprint of a workload to one-fifth of that of running the same workload in on-premises data centres. As the power generation in these countries still relies heavily on fossil fuels (such as coal and natural gas), this compressed workload energy also results in significantly lower Scope 2 emissions (indirect emissions linked to energy consumption).¹⁷

For a 1-megawatt enterprise data centre, assuming average utilisation of 30%, moving to the cloud can reduce workload-related Scope 2 greenhouse gas emissions by nearly 2,400 metric tonnes per year. That is the equivalent of removing over 2,000 cars from the road.

According to our model calculations, across the APAC countries surveyed, for a 1-megawatt enterprise/public sector data centre (equivalent to about 500-1,000 square metres of IT space, or a ‘moderate size’ cloud migration project), assuming average utilisation of 30%, switching all applications to cloud services could reduce workload-related Scope 2 greenhouse gas emissions by nearly 2,400 metric tonnes per year. This is the equivalent of removing over 2,000 cars from the road.¹⁸

Figure 5: APAC Cloud Carbon Reduction Potential

Source: 451 Research/S&P Global Market Intelligence



17. https://ghgprotocol.org/scope_2_guidance

18. Estimated assuming 2300 grams of CO₂ produced per litre, with an average distance traveled of 10,000 km per year at 20 km to the litre.


BLACK & WHITE | THE CARBON REDUCTION OPPORTUNITY OF MOVING TO THE CLOUD FOR APAC

Greater access to corporate renewable energy options would boost the cloud’s carbon reduction potential even further in APAC. Even though APAC electric power grids are making efforts to improve efficiency and sustainability, as are enterprises and public sector organisations, our research shows that a boost in renewable energy for hyperscale cloud providers – all of whom have made commitments to reach 100% renewable energy across their global operations within the next decade – would lead to an even greater reduction in carbon emissions produced by IT workloads.

This is because hyperscale cloud operators have the organisational skill set, long-term view, and large-scale demand to drive more renewables projects than is possible for most enterprise or public sector customers. For example, renewable energy projects commissioned by hyperscale cloud providers typically start at tens of megawatts of capacity, with some cloud operators contracting for new renewable power output well in excess of 100 megawatts (the equivalent of dozens of offshore wind turbines) under a single contract.¹⁹

Figure 6: APAC Cloud Carbon Reduction Potential at 100% Renewable Corporate PPA Mix

Source: 451 Research/S&P Global Market Intelligence



19. See, for example, <https://blog.aboutamazon.eu/sustainability/amazon-announces-its-largest-single-renewable-energy-project-yet>

BLACK & WHITE | THE CARBON REDUCTION OPPORTUNITY OF MOVING TO THE CLOUD FOR APAC


In a hypothetical scenario in which cloud operators are able to source 100% renewable power for their cloud infrastructure in each given market, the carbon emissions produced by IT workloads would be even more significantly reduced across all of the APAC markets covered in this study.

In this scenario, using the same example of the 1-megawatt on-premises data centre from before, surveyed APAC organisations would, on average, save an additional 587 metric tonnes per year of emissions by shifting workloads to the cloud – on top of the already sizeable reduction attributable to improved server and data centre efficiency described above. This additional carbon reduction is equivalent to removing an extra 475 cars from the road on average²⁰ with each cloud migration project of moderate size. Yet in many cases, the lack of accessible or affordable corporate renewable energy options in APAC still leaves a significant amount of carbon reduction potential on the table. Multiplied to thousands of on-premises data centres across APAC, the collective greenhouse gas reduction from moving workloads to the hyperscale cloud could be equivalent to reductions from millions of households.

20. Estimated assuming 2300 grams of CO2 produced per litre, with an average distance traveled of 10,000 km per year at 20 km to the litre.

Appendix: Country Profiles – Some Differences, But Similar Conclusions

Country Profile: Japan



Equivalent Household's Electricity Use Saved if 25% of organisations moved 1MW of IT to the Cloud



1 house = 10,000 households


In terms of market sizing, 451 Research/S&P Global Market Intelligence's Datacenter KnowledgeBase estimates that the Japan data centre market will have 1.23 million square metres (m²) of operational space by 2026, with a compound annual growth rate (CAGR) of 4% throughout the forecast period (2020-2026).

Based on our survey results, Japanese enterprise and public sector IT energy efficiency is below that of the APAC average as a result of longer server lifecycles and a slower adoption of newer server platforms compared with peers. This results in typically older server clusters on average. Virtualisation rates are also relatively lower than, for example, South Korea or India. This slight disadvantage is more than offset by efficiently run data centre facilities, which help Japanese enterprises and public sector organisations to be more energy efficient in their overall digital infrastructure than the APAC average, according to our calculations.

The average Japanese core enterprise and public sector data centre in our sample would see a carbon emissions reduction of 1,885 metric tonnes for every megawatt of data centre capacity migrated to the cloud (the equivalent of between 2,000-4,000 servers). This is attributed to a 77% reduction in workload energy on average. If cloud providers were able to source 100% renewable energy, the total average emission reduction value would amount to 2,378 metric tonnes per year for every megawatt of enterprise/public sector data centre capacity displaced by cloud services. With an estimated number of 11,000 businesses in Japan that have more than 250 employees,²¹ if even only 25% of these firms put 1MW of IT load into the cloud (a cloud migration project of moderate size), this would save the equivalent of roughly a year's worth of emissions from 328,000 Japanese households' electricity use.

21. Sources: OECD 2020 (data from 2017) and IEA Atlas of Energy (data from 2018). Also, note that in our survey, the mean amount of IT equipment for large enterprises and public sector organisations was 4 MW.

Country Profile: South Korea




While Korean businesses are still warming to the idea of digital transformation (according to the Korea Industrial Technology Association, less than 10% of Korean businesses have actively pursued a move to the cloud), the continued development of 5G and other innovative technologies will give rise to edge cloud and distributed data centres. 451 Research/S&P Global Market Intelligence's Datacenter KnowledgeBase estimates that the data centre market will reach 407,000 m² of operational space by 2026, with a CAGR of 3% throughout the forecast period (2020-2026).

Our survey reveals that South Korean organisations have seen their server infrastructure expand at the highest rate among APAC counties studied, and they are also among the fastest to adopt the latest server technology. Yet data centre infrastructure energy efficiency is still below APAC average due to a more conservative approach to workload consolidation. Relatively speaking, South Korean organisations seem to prefer guaranteed application performance over concerns about energy consumption.

On average, South Korean enterprises and public sector organisations would see workload energy savings of 80% from migrating to hyperscale cloud services. For every megawatt of enterprise/public sector data centre capacity, we estimate this translates into 2,123 metric tonnes of carbon reduction annually, considering the grid emission factor in South Korea. If cloud service providers were able to source 100% renewable energy for their infrastructure, this would amount to 2,605 metric tonnes in total per megawatt of enterprise data centre capacity per year, compared with the current baseline. With an estimated number of 2,400 businesses in South Korea that have more than 250 employees,²² if even only 25% of these firms put 1MW of IT load into the cloud (a cloud migration project of moderate size), this would save the equivalent of roughly a year's worth of emissions from 53,000 Korean households' electricity use.

22. Sources: OECD 2020 (data from 2017) and IEA Atlas of Energy (data from 2018). Also, note that in our survey, the mean amount of IT equipment for large enterprises and public sector organisations was 4 MW.

Country Profile: Singapore




The growing appetite for technologies such as big-data analytics, AI, and IoT has been a driving force for the adoption of cloud infrastructure services in Singapore, and there is broad use of cloud-based offerings to accelerate sectoral transformation in the country. However, especially after accounting for the current moratorium on new data centre builds, according to 451 Research/S&P Global Market Intelligence's Datacenter KnowledgeBase, the Singapore data centre market is expected to remain relatively flat, reaching 386,000 m² of operational space by 2026, with a CAGR of 2% throughout the forecast period (2020-2026).

In our survey of APAC enterprises and public sector organisations, only Singaporean organisations managed to outperform the APAC average for both data centre server-level *and* facility-level efficiency. This is due to the fact that enterprises and public sector organisations in Singapore consolidate workloads more (to reduce the number of servers needed), which drives up utilisation and efficiency. Singaporean organisations have also been relatively early adopters of the latest server technology, and refresh their servers faster on average (roughly every three-and-a-half years). Despite the tropical climate, on-premises data centres in Singapore are among the most efficient in APAC, although at an average self-reported PUE of 1.89, they are still far off from the state-of-the-art cloud facilities. While still low in terms of overall renewable energy penetration, the grid in Singapore is currently the least carbon-intensive in surveyed APAC countries, due to heavy reliance on natural gas as opposed to coal.

Yet Singaporean enterprises and public sector organisations moving to the cloud would still see an average workload energy reduction of 76%, which amounts to 1,542 metric tonnes of carbon per megawatt of data centre capacity per year. Greater access to renewable energy resources for cloud providers would result in additional emissions removal of up to 434 metric tonnes per year, compared with the grid baseline. Powering cloud services with 100% renewable energy would bring total workload energy-related carbon reductions to 1,976 metric tonnes per megawatt per year. With an estimated number of 1,300 businesses in Singapore that have more than 250 employees,²³ if even only 25% of these firms put 1MW of IT load into the cloud (a cloud migration project of moderate size), this would save the equivalent of roughly a year's worth of emissions from 23,500 Singaporean households' electricity use.

23. Sources: OECD 2020 (data from 2017) and IEA Atlas of Energy (data from 2018). Also, note that in our survey, the mean amount of IT equipment for large enterprises and public sector organisations was 4 MW.

Country Profile: Australia



It has been more than two years since the Australian government first announced its national digital transformation strategy. Aside from establishing a task force to advance the country's digital agenda, the government has put together an ambitious roadmap to achieve its goal of becoming one of the world's top three digital governments by 2025.


According to 451 Research/S&P Global Market Intelligence's Datacenter KnowledgeBase, the data centre market in Australia is expected to grow to 448,000 m² by 2024, with a CAGR of 7% throughout the forecast period (2019-2024).

While Australian organisations score higher on server-level efficiency, largely by maintaining relatively newer server infrastructure (virtualisation rates are below average), their overall performance is hindered by below-average facility-level efficiency. Australian on-premises data centres are the second-least efficient (after India) in terms of facility-level efficiencies. One potential explanation could be the Australian population's geographic distribution, which is highly dispersed, resulting in relatively smaller enterprise and public sector data centres that are comparatively less optimised.

Australian businesses and public sector organisations can expect an average energy savings of around 80% from moving workloads to cloud infrastructure, a carbon reduction of 3,149 metric tonnes of CO₂e per megawatt of data centre capacity, as shown in our scenario modelling. Thanks to favourable policy and market structures, corporate renewable energy options are greater in Australia than in other parts of APAC. When reaching 100% renewable energy for their infrastructure in Australia, cloud operators will be able to curb emissions further, by up to 746 metric tonnes, leading to total workload energy-related emission reductions of 3,895 metric tonnes. With an estimated number of 2,900 businesses in Australia that have more than 250 employees,²⁴ if even only 25% of these firms put 1MW of IT load into the cloud (a cloud migration project of moderate size), this would save the equivalent of roughly a year's worth of emissions nearly 60,000 Australian households' electricity use.

24. Sources: OECD 2020 (data from 2017) and IEA Atlas of Energy (data from 2018). Also, note that in our survey, the mean amount of IT equipment for large enterprises and public sector organisations was 4 MW.

Country Profile: India



According to 451 Research/S&P Global Market Intelligence's Datacenter KnowledgeBase, the data centre market in India is expected to grow to 580,000 m² by 2024, with a CAGR of 16% – the highest among APAC markets surveyed – throughout the forecast period (2019-2024). As data centre activity continues to surge, so will energy consumption, which will make energy efficiency a focal point for enterprise/public sector users and data centre providers.

In our study, the server-level efficiencies of Indian businesses and public sector organisations exceeded their peers in other surveyed APAC countries, as a result of higher rates of virtualisation and a more aggressive stance towards workload consolidation. Indian organisations drive their systems somewhat harder, and their server infrastructure (thanks to constantly growing requirements for more capacity) are among the youngest on average in APAC. However, much of this is offset by inefficiencies at the facility level. Even considering the higher levels of humidity and warmer climate throughout Indian IT clusters such as those in Mumbai, Bangalore, and Chennai, India's average reported PUE of 2.3 leaves ample room for improvement.


Due to coal-heavy electricity generation in India, moving the workload equivalent of one megawatt of IT data centre capacity to the cloud would yield a massive carbon reduction of 3,449 metric tonnes per year on average. If cloud providers were able to provide services powered by 100% renewables, a total of 4,235 metric tonnes of CO₂e emission could be avoided by running the same enterprise and public sector workloads in the cloud. With an estimated number of 1,200 large publicly-traded businesses in India,²⁵ if even only 25% of these firms put 1MW of IT load into the cloud (a cloud migration project of moderate size), that would save the equivalent of roughly a year's worth of emissions over 160,000 Indian households' electricity use.

25.Sources: OECD 2020 (data from 2017) and IEA Atlas of Energy (data from 2018). Also, note that in our survey, the mean amount of IT equipment for large enterprises and public sector organisations was 4 MW.

Methodology

The focus of this model is to capture the carbon impact of key design and operational features of enterprise and public sector organisation data centres compared with the hyperscale cloud, and to understand how key components impact the overall efficiency picture. The core of the model centres on Scope 2 emissions (utility grid electricity), and does not include Scope 1 emissions (direct emissions from site operations, such as vehicle emissions, cooling system refrigerants, diesel engine power generator emissions, etc.), nor Scope 3 (embodied in buildings and IT products).

We decided to focus on Scope 2 because enterprises and public sector organisations have direct control over key factors to influence energy consumption, which substantially affects data centre energy efficiency and carbon footprint. In our view, Scope 1 emissions do not reflect the core operational efficiency of a data centre, as there is little room for differentiation (for example, virtually all operators need generators that run tests or are utilised during grid failures). Scope 3 emissions are more significant, and future versions of the model may include embodied carbon calculations for facility and IT systems, but we do not expect these to meaningfully alter the conclusion of the analysis.



The carbon emissions model consists of five major areas: embedded (or embodied) emissions of both data centre facilities and IT hardware; carbon intensity of the grid; then facilities and IT operational emissions. The objective of the model is to show the difference between on-premises and cloud operations.

The model produced a ratio showing the relative energy and carbon efficiency difference between enterprise/public sector data centres and the cloud. The model incorporated APAC survey data, data from S&P Global Market Intelligence on APAC data centre and energy markets, third-party industry sources, and data from cloud operators. Some survey questions were not directly used in the model calculations, but for further background analysis.

Our survey focused on understanding some of the characteristics of APAC public sector organisations and enterprises (between \$10m and \$1bn in annual revenue) that influence efficiency metrics. These included policies and attitudes such as consolidation levels, speed of server technology adoption, and typical server lifespan. We believe that such an approach, while requiring careful tuning of some assumptions, creates a much more robust picture and provides better context than asking exclusively for technical specifications and operational metrics, many of which may not be tracked with the 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 ask for average lifespan, speed of adoption of new tech (to account for additional lag compared with the cloud) and capacity change (skew of distribution).
- *Server utilisation*: In our experience, it is very difficult to obtain accurate data when enterprises/public sector organisations are asked to share estimated server utilisation. Instead we ask 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/S&P Global Market Intelligence used these responses to estimate average utilisation based 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 manufacturers do aggressively fine-tune hardware and software specifically for the *specpower_ssj2008* benchmark to attain the best possible result in ways that may 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/public sector IT environment.

While SPEC data is in 10% increments, we needed finer granularity of 1% for our calculations as we modelled IT operational efficiency of the surveyed 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 or public sector organisation, then weighted such efficiency readings by distribution of server generations per organisation. 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.



About AWS Institute

The AWS Institute is a thought leadership and executive education program to accelerate digital transformation for public sector executives. Find our content and programs here <https://aws.amazon.com/institute/>

BLACK & WHITE | THE CARBON REDUCTION OPPORTUNITY OF MOVING TO THE CLOUD FOR APAC

451 Research

S&P Global
Market Intelligence

451 Research

S&P Global

Market Intelligence

About 451 Research

451 Research is a leading information technology research and advisory company focusing on technology innovation and market disruption. More than 100 analysts and consultants provide essential insight to more than 1,000 client organizations globally through a combination of syndicated research and data, advisory and go-to-market services, and live events. Founded in 2000, 451 Research is a part of S&P Global Market Intelligence.

© 2021 S&P Global Market Intelligence (S&P). All Rights Reserved. Reproduction and distribution of this publication, in whole or in part, in any form without prior written permission is forbidden. The terms of use regarding distribution, both internally and externally, shall be governed by the terms laid out in your agreement with S&P. The information contained herein has been obtained from sources believed to be reliable. S&P disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although the information contained herein may include legal issues related to the information technology business, S&P does not provide legal advice or services and its research should not be construed or used as such. S&P shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.



NEW YORK

55 Water Street
New York, NY 10041
+1 212 505 3030



SAN FRANCISCO

One California Street, 31st Floor
San Francisco, CA 94111
+1 212 505 3030



LONDON

20 Canada Square
Canary Wharf
London E14 5LH, UK
+44 (0) 203 929 5700



BOSTON

75-101 Federal Street
Boston, MA 02110
+1 617 598 7200

