Corporate Cloud and sustainability strategies

An introduction to the issues around decarbonization and achieving net zero in cloud computing

thoughtworks
Strategy. Design. Engineering.

aiven

Introduction

"It is international scientific consensus that, in order to prevent the worst climate damages, global net human-caused emissions of carbon dioxide (CO_2) need to fall by about 45 percent from 2010 levels by 2030, reaching net zero around 2050. Global warming is proportional to cumulative CO_2 emissions, which means that the planet will keep heating for as long as global emissions remain more than zero."

Oxford Net Zero Interdisciplinary Research Initiative, University of Oxford

Mitigating a business's impact on the environment is a key element of current corporate practice, and an increasingly important element of any future-readiness or future-proofing strategy. Business continuity, compliance, talent attraction and staff retention are all dependent on environmental and sustainability performance to varying degrees.

For example, in a recent Deloitte survey, 49% of the Gen Z age group (those born between 1997 and 2012) and 44% of millennials (1981 to 1996) said that they had made career choices based on their personal ethics. A Yale School of Management study of more than 2,000 students from 29 business schools found that 51% would accept a lower salary to work for an environmentally responsible company — up from 44% five years previously. At Harvard Business School, 600 students took second-year elective courses related to social enterprise in 2020, compared with 251 in 2012.

As customers, members of Gen Z also place greater emphasis on making consumer choices that are aligned with their values, with 73% saying they are willing to spend more on sustainable products, while 62% prefer to buy from sustainable brands. These numbers are a strong indication that environmental performance will become an increasingly important element in securing competitive advantage.

At the same time, the growing prominence of environmental, social and governance (ESG) measures and indices throughout the world's capital markets reflect the concerns of the investment community and their clients — and a direct link between financial and environmental performance.

Higher energy prices, supply chain pressures, increased interest rates and higher input costs all affect the value calculations and cost benefits of emission reduction programs. Nonetheless, analysts from McKinsey & Company estimate that energy-intensive industries in Europe could create anywhere from €3 billion to €12 billion in value by deploying energy-efficiency measures such as advanced analytics.

However, the relationship between a business's operations and the broader ecosystem is an intricate one. Understanding the environmental impacts of an individual business can be a complex process that covers not just its immediate operations, but its wider value chain.

This paper explores the decarbonization agenda as it applies to businesses, particularly those dependent on cloud-based services, including:

- The expansion of the types of emissions they are being required to measure
- The voluntary standards and legally binding obligations in place
- The challenges of measuring and managing emissions from cloud operations
- New solutions from Aiven and Thoughtworks to improve and standardize measurements

The Greenhouse Gas Protocol: direct and indirect types of emissions

There is as yet no universally applicable methodology for quantifying a business's environmental impact, although there are attempts by international bodies to offer advice and guidance. The most prominent of these is the Greenhouse Gas (GHG) Protocol, an international framework that builds on a 20-year partnership between the World Resources Institute (WRI) and the World Business Council for Sustainable Development (WBCSD).

The GHG Protocol provides standards, guidance, tools and training for business and governments to measure and manage climate-heating emissions from CO2, methane and other greenhouse gasses. Within its comprehensive and globally standardized series of frameworks it defines three distinct categories for measuring and managing emissions. These are: direct operations (known as Scope 1); indirect emissions from purchasing energy (Scope 2); and other indirect emissions in the value chain (Scope 3).

Scope 3 is considered the hardest to measure, because these emissions are outside an organization's direct control. For most companies, this is also the category into which emissions associated with cloud computing fall, since they are derived from outsourced services. Although the GHG Protocol does not have specific guidance on measuring cloud-based emissions, Scope 3 guidelines are therefore the ones to follow.

In 2020, 239 companies signed up to the Science Based Targets Initiative (SBTi), an independent organization promoting climate action in the private sector, with 94% saying they would reduce emissions linked to their customers and suppliers. The number of companies who had signed up by March 2024 was just under 8,000 — most of whom acknowledged that Scope 3 is their biggest barrier to achieving net zero.

The SBTi has provided a substantial framework for setting emissions-reduction targets in line with the latest climate science. Its Corporate Net-Zero Standard follows three key stages and is designed to enable businesses to achieve and remain at net zero:

- Near-term targets — five to 10 years. Emission reduction targets in line with 1.5°C pathways
- Long-term targets. Reducing emissions to a residual level in line with 1.5°C scenarios by no later than 2050
- Neutralization. Counterbalance of emissions by permanent removal and storage of carbon from atmosphere

While SBTi has been a voluntary framework, businesses can also expect increasing levels of regulation too.

The EU's new Corporate Sustainability Reporting Directive (CSRD) enforces standardized sustainability reporting with hefty fines for non-compliance. This applies not just to EU firms, but also to some international firms with subsidiaries or suppliers in the bloc. Mimicking financial reporting, the CSRD mandates disclosures on carbon emissions and other sustainability impacts, with details varying by company size and listing status. Companies will face a double materiality assessment, requiring them to report on both their environmental impact and how sustainability trends affect their business. This could include data on Scope 3 cloud emissions, potentially exceeding 1,000 data points per firm. The CSRD signifies a significant shift towards corporate accountability for environmental performance.

In brief, the first step to making a meaningful contribution to achieving net zero for most businesses is by quantifying the total amount of greenhouse gas emissions associated with their activities to establish their current carbon footprint. The move to decarbonization can then begin by identifying the greatest sources of emissions and reducing the material emissions from these ‘hotspots’. Reducing emissions within the value chain in this way is known as abatement.

For businesses who are dependent on cloud computing services, there are currently very few tools that accurately measure Scope 3 cloud emissions in a cost-efficient way. Without standards, most cloud providers have developed their own, often proprietary, measurement techniques. Even when such tools are deployed, the complex and distributed nature of typical cloud infrastructures means that measuring emissions accurately and reliably can be extremely challenging. It also means that this crucial area of cloud operations does not benefit from the collaborative approach that makes open-source technologies so appealing.

Internal management of sustainability and carbon reduction

The impacts of the drive to net zero, sustainability demands and decarbonization on cloud-based services are extensive, and will therefore raise significant questions around internal organization, incentivization and data sharing to name just three.

Most companies will need a corporate sustainability strategy, with associated management and assigned responsibilities. The strategy itself will encompass the measurement and reduction of greenhouse gas emissions and will typically be approved by C-level executives since it is likely to occupy an increasingly high-priority position within the overall business strategy.

For day-to-day purposes, the sustainability strategy is often owned by a sustainability team that consists of subject-matter experts who work with the operational teams who will execute it across the business. However, many individual contributors sit outside that team, and, for them, the sustainability strategy can feel very abstract, which makes it difficult to relate and contribute to. Software developers and other users of the Aiven platform within an organization often fit into this category, despite regular interaction with cloud-based services and carbon-emitting infrastructure.

Getting to grips with cloud carbon emissions

During a period of heightened awareness among Thoughtworkers about the intersection of digital services and climate change, Cloud Carbon Footprint emerged as a tool to address the environmental impact of cloud computing. Recognizing the potential for significant emissions in the cloud computing sector, Thoughtworks sought to quantify its own carbon footprint but faced obstacles due to the absence of suitable tools and methodologies.

Motivated by Etsy's innovative Cloud Jewels methodology for estimating the energy consumption of Google Cloud compute resources, Thoughtworks embarked on developing Cloud Carbon Footprint. Leveraging its expertise in custom software development and open-source contributions, Thoughtworks released Cloud Carbon Footprint as an open-source tool in March 2021. The aim was to provide organizations with the means to assess their cloud energy usage and carbon emissions, thereby fostering awareness and more thoughtful usage of cloud services alongside necessary transparency to make data-driven optimizations. This initiative also aimed to encourage cloud providers to share this data with their customers.

As of today, Cloud Carbon Footprint stands as a pioneering solution for multi-cloud organizations to track both their historical and real-time carbon emissions using a uniform and transparent approach. It has expanded its scope to encompass additional features such as on-premise machines, marginal emission factors as well as embodied emissions, continuing to facilitate environmental accountability in cloud computing.

How Cloud Carbon Footprint helps you measure CO2e

Whether Thoughtworks' cloud carbon footprint tool is being applied to compute or storage — Aiven's two main current usage focus areas — networking, memory or other area, CCF's methodology is based on the following **formula**:

Total CO2e = Operational emissions + Embodied emissions

Where:

Operational emissions = Cloud provider service usage × Cloud energyconversion factors [kWh] × Cloud provider Power Usage Effectiveness (PUE) × Grid emissions factors [metric tons CO2e]

and:

Embodied Emissions =
estimated metric tons CO2e emissions from the manufacturing of datacenter servers, for compute usage

First, the application establishes cloud usage via provider billing data. Second, it estimates energy used by leveraging a similar process to the Cloud Jewels methodology. Finally, the application converts estimated kilowatt hours into estimated CO2e. These processes are inevitably improved and refined on an ongoing basis.

Data can then be visualized in a dashboard for developers, sustainability leaders and other organizational stakeholders to view and then take action. CCF currently supports AWS, Google Cloud and Microsoft Azure.

A starting point for each individual client use case

Given how each business's IT operations context are unique — for example, each involves a different cloud setup and tech stack — the bespoke specifics of approach for each client use case will naturally depend on variables identified at the all-important consulting stage.

We use domain-driven design to separate the estimation logic from both the data input source (e.g. cloud APIs, on-premise or co-located data centers) and the output source (e.g. front-end dashboard, CSV, etc.) so new inputs and outputs can easily be added.

Here we explore CCF's cloud usage, energy usage and estimated CO2e processes in greater detail.

Step 1: Establishing cloud usage via provider billing data

By default, the application queries the following cloud provider billing and usage reports to provide a holistic understanding of client emissions based on each resources' usage type and service

- AWS Cost and Usage Reports with Amazon Athena
- GCP Billing Export Table using BigQuery
- Azure Consumption Management API

This approach provides a more holistic estimation of cloud energy and carbon consumption, but as a trade-off uses a less accurate average constant (rather than measured) CPU utilization.

Step 2: Estimating energy consumption

CCF estimates energy used by cloud provider resources to determine energy coefficients (kWh as summarized here, that we combine with specifics for the region and its energy production characteristics:

aws

Amazon Web Services

- Average Minimum Watts (0% CPU Utilization): 0.74
- Average Maximum Watts (100% CPU Utilization): 3.5
- Average CPU Utilization for hyperscale data centers: 50%
- HDD Storage Watt Hours / Terabyte: 0.65
- SSD Storage Watt Hours / Terabyte: 1.2
- Networking Kilowatt Hours / Gigabyte: 0.001
- Memory Kilowatt Hours / Gigabyte: 0.000392
- Average PUE: 1.135

Google Cloud Partner

- Median Minimum Watts (0% CPU Utilization): 0.71
- Median Maximum Watts (100% CPU Utilization): 4.26
- Average CPU Utilization for hyperscale data centers: 50%
- HDD Storage Watt Hours / Terabyte: 0.65
- SSD Storage Watt Hours / Terabyte: 1.2
- Networking Kilowatt Hours / Gigabyte: 0.001
- Memory Kilowatt Hours / Gigabyte: 0.000392
- Average PUE: 1.1

Azure

- Average Minimum Watts (0% CPU Utilization): 0.78
- Average Maximum Watts (100% CPU Utilization): 3.76
- Average CPU Utilization for hyperscale data centers: 50%
- HDD Storage Watt Hours / Terabyte: 0.65
- SSD Storage Watt Hours / Terabyte: 1.2
- Networking Kilowatt Hours / Gigabyte: 0.001
- Memory Kilowatt Hours / Gigabyte: 0.000392
- Average PUE: 1.125

The servers used by cloud providers are taken into account alongside their energy usage, as detailed in the SPECPower database and 2016 US Data Center Energy Usage Report. Etsy investigated GCP servers, and Thoughtworks has additionally analyzed AWS and Azure servers. Understanding of the processors used reflects the best information publicly available but does not take account of any custom processors, such as those deployed by AWS.

A typical compute estimation will follow the same, two-step formula as Cloud Jewels.

Firstly, Average Watts are calculated—the average compute energy at a moment in time. When a server is idle, it still takes some power to run it (Minimum Watts). As the server utilization increases, the amount of power consumed increases too. The total energy used is the Minimum Watts figure plus the watts from additional server usage (average per hour).

Average Watts

= Min Watts + Avg vCPU utilization * (Max Watts - Min Watts)

Secondly, this is translated into total Watt Hours based on the amount of time servers are being used, or virtual CPU hours.

Compute Watt-Hours = Average Watts * vCPU Hours

Input data sources for the formula variables and source context

Min Watts (constant)
depends on the CPU processor used by the Cloud provider to host the virtual machines. The constant per processor microarchitecture was determined using the SPECPower database.

Max Watts (constant)
calculated as per Min Watts, above.

Avg vCPU Utilization (variable or constant)
either pulled from the cloud provider APIs or falls back to a 50% average taken from the 2016 U.S. Data Center Energy Usage Report.

vCPU Hours (variable)
pulled from the cloud provider billing and usage data.

When the underlying processor micro-architecture or group of micro-architectures used for a given cloud provider virtual machine are known, the corresponding minimum and maximum watts are used. When a group of micro-architectures includes either Ivy Bridge or Sandy Bridge, they are treated as outliers using the median of that group—in order not to overestimate.

When the underlying processor micro-architecture is unknown, the average or median of all micro-architectures used by that cloud provider is used.

Why Graphic Processing Units (GPUs) demand a different approach

As the SPECPower Database doesn't include energy data for the minimum and maximum watts of GPUs, a different approach is required.

This applies the same compute estimation formula. However, because the cloud providers provision entire physical GPUs to customers, GPU Hours are used instead of virtual CPU Hours. Minimum and maximum watts of GPUs are therefore based on a data set published by Teads that applies to AWS GPUs—ratios we have also assumed for GCP and Azure.

Step 3: Estimating carbon (CO2e)

Once kilowatt hours for usage of a given cloud provider are estimated, they are then converted into estimated CO2e using publicly available data on emission factors for a given electricity grid based on the mix of local energy sources. This is based on the cloud provider datacenter region in which each service is running. Sources are as follows:

Google Cloud Partner

Google Cloud Partner

- Google figures for grid carbon intensity by GCP regions are here.

AWS and Azure

- In the US, the application uses the EPA's eGRID2020 Data that provides NERC region specific emission factors annual for CO_2e — a better representation of energy consumed by data centers than more granular eGRID subregion or state emissions metrics.
- Outside the US, the application uses the carbonfootprint.com country specific grid emissions factors report. For Europe specifically, it draws on EEA emissions factors.

As these data sources are averages over a given year that is pre-2020, and also don't take into account time of day, there is an option to use Electricity Maps API. This provides real-time, historical, and forecasted electricity emissions data.

Integrating Thoughtworks' Cloud Carbon Footprint with Aiven services

In light of all these challenges, Aiven is developing new products that will allow customers to manage their Scope 3 cloud emissions. The goal is to give sustainability managers and executives access to more accurate calculations on cloud emissions for both decarbonization and reporting purposes, and to increase visibility of emissions data across the organization, while enabling software developers to contribute to emission reduction efforts.

This development aligns strongly with Aiven's vision for ESG: To be a driving force for positive change, creating sustainable data solutions for application builders and for the planet. Because Aiven offers open-source, multi-cloud solutions to manage infrastructure efficiently, it is uniquely positioned to help its customers to measure, visualize, and manage their cloud emissions in a transparent and standardized way.

To achieve this, Aiven is working in partnership with Thoughtworks to build on Cloud Carbon Footprint, Thoughtworks' free and open-source tool for measuring and monitoring carbon emissions arising from the use of cloud resources.

Thoughtworks' Cloud Carbon Footprint methodology is well suited for calculating the emissions associated with a specific cloud infrastructure setup. However, to really serve the Aiven client base, it needs to achieve the same level of availability, granularity, accessibility, and standardization in emissions data from resources across multiple cloud regions or providers.

Aiven maintains a detailed record of all the resources used in the implementation of each of its services — for example, the number of virtual machines running, their utilization, and their network use over a specific time period. Emissions reporting needs to be at a similar level of granularity to provide meaningful support for emissions reduction and infrastructure optimization strategies. It therefore needs to be at a single service instance level, and to show how emissions fluctuate over the course of a day, week or month.

To get to this point, Aiven tapped into Thoughtworks expertise to create a Python-based port of the Cloud Carbon Footprint tool, known on GitHub as the Cloud Emissions Estimator and published under an open-source license. The port replicates the calculations made by Thoughtworks' Cloud Carbon Footprint, and produces the same energy and carbon emission estimates, but those are based on the input of any designation group of cloud resources.

As a result, Aiven's solution supports reporting on estimated emissions based on past usage, and has potential to guide users to select less carbon-intense locations for their own workloads. In the future, Aiven wants to explore the implementation of functionality to provide active recommendations and suggestions for right-sizing solutions and optimizing locations to improve environmental performance.

Transparent calculations are essential for allowing everyone to verify results and suggest improvements — or even suggest alternative use cases — which is why Thoughtworks' Cloud

Carbon Footprint and Aiven's Carbon Footprint Tool are released under permissive, open-source license.

Let us examine a hypothetical scenario to provide further practical details: Consider a virtual machine hosted on Google Cloud Platform for a duration of one month, equivalent to 730 hours for the sake of simplicity. This virtual instance, situated in Europe West 1, Belgium, is configured with eight CPUs featuring Intel Ice Lake architecture, each operating at a utilization rate of 20%.

Utilizing the established formula for watt-hour estimation - Watt-hours = CPU Count * (Min Watts + CPU utilization * (Max Watts - Min Watts)) * Running Hours, with Ice Lake's parameter values set at 0.77 and 3.97 - we derive an estimated total energy consumption of 8 * (0.77 + 0.20 * (3.97 - 0.77)) * 730, resulting in 8234.40 watt-hours.

Subsequently, incorporating the estimation of the provider's Power Usage Effectiveness (PUE), which stands at 1.1 for Google Cloud Platform, we refine our calculation. The total estimated energy consumption now scales up to 8234.40 * 1.1, yielding 9057.84 watt-hours.

In the final step, we convert this estimated energy consumption into carbon emissions by juxtaposing it with the Grid Carbon Intensity. For the Europe West 1 region of Google Cloud Platform, this intensity is set at 123 grams of CO2 equivalent per kilowatt-hour of energy consumed. Consequently, in our illustrative case, this translates to estimated CO2 emissions of 1.11 kilograms.

It is essential to emphasize that this example focuses solely on CPU utilization. For a comprehensive assessment of actual usage, one must consider additional factors such as memory utilization, allocated storage, and networking impact.

Use case 1: Reliable, standardized emissions reporting

Accurate data on cloud emissions is foundational to high-quality accounting and disclosure practices. For companies who opt for a single cloud provider to build applications, the choice is either to rely on the provider's reported figures, which they manually integrate into their own footprint, or to use the estimation methods associated with their spending on cloud usage. This is known as the spend-based method, which has the advantage of simplicity — but at the expense of precision.

For larger enterprises who engage with multiple cloud providers, data gathering becomes more complex. Each provider employs their own methodology for calculating emissions and presenting information to users, although all typically synchronize data presentation with their billing cycles.

For example, Google reports comprehensively on all three Scopes of the GHG Protocol through its Cloud Carbon Footprint dashboard. However, this has a time lag of up to 21 days for access to the preceding month's data.

Both AWS and Microsoft Azure focus exclusively on users' Scope 2 emissions and emphasize the indirect emissions that stem from the electricity generated to power their data centers. AWS also has a three-month delay in displaying cloud emissions data through its own Customer Carbon Footprint tool.

In contrast, the Carbon Footprint Tool developed by Aiven with Thoughtworks, draws on data inputs from multiple cloud providers to meticulously calculate a highly granular, multi-cloud footprint. Anchored in leading open-source methodology, this tool estimates project-level cloud emissions on an hourly basis. This data is seamlessly integrated and accessible directly through the Aiven Console or via API integration, and is independent of billing cycles, to give users granular, accurate, standardized, and timely emissions calculations.

This approach goes beyond straightforward compliance to propel cloud-dependent businesses toward their sustainability targets.

Use case 2: Efficient and inclusive decarbonization

The challenge of measuring Scope 3 emissions is partly due to the inherent complexity of measuring the upstream and downstream activities of suppliers, and partly because Scope 3 accounts for such a large percentage of overall emissions. Enterprises who wish to be seen as agents of change, especially those with a mature sustainability strategy, must contend with their suppliers' operations and the influence they can apply.

This transformative role can involve requesting suppliers to adhere to specific ESG standards, demanding transparency regarding emission factors associated with their products, or urging the decarbonization of supplier operations. As regulations like the CSRD are enacted, companies who are obliged to report their own emissions under its terms, are also those who need to encourage their suppliers to go the extra mile.

As noted above, there is no single or universal approach to take. For industries with high emissions, decarbonization requires a significant financial commitment to re-shape their operations. In sectors where carbon intensity is already lower, companies are redefining their business models, crafting new strategies, and embracing values that resonate with their stakeholders.

But regardless of their path to decarbonization, these businesses have one thing in common: optimizing operational costs can also substantially reduce emissions. Whereas the move to sustainability was once seen as a financial burden — for example, using data centers located in regions powered by renewable energy — such investments now often yield disproportionately high impact. As noted above, these impacts can include, but are not limited to, averting non-disclosure fines, retaining and acquiring top talent, attracting high-quality investments, meeting customer expectations, and minimizing costs associated with neutralizing residual emissions.

All this puts C-level executives and sustainability teams under pressure to find ways to foster a holistic and inclusive approach to climate action. Employees across functions should be empowered to contribute, even if they are not directly associated with sustainability teams or work in operations with the highest emissions sources. This is particularly evident in tech companies, where a significant proportion of employees are software engineers. Excluding individual contributors from reduction strategies can not only stymy those specific efforts, it can also affect morale and misses an opportunity to align company action with employee values.

By exposing accurate and granular emissions data at the service-creation stage, Aiven could provide developers with the opportunity to actively participate in their company’s climate journey. It also enables developers to make environmentally conscious choices regarding infrastructure setups, and to factor in multiple cloud providers, locations, performance considerations, and forecast climate impacts. At the same time, senior executives gain visibility of the multiple functions within the organization that contribute to indirect emissions.

Conclusion: Innovation and potential for substantial change

The urgent need for rapid supplier decarbonization, the demand for cost optimization during the transition to net-zero, and the need for inclusive decarbonization strategies are a perfect storm of both drivers to action and threats to success. But the potential of Aiven's Carbon Footprint Tool, built on Thoughtworks' Cloud Carbon Footprint, to address a segment of Scope 3 emissions is a cause for some optimism. It harnesses the power of true partnership between sustainability teams and developers to foster efficient and inclusive climate action.

By improving collaboration, companies can optimize their resource consumption and cost management related to their net-zero program. Given the central and expanding role that cloud plays in the operations of organizations of all kinds, it has the potential to achieve a significant impact on Scope 3 emissions.

About Thougtworks

Thoughtworks is a global technology consultancy that integrates strategy, design and engineering to drive digital innovation. We are over 11,500 Thoughtworkers strong across 51 offices in 18 countries. For 30 years, we've delivered extraordinary impact together with our clients by helping them solve complex business problems with technology as the differentiator.

thoughtworks.com

About Aiven

Aiven is a leading global data and AI platform company dedicated to helping organizations extract maximum value from their data. Their cloud platform seamlessly integrates a range of open-choice services for streaming, storing, and serving data, ensuring simplicity, security, and speed across major cloud providers. Trusted by thousands of customers worldwide, Aiven enables the swift and confident creation of next-generation applications.

aiven.io

Strategy. Design. Engineering.