IIIEE Theses 2026:17

The Weight of the Cloud: Who Owns the Rain?

Cloud Sustainability Beyond Metrics: Provider Transparency and Customer-Side Feasibility in SaaS Deployment Decisions

Neo4j as a Case Study

Aoife Michelle Ní Uanacháin

Supervisors

Matthias Lehner

Paulina Aguilera Salgado

Thesis for the fulfilment of the

Master of Science in Environmental Management and Policy

Lund, Sweden, May 2026

Published in 2026 by IIIEE, Lund University, P.O. Box 196, S-221 00 LUND, Sweden,

Tel: +46 – 46 222 02 00, Fax: +46 – 46 222 02 10, e-mail: iiiee@iiiee.lu.se.

ISSN 1401-9191

Stop Saying the Cloud is Just Someone Else's Computer,
Because it's Not

– Mary Branscombe

Acknowledgements

I wish to express my deepest gratitude to the many individuals who have supported me throughout this academic journey and my time at the Institute.

To my esteemed supervisor, Matthias, thank you for being the first professor to welcome me into the Insti, marking the beginning of this incredible journey, and for your steady guidance in bringing this thesis to completion. To my esteemed supervisor, Paulina, I am profoundly grateful for your mentorship, which began with our consulting project and continued through my enriching internship experience at Neo4j; your ability to connect theory with practice has significantly shaped my professional development.

To the staff at the Insti, my sincere gratitude for the diverse opportunities you facilitated; from industry study visits to conferences and applied projects, these experiences offered a meaningful bridge between academic learning and professional application. It has been an honour to represent and study at Lund University, recognised as the world leader in the QS World University Rankings: Sustainability 2026.

To the wonderful people who made my personal life here so rich and rewarding, the inspiring members of Batch 31. Thank you for the shared experiences, friendships, and moments of joy that enriched this journey; you provided a sense of community that I will value forever. It was a privilege to serve as your StuCo.

Finally, to my family and friends back home, thank you for your constant encouragement and support. Your belief in me has been a continuous source of motivation and strength, allowing me to fully embrace this experience.

And last but not least, I want to thank myself.

Abstract

The rapid expansion of cloud computing and AI-driven digital services is reshaping global digital infrastructure, while simultaneously intensifying concerns surrounding the energy use, carbon emission, and broader environmental impacts of data centres. Although hyperscale Cloud Service Providers (CSPs) increasingly position cloud computing as sustainable and energy-efficient, it remains unclear whether the sustainability information they disclose is sufficiently transparent, comparable, and actionable to support real-world deployment decision by cloud-dependent organisation. This thesis addresses the gaps by examining both the availability of cloud sustainability information and its organisational feasibility and applicability in practice.

The study adopts a qualitative, interpretative research design structured in two phases. Phase 1 focuses on evaluating the transparency, comparability, and decision relevance of sustainability disclosures from Amazon, Microsoft, and Google; Phase 2 investigated how such information is interpreted and acted upon within Neo4j, a mid-sized SaaS organisation chosen as case study.

The thesis contributes to literature by linking provider-level transparency with customer-side feasibility, demonstrating that the key challenges in cloud sustainability lies in the translation of fragmented information into actions. The study further providers practical recommendations for improving disclosure practices and enabling more effective sustainability integration in cloud-dependent organisations.

Keywords: Cloud Sustainability; Cloud Deployment Decision; Data Centre Environmental Impacts; Transparency; Bounded Rationality; Mid-Sized SaaS Firms.

Executive Summary

Cloud computing is increasingly positioned as central to both digital transformation and sustainability transition. Hyperscale Cloud Service Providers (CSPs) such as Amazon, Microsoft, and Google frame cloud infrastructure as environmentally efficient, carbon-aware, and capable of supporting decarbonisation goals through renewable energy procurement and energy-efficient data centres. Simultaneously, the rapid expansion of artificial intelligence (AI), cloud service, and digital infrastructure has intensified concerns regarding rising electricity demand, carbon emissions, water consumption, and broader environmental impacts. Within this context, organisations are increasingly expected to make sustainability-informed infrastructure decisions, despite relying on cloud ecosystems they do not directly control.

This thesis examines cloud sustainability from the perspective of both providers and cloud customers. The key challenge addressed is not whether greener cloud computing is technically possible, but whether sustainability can meaningfully be operationalised by cloud customers under organisational, informational, and infrastructural constraints. While the technical literature demonstrates the potential for carbon-aware computing, workload shifting, renewable energy matching, and energy-efficient infrastructure, much less is known about how these possibilities translate into actionable deployment decision for cloud-dependent organisations.

The research was motivated by a clear gap in the literature: existing research on cloud sustainability and AI infrastructure largely focuses on technical optimisation models, energy-efficiency metrics, and infrastructure engineering. Studies frequently examine data centre metrics such as Power Usage Effectiveness (PUE), Carbon Usage Effectiveness (CUE), Water Usage Effectiveness (WUE), workload shifting, and carbon-aware Kubernetes orchestration. However, limited attention is given to the organisations' feasibility of acting upon sustainability information, particularly within mid-sized Software-as-a-Service (SaaS) firms dependent on hyperscale cloud infrastructures. Literature also demonstrates significant concern regarding transparency, comparability, and governance asymmetries within hyperscale cloud markets, yet empirical work connecting these issues to actual deployment decision-making remains limited.

To examine this broader gap, the thesis investigates two interconnected dimensions:

(1) The external information environment, namely what sustainability information CSPs publicly disclose and how actionable it is for customers
(2) The internal decision environment, namely the extent to which a cloud-depended SaaS firm can realistically integrate sustainability considerations into deployment decisions

The thesis addresses these dimensions through the following research questions (RQs):

RQ1(a) *How transparent and comparable are the sustainability reports and associated public environmental disclosures of Amazon, Microsoft, and Google, regarding data centre energy efficiency and carbon performance?*

RQ1(b) *What provider-level and selected region-level-information on data centre energy efficiency and carbon performance is available for AWS, Azure, and GCP, and how actionable is this information for customer decision-making?*

RQ2 *To what extent are sustainability-informed cloud deployment decisions possible for Neo4j within existing organisational, technical, and contractual constraints?*

To resolve these questions, the thesis adopts a qualitative exploratory case-study approach consisting of two analytical phases.

Phase 1 examines the external information environment through a document-based analysis of the sustainability disclosures of Amazon, Microsoft, and Google, with focus on Amazon Web

Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). The data corpus included annual sustainability reports, environmental disclosures, technical sustainability documentation, carbon tools, region-specific public datasets, and related public materials. The analysis was guided by the European Commission Joint Research Centre's (JRC) *Best Practices Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency*. The JRC framework was adopted not as a benchmark of performance, but an analytical lens for evaluating the adequacy of sustainability disclosure in terms of transparency, comparability, granularity, and decision relevance.

Phase 2 examined the internal decision environment through an exploratory organisation case study at Neo4j. Empirical material was gathered through consultations, internal organisational insights, and analysis of deployment-related consideration surrounding cloud infrastructure decision-making. The analysis was guided by bounded rationality theory, operationalised through four dimensions of feasibility:

(1) Informational feasibility
(2) Technical feasibility
(3) Organisational feasibility
(4) Governance feasibility

The literature review established four core themes underpinning the research: data centre sustainability metrics and their limitations; carbon-aware cloud deployment and workload scheduling; provider transparency and sustainability reporting; and the role and agency of cloud customers within hyperscale infrastructure ecosystems.

The review demonstrates that while metrics such as PUE, WUE, CUE and renewable energy claims remain dominant within industry practices, they are frequently criticised for aggregated reporting, inconsistent methodology, weak regional comparability, and limited decision-relevance for customers. Research on carbon-aware scheduling and workload shifting demonstrates theoretical potential for reducing emission by aligning workloads with lower-carbon electricity periods or regions. However, these approaches are proven highly conditional upon workload flexibility, infrastructure access, service level agreement (SLA) constraints, and real-time environmental data availability. At the same time, transparency literature consistently highlights the absence of standardised, granular, and regionally comparable environmental disclosure from hyperscale CSPs. Finally, multi-cloud and governance literature demonstrates that cloud customers possess limited but uneven agency, constrained by portability issues, interoperability challenges, vendor lock-in, engineering complexity, and provider dependencies.

Findings of Phase 1 reveal that CSP sustainability disclosure is extensive but only partially actionable for customer-side decision-making.

All three providers disclose meaningful sustainability information, including global efficiency indicators, renewable-energy procurement strategies, net-zero ambitions, and broad operational sustainability initiatives. However, substantial differences emerged regarding methodological clarity, reporting boundaries, temporal granularity, and region-specific transparency. The findings demonstrate that providers frequently disclose different sustainability objects, at different scales, and using different methodologies, making rigorous cross-provider comparison difficult. Among the providers analysed, Google providers the strongest public basis for region-level sustainability interpretation through carbon-free energy data and location-linked disclosure. Amazon and Microsoft both disclose substantial sustainability information and advanced ambitions, with Microsoft particularly demonstrating comparative sophistication in customer-facing tooling. Nevertheless, most disclosures remain aggregated at global, corporate, or fleet wide levels, and poorly structured around customer-relevant deployment units.

The findings therefore support the argument that current CSP sustainability reporting primarily functions as corporate signalling and governance communication; sustainability information exists, but not in forms sufficiently granular, standardised, or comparable to support systemic region-level optimisation by cloud customers.

Findings of Phase 2 demonstrated that sustainability-informed deployment decision at Neo4j are possible, but only within a bounded and conditional decision space.

Sustainability is recognised as increasingly important, yet deployment decisions remain primarily governed by cost, latency, service reliability, customer expectations, contractual obligations, compliance requirements, and operational stability. Opportunities for experimenting sustainability-informed optimisation are found in internal, development, canary, or non-production environments, where workload flexibility and organisational tolerance for are higher. New deployments also represent more feasible opportunities than migrations of existing production systems. By contrast, customer-facing production environments are found to be strongly constrained by latency sensitivity, governance requirements, SLA commitments, engineering complexity, and provider dependencies. These constraints limit the feasibility of acting upon sustainability considerations even where lower-carbon deployment options existed.

The findings therefore strongly support the bounded rationality framework adopted in the thesis. Neo4j struggle to optimise cloud sustainability in an unconstrained manner; sustainability becomes feasible only where informational, technical, organisational, and governance conditions align sufficiently to permit action without unacceptable operational trade-offs.

The thesis further identified a structural misalignment between information supply and decision demand. At **provider level**, sustainability disclosure is increasing in volume but not necessarily in cloud customer decision relevance. Metrics remain aggregated, methodologically inconsistent, and detached from operational deployment realities. At the **organisational level**, cloud customers are increasingly expected to contribute to sustainability goals while operating within infrastructures largely controlled by hyperscale providers. This creates a governance asymmetry in which providers control the production and disclosure of environmental information, while customers are expected to act upon it despite incomplete visibility and limited control.

Conceptually, this thesis contributes to the literature by linking provider transparency and customer feasibility into a single analytical framework. Existing research often treats technical optimisation and organisation decision-making separately. This thesis demonstrates that sustainability in cloud computing emerges from the interaction of infrastructure governance, disclosure systems, organisation bounded rationality, and hyperscale market structures. **Practically**, this thesis contributes insights for cloud customers, providers, and policymakers. For cloud customers, the findings demonstrate that sustainability-informed deployment requires selective and incremental integration rather than idealised optimisation. For providers, the findings highlight the growing importance of regionally granular, methodologically transparent, and customer-relevant sustainability disclosure. For policymakers, the research demonstrated that current sustainability expectations placed upon cloud customers may exceed the informational governance condition necessary to support meaningful action.

This thesis argue that sustainability in cloud computing is primarily constrained by the fragmented transparency, infrastructural asymmetries, governance limitations, and bounded organisational decision-making criteria. Progress therefore depends less on discovering new optimisation techniques and technical solutions, and more on enabling sustainability-informed action through improved transparency, standardisation, coordination, and governance structure.

Sustainability in the cloud will not be determined by who reports the most, but by who enables decisions where information, control, and responsibility converge. Future goals for improvement must focus on redesigning cloud governance, so that sustainability becomes a decision, not a disclaimer.

Table of Contents

List of Tables

Abbreviations

The following table informs the significance of various abbreviations or acronyms utilised throughout this thesis.

AI - Artificial Intelligence

API - Application Programming Interface

AWS - Amazon Web Services

CO_2 - Carbon dioxide

CFE - Carbon Free Energy

CoC - Code of Conduct

CPU - Central Processing Unit

CSP - Cloud Service Provider

CSRD - Corporate Sustainability Reporting Directive

CUE - Carbon Usage Effectiveness

DCCSE - Data Centre Cooling System Efficiency

DCeP - Data Centre Energy Productivity

DCiE - Data Centre Infrastructure Efficiency

E-Waste - Electronic Waste

EC - European Commission

EED - Energy Efficiency Directive

EN - European Norms

ERF - Energy Reuse Factor

ERE - Energy Reuse Effectiveness

ESG - Environmental, Social, and Governance

EU - European Union

FP&A - Financial Planning & Analysis

GCP - Google Cloud Platform

GEC - Green Energy Coefficient

GHG - Greenhouse Gas

GPU - Graphics Processing Unit

HPC - High-Performance Computing

HVAC - Heating, Ventilation, & Air Conditioning

IaaS - Infrastructure as a Service

IEA - International Energy Agency

IEC - International Electrotechnical Commission

ICT - Information and Communication Technology

IT - Information Technology

ITEU - IT Equipment Utilisation

ISO - International Organisation for Standardisation

JRC - Joint Research Centre

KPI - Key Performance Indicator

kWh - Kilowatt-Hours

LCA - Life-Cycle Analysis

PaaS - Platform as a Service

PDU - Power Distribution Unit

PPA - Power Purchase Agreements

PUE - Power Usage Effectiveness

REF - Renewable Energy Factor

RQ - Research Question

SaaS - Software as a Service

SLA - Service Level Agreements

SPUE - Server Power Usage Effectiveness

T-C-D - Transparency, Comparability & Decision Relevance

UPS - Uninterruptible Power Supply

WUE - Water Usage Effectiveness

1 Introduction: Cloud Computing, Energy Use and Regulation

1.1 Technical Contextualisation

The cloud market is experiencing rapid growth, only further amplified with the commercialisation of 5G, the rise of Artificial Intelligence (AI), hybrid and multi-cloud adoptions, and global data acceleration. The global cloud computing market, currently valued at US $943 billion, is expected to surpass $1 trillion by 2026 (Alexandre, 2025). As this growth continues, the material implications of cloud computing become increasingly visible, with energy consumption and associated carbon emissions of data centres emerging as critical concerns. This raises important questions regarding how environmental performance is measured, disclosed, and acted upon; especially for organisations that depend on cloud infrastructure without controlling it (Alexandre, 2025).

According to the National Institute of Standards and Technology (Mell & Grance, 2011, p.2), "cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources", e.g. networks, servers, storage, applications, and services. It underpins a growing share of digital services, for it permits enterprises to access computing resources hosted by third parties on the Internet, otherwise known as the cloud. Cloud computing is a utility-based model where computing resources are accessed and paid for on demand over the Internet, removing the need for physical infrastructure ownership. Data centres, hosting the necessary servers, networking, storage, and cooling systems, form the backbone of this service, processing and storing digital data. (Mell & Grance, 2011).

To understand the distribution of responsibility and control in cloud computing, it is necessary to distinguish between how cloud infrastructure is organised, i.e. deployment models, and how services are delivered, i.e. service models. According to Mell & Grance (2011), there are four deployment models, namely: private-; public-; community-; and hybrid- cloud provisioning. Private infrastructure is for a single organisation's exclusive use, often managed on premises. Public cloud is open use, managed by cloud providers. Community cloud serves multiple organisations with shared interests, managed by one or more of them. A hybrid cloud combines two or more of these types. The service models are three, namely: Infrastructure as a Service (IaaS); Platform as a Service (PaaS); and Software as a Service (SaaS). IaaS offers fundamental computing resources, e.g. storage and networking, allowing consumers to manage operating systems and deploy software. PaaS allows consumers to deploy applications using provider-supported tools without managing the underlying infrastructures. SaaS, most relevant to this research, allows consumers to use provider applications running on a third-party cloud infrastructure from multiple devices (Mell & Grance, 2011). Distinguishing between deployment models, service layers, and degrees of control enables further comprehension on how environmental responsibility and authority is distributed within cloud computing.

Within this structure, a small number of actors dominate the provision of cloud services. Borra (2024) defines Cloud Service Providers (CSPs) as information technology (IT) businesses that make available over the Internet computer power, data storage, and applications. Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) are identified as some of the most prominent CSPs. Together, these three providers dominate approximately 63% of global cloud infrastructure spend according to Synergy Research Group, with AWS holding ≈29% of global market shares, Azure ≈20% and GCP ≈13% as of Q3 2025 (Kumar, 2026). Choosing among these providers is a key infrastructure decision for organisations. For this research, these three CSPs define the empirical scope, due to their market dominance and relevance to the research case study.

The scale at which service providers operate introduces an additional layer of complexity: the global distribution of infrastructure. Large CSPs are often referred to as *hyperscale providers* due to the scale of their global distribution, operating across large-scale infrastructures known as *hyperscale data centres* (Bandhavi Sakhamuri, 2025). These distributed units are commonly referred to as *cloud regions*, each consisting of multiple data centres within a specific geographic area, designed to provide service with redundancy and low latency. The regions are embedded in specific national or regional electricity grids, with local energy systems shaping their environmental impact profile. The differences in electricity generation mixes, like the reliance on different energy sources, cause variations in the carbon intensity results. While organisations select regions based on factors like proximity, regulatory requirements, performance, and cost, the environmental impact is also a vital factor to keep in consideration, yet often gets overlooked (Bandhavi Sakhamuri, 2025). The scale and geographic spread of hyperscale infrastructure is substantial: AWS operates in 33 regions, Azure 60+, and GCP 40+ globally (Kumar, 2026). A present, the U.S. holds the largest number of data centres globally, responsible for 45% of global data centres electricity consumption; following the lead are China at 25% and Europe 15% (Kandasamy, 2025; Net Zero Insights, 2025). The strain on local power grids will only continue to intensify the more data centres are built, as these construction timelines outpace upgrades to grid and transmission infrastructures (Net Zero Insights, 2025; Safari et al., 2025).

CSPs play a dual role in this context; while enabling the rapid scaling of digital services and the growth of cloud-first business models, CSPs' data centres constitute energy-intensive infrastructure whose environmental performance has significant multi-system-level implications. *Hyperscale* CSPs often operate more energy-efficient facilities than smaller, on-premises data centres, due to economies of scale (Al-Sharafi et al., 2023). Nonetheless, their reliance on geographically distributed local electricity grids means that the environmental impact of cloud services remains highly dependent on regional energy mixes and regulatory contexts and thus cannot be understood solely in terms of efficiency gains (Safari et al., 2025).

The rapid expansion of cloud infrastructure, its dependence on regional energy systems, and the concentration of control among a small number of providers, create a complex environment in which environmental responsibility is not clearly actionable.

1.2 Defining the Problem

1.2.1 Energy Use and Environmental Impact in Data Centres

Cloud computing is a significant factor in streamlining business operations. Concomitantly, the rapid rise of cloud-based workloads raises serious environmental concerns regarding data centre energy use and carbon emissions. Although perceived as virtual and abstract, the cloud requires significant material and energy inputs.

Data centres and transformation networks account for approximately 1% of global energy-related greenhouse gas (GHG) emissions (Net Zero Insights, 2025). This level of consumption places substantial pressure on the energy infrastructure, raising many concerns surrounding the topics of grid stability and long-term cost efficiency; concerns only further amplified by the current instabilities of the global market, seeing electricity prices become increasingly volatile, influenced by geopolitical conflicts and fuel price fluctuations (Safari et al., 2025). Reducing embodied and operational carbon emissions, and energy usage, is essential for achieving the ambitions and imminent sustainability goals, like maintaining a 1.5-degree℃ limit on world temperature rise set by the Paris Agreement or reaching the 17 Sustainable Development Goals within the United Nations' 2030 Agenda (Hankendi et al., 2025).

The surge of data centres and their increasing local electricity demands contribute to slowing down the transition to clear energy, particularly in those countries with the highest concentrations of data centres. Moreover, modern AI systems require massive computational resource booths during training and inference phases, increasing the carbon- and energy-intensity of data centres, and posing a further bane for the environment. The demands of energy have driven some utilities to delay the shuttering fossil fuel power plants and have prompted proposals to revive others currently out of operation (Kandasamy, 2025).

The growing demand for energy poses a substantial environmental and climate challenge due to data centres' reliance on non-decarbonised power, resulting in carbon dioxide (CO_2) emissions, high water consumption, noise and light pollution, and significant electronic waste (Directorate-General for Energy, 2025; Net Zero Insights, 2025). These impacts are further elaborated below.

Electricity

According to Hankendi et al. (2025) data centres consume roughly 3% of the world's electricity resources to power this equipment. Due to increasing demands, further catalysed by the rise of AI, this figure is expected to reach 9% of the global electricity demand by 2030 (Safari et al., 2025). This reflects a projected rise to nearly 945 TWh within that timeline (IEA, 2025).

Highlighting the AI-driven exponential growth in carbon emissions, the International Energy Agency (IEA) provides global electricity consumption projections for AI operations reaching 800TWh by 2026. AI-optimised data centres can consume as much annual electricity as 100,000 households, or 350,000-40,000 electric vehicles (Net Zero Insights, 2025). Google's electricity usage alone for AI operations is projected to exceed its current data centre total electricity consumption by 2034 (Hankendi et al. 2025).

The construction timelines of new data centre projects outpace upgrades to grid and transmission infrastructure; hence several jurisdictions have instituted temporary moratoriums on development to process connection backlogs and evaluate the capacity of the local grids to hold these intensified strains of demand (Net Zero Insights, 2025).

Water

With large and new-AI focused data centres, alongside energy usage and carbon emissions is reported a vast consumption of water. These quotas are closely correlated: the more the data centre operations and AI requests consume larger amounts energy, typically they also consumer large amounts of water to cool processor chips. Data centre developers are increasingly tapping into freshwater surface and underground aquifers to cool these facilities, putting the nearby communities at risk. Freshwater is a critical resource; for perspective, only 3% of the Earth's water is fresh, 0.5% of which is safely accessible for human consumption.

Large data centres can consume up to 5 million gallons of water per day, or 1.8 billion per year, the equivalent use of a town populated by 10,000-50,000 people; a medium-sized data centre can consume roughly 110 million gallons of water per year, equivalent to the annual water usage of about 1,000 households (Pavlinich, 2026; Yanez-Barnuevo, 2025).

Water-cooled data centres consume less energy to cool systems and emit roughly 10% less CO_2 emissions than air-cooled ones, however they place immense stress on freshwater. A Cornell University study estimates that AI demand alone could require 4.2-6.6 billion cubic meters of fresh water annually by 2027. More generically this demand exacerbates the water crisis in water-sensitive regions, already suffering the effects of climate-related water shortages resources (Net Zero Insights, 2025). Data centres increase the threat of water insecurity. For instance, Google has seven active data centres in water-scare areas of the U.S. and was planning to build six more. In 2023, the state of Arizona was forced to revoke construction permits due to scarcity of groundwater in Maricopa County, where Meta resides one data centre, Microsoft two, and Google one with a second in development. (Pavlinich, 2026).

Air Pollution

Data centres, especially gas-powered, emit significant pollutants including nitrogen oxides, methane, volatile organic compounds, and fine particulate matter. These increase rates of respiratory diseases and cardiovascular conditions, elevating risks of cancer. Model studies for 2030 projections indicate the potential for U.S. data centres to cause approximately 600,000 asthma cases and 1,300 premature deaths (Pavlinich, 2026). Hence, many researchers call for health-informed AI, that factors in health costs to CO_2 emissions when deciding on data centre locations. That said, any location could pose problems to wildlife and the surrounding ecosystem (Kandasamy, 2025).

Noise and Light

Underestimated impacts are noise and light pollution, increasingly affecting nearby communities. The persistent hum of servers, networks, switches, routers, air conditioning, and industrial cooling fans are capable of generating sound levels exceeding 80-90 decibels (Net Zero Insights, 2025). Hyperscale facilities also require all-night lighting, disrupting the natural circadian rhythms of the body, including melatonin production and sleep-wake cycles. Chronic exposure can cause sleep disturbance, hearing loss, elevated stress hormone levels, hypertension, anxiety, and even cardiovascular risks, for an overall decreased quality of life (Net Zero Insights, 2025; Pavlinich, 2026). These negative externalities also follow nature, disrupting animal communication and forcing wildlife to change migration patterns (Net Zero Insights, 2025). AI data centres have been deemed potential sensory danger zones, due to the light and noise exceeding thresholds having fitness consequences for species (Kandasamy, 2025).

Electronic-Waste

While server longevity is shaped by engineering design, workload intensity, maintenance, and ambient operation conditions, and should typically require replacement every 3-5 years, a 2020 survey reported that 42% of IT managers replace servers every 2-3 years. These short upgrade cycles typical in high performing facilities produce in significant and unnecessary volumes of e-waste. Indeed, beyond wear-and-tear many operators choose upgraded equipment to capitalise on superior performance, energy efficiency, and capability of newer technologies, disregarding the impact of staggering electronic-waste volumes (Net Zero Insights, 2025).

According to Long et al. (2022), most of a data centre is made up of the physical infrastructure, i.e. power supply and temperature control systems, and IT equipment, i.e. servers arranged in racks, storage and networks. IT equipment must be constantly operating to maintain extraordinary dependability. Heating, Ventilation, & Air Conditioning (HVAC) in control of temperature, is primarily employed to avoid overheating created by the IT equipment during operation. Power supply facilities include Uninterruptible Power Supply (UPS), Power Distribution Unit (PDU) transformers, connectivity, and backup generations, etc. Data centre operators spend a lot of money on electricity. More effective measures are needed to establish a healthier link between human development and environmental preservation, and the establishment of effective energy efficiency evaluation methods and standards are required to achieve this objective.

The energy sector is at the heart of one of the most important technological revolutions. Affordable, reliable, and sustainable energy supply will be crucial determinants for countries to deliver the energy currently demanded at speed and scale (IEA, 2025). For instance, energy consumption can be dropped by up to 50% when renewable energy resources are incorporated; this efficiency can be further enhanced with dynamic resource provisioning and allocation of computing resources based on real-time demand. However, due to the lack of universal green compliance standards, it becomes difficult to establish consistent sustainability benchmarks for CSPs. Moreover, the need for an upfront investment in renewable energy infrastructures poses financial barriers to smaller enterprises (Beena et al., 2025).

If solutions are to be implemented and evaluated, it is vital to retrieve accurate energy consumption data on data centres, and the tech companies using them: a task far from simple (Kandasamy, 2025).

The sheer magnitude and complexity of the global sustainability challenge demand a fundamental reassessment of traditional economic mechanisms. Clearly, relying solely on the self-regulating forces of the free market, often symbolised by Adam Smith's invisible hand, will prove insufficient to drive the necessary systemic change. Market failures, such as externalities like pollution and the under-provision of public goods like a stable climate, mean that the private cost-benefit calculus does not align with the true societal costs.

Therefore, decisive and well-designed governmental policy intervention is not merely helpful, but a crucial and integrated component of any successful strategy to transition to a sustainable economy. This policy integration must span a wide spectrum of tools, including carbon pricing mechanisms, regulatory standards for resource efficiency, incentivizing sustainable innovation through research and development support and tax credits, and directing public investment toward green infrastructure. Without this active, intentional, and coordinated policy framework, the market will inevitably prioritise short-term profit over long-term planetary and social well-being, leaving the sustainability chasm unaddressed.

1.2.2 Regulations and Governance: The Current Challenges

As digitalisation accelerates, regulatory frameworks are increasingly shaping the sustainability expectations placed on data centre operators and cloud-dependent organisations. Governments and supranational bodies are moving beyond voluntary disclosure towards mandatory reporting and performance benchmarking, under means of Environmental, Social, and Governance (ESG) reporting. ESG puts stringent requirements on companies to reduce their impact, no longer as a desirable outcome, but as a business imperative. Indeed, some regulator frameworks respond with financial implications to inadequate environmental performance. Internationally, the European Union (EU) has been instrumental in promoting the global pledge to double the rate of energy efficiency improvements worldwide, from 2% to 4% by 2030. Nonetheless, there is still a widespread issue concerning data disclosure, making it difficult for policymakers to access and evaluate energy and sustainability performance (Kandasamy, 2025).

The Energy Efficiency Directive (EED) recast, Directive 2023/1791, introduces explicit obligation for large data centres to report energy consumption, water usage, and key performance indicators (KPIs). This includes the establishment of a European database on data centre performances, aimed at improving transparency and enabling cross-facility comparison. These measures are completed by the planned EU-wide data centre rating scheme, which seeks to standardise sustainability metrics and provide clearer signals to both regulators and market actors (European Commission, n.d.).

The EU Corporate Sustainability Reporting Directive (CSRD), Directive (EU) 2022/2464, mandates climate impact disclosure, and penalises non-compliance. Market pressures and investor scrutiny is emerging as equally significant, with environmental impacts being integrated within vendor selection processes and climate-related shareholder resolutions increasing. These conditions combined are transforming sustainability into a core business requirement, driving demand for technological innovations and solutions (Bandhavi Sakhamuri, 2025).

Such regulator developments directly respond to persistent data gaps. Existing research highlights that limited availability of granular, comparable data remains a central barrier to effective governance, concerning both policy evaluation and organisational decision-making (Aslan et al., 2025). This lack of transparency is a technical and structural issue, reflecting misaligned incentives between providers, regulators, and customers.

In parallel, voluntary initiatives such as the EU Code of Conduct (CoC) for Data Centres continue to prompt best practices in energy efficiency, including cooling optimisation, server utilisation, and renewable energy integration (Acton et al., 2025). The EU Best Practice Guidelines for the EU CoC on Data Centres Energy Efficiency is a regularly updated document to assist data centre operators identify measures to improve the energy efficiency of their data centres (Directorate-General for Energy, 2025). However, voluntary schemes alone have proved insufficient to ensure consistent adoption, leading to a gradual shift towards hybrid governance models, combining mandatory reporting with soft regulatory guidance.

Beyond Europe, broader governance challenges emerge from the material and systemic nature of digital infrastructure. Recent research emphasises that digital systems, including data centres, are often treated as immaterial in policy disclosure, obscuring their dependence on physical resources, energy systems, and global supply chains. This framing limits the effectiveness of existing governance approaches, which tend to prioritise efficient metrics while overlooking lifecycle impacts and embedded emissions (Pardalis et al., 2026).

Consequently, governance on data centre sustainability is evolving along three key dimensions:

(1) Transparency and standardisation, e.g. EED reporting and EU databases
(2) Performance benchmarking and accountability, e.g. rating schemes, standardised KPIs
(3) Lifecycle and system-level consideration, e.g. energy sourcing, infrastructure durability, supply chain impacts, etc.

Despite these advances, a critical gap remains; the current regulatory frameworks are primarily provider-focused, whereas sustainability responsibilities increasingly extend to cloud customers. This creates a misalignment between who reports environmental performance, i.e. CSPs, and who is expected to act on it, i.e. cloud users. This reinforced the central problem addressed in this thesis.

1.2.3 Organisational Reality for Mid-Sized SaaS Firms

Cloud computing has become a foundational enabler for organisations of all sized, particularly small and medium enterprises (SMEs) and mid-sized SaaS firms, exercised as cloud integration (Al-Sharafi et al., 2023). By externalising infrastructure, cloud adoption reduces upfront capital expenditure and enables scalable, flexible operations, supporting innovation and digital growth (Jaensson & Wiklund, 2022). Moreover, enterprises increasingly adopt multi-cloud or hybrid architectures to spread risk, meet compliance and locality demands, and avoid vendor lock-in effects. This gives smaller providers a chance to increase shares, especially for workloads less suited to the mega-cloud model (Alexandre, 2025).

This model fundamentally reshaped the distribution of control and responsibility. While cloud computing is often associated with efficiency gains and reduced on-premises infrastructure, environmental impacts do not disappear, they are displaced upstream into hyperscale data centres. These infrastructures remain highly energy- and resource-intensive (Ewim et al., 2023).

For mid-sized SaaS firms, this creates a structural tension. Organisations are increasingly required to meet sustainability targets and ESG reporting requirements, but they operate within infrastructural systems they neither own nor fully control. This manifests into a broader phenomenon identified in digital infrastructure research: users of digital systems are dependent on underlying material infrastructures whose environmental impacts are largely opaque and externally governed (Pardalis et al., 2026).

Multi-cloud and hybrid strategies are sometimes adopted to address operational concerns such as availability, latency, reliability, and vendor lock-in. However, these approaches introduce additional complexity, including portability challenges, interoperability issues, and increased governance overhead (Jaensson & Wiklund, 2022). Furthermore, the increasing integration of AI and data-intensive workloads amplifies these challenges. AI-driven services significantly increase demand for computational resources, thereby intensifying reliance on energy-intensive data centre operations (Aslan et al., 2025).

Ultimately, mid-sized SaaS organisations are highly constrained to operating decisions outside of their domain, not in control of the emissions that reflecting their performances. Indeed, these firms operate under the condition of *bounded agency*: this mismatch between responsibility and capability that lies at the core of the research problems addressed in this thesis.

1.2.4 Research Gap and Problem Statement

Research Gap

As cloud computing continues to expand, the environmental implications of data centre energy use have become increasingly salient for providers, regulators, and cloud-dependent organisations. However, while infrastructure level efficiency improvements are documented, it remains unclear how customer-side deployment decisions can meaningfully influence emission outcomes within these structural constraints.

Problem Statement

The core problem addressed in this thesis is the gap between infrastructure-level sustainability performance and customer-level decision capacity. While hyperscale providers increasingly disclose environmental performance indicators, these are often aggregated, methodologically heterogeneous, and limited in regional granularity (Hankendi et al., 2025). Mid-sized SaaS firms are therefore expected to reduce their operational emissions, while relying on infrastructure whose environmental characteristics they neither control nor observe fully. This creates a structural decision-making challenge that remains insufficiently examined in the literature.

1.3 Case Study Contextualisation

Neo4j (n.d.) is the world-leading graph data platform, providing graph databases for uncovering relationships in data, powering AI, fraud detection, recommendations, and knowledge graphs.

The scope of this research places specific focus on the development of Aura, Neo4j's cloud-native graph database-as-a-service, and other internal development workloads.

This thesis examines energy efficiency-informed cloud deployment decisions through the case of Neo4j, whose cloud-native offering provides a relevant setting for studying sustainability decision-making in a cloud-dependent SaaS context. Neo4j describes itself as a graph data platform, and more recently, a graph intelligence platform, positioning graph technology as particularly valuable for connected data use cases such as knowledge graphs, fraud detection, recommendations, analytics, and AI-enabled applications. (Neo4j, n.d.-a; Hasbe, 2026). Neo4j is especially relevant to this research because it operates in a digital environment where cloud infrastructure is strategically important, but where the organisation does not directly control the underlying data centre operations criteria (Neo4j, Sustainability Specialist, personal communication, January 28, 2026). This places it in a structurally similar position to many other mid-sized cloud dependent firms seeking to align business with sustainability ambitions.

The case study focused specifically on Neo4j Aura. Neo4j describes Aura as a fully managed, cloud-native graph service, designed to let customers build and scale graph-powered applications without managing the underlying infrastructure themselves (Neo4j, n.d.-b). Aura is particularly relevant because it is delivered through the very cloud model at the centre of this study: it runs as a managed service on major hyperscale cloud platforms and is offered through their marketplaces, including AWS, Azure, and GCP (Neo4j, n.d.-c, n.d.-d). This makes Aura an analytically strong focal point for examining how environmental information disclosure by CSPs may or may not be translated into deployment decisions by a cloud customer operating downstream in the value chain.

Neo4j is also a suitable case because its product context sits at the interaction of cloud computing, data intensive services, and AI-related growth, all of which are increasingly putting pressure on digital infrastructure. Neo4j explicitly positions graph technology and knowledge graphs as enabling more accurate and explainable AI applications, making its cloud operations relevant not only as generic SaaS infrastructure, but as part of a broader shift towards more computationally demanding digital services (Hasbe, 2026; Neo4j n.d.-e). At the same time, this does not make Neo4j a hyperscale infrastructure operator; rather it remains dependent on cloud providers for the environmental performance on the infrastructure on which Aura runs. This dependency is precisely what makes the case analytically useful: Neo4j occupies the position of a cloud customer that makes deployment choices under conditions of indirect control, partial information, and competing business constraints criteria (Neo4j, Sustainability Specialist, personal communication, January 28, 2026).

A further reason for selecting Neo4j is that the case combines practical relevance with analytical generalisability. Neo4j's cloud operations provide meaningful empirical material, while representing a broader class of firms that rely heavily on hyperscale infrastructure without shaping its design. Aura's availability across the leading cloud marketplaces also reflects the commercial and operational realities discussed elsewhere in this thesis: cloud deployment choices are influenced not only by technical feasibility, but also by procurement models, marketplace arrangements, customer location, and service-level considerations (Neo4j, n.d.-c, n.d.-d). Neo4j's reliance on multiple hyperscale providers reflects both customer preference and commercial considerations, including marketplace agreements and scalability requirements (Neo4j, Sustainability Specialist, personal communication, January 28, 2026). As such, Neo4j offers a case through which to explore how sustainability-related information may enter decision-making in practice, and where structural limits remain.

Accordingly, the empirical scope of this study is restricted to the three primary hyperscale CSPs most relevant to Aura, i.e. AWS, Azure, and GCP, and to Neo4j's internal and development-oriented cloud use, rather than customer-facing production optimisation in general. This focus is deliberate. Internal and development environments offer a more realistic space in which sustainability-oriented deployment choices can be examined without overstating the firms' degree of control over production infrastructure. In this sense, Neo4j is not treated as an exceptional case, but as an analytical valuable one: an organisation navigating cloud growth, AI-relevant workloads, and increasing sustainability expectations while remaining downstream from the infrastructure decisions that shape much of its environmental footprint.

The Neo4j case therefore provides a concrete setting in which to examine whether region-level differences in cloud sustainability performance can be identified, interpreted, and incorporated into organisational decision-making under real-world operational constraints.

1.4 Aim and Research Questions (RQs)

The aim of this research is to explore how more sustainable cloud deployment decisions can inform mid-sized SaaS companies, integrating energy efficiency and environmental transparency considerations into regional deployment choices. Specifically, this study seeks to advise internal decision-making at Neo4j, by supporting sustainability aware trade-offs between energy efficiency and operational constraints.

While prior research has primarily focused on infrastructure efficiency or carbon-aware computing at system level, limited attention has been paid to the decision-making constraints faced by cloud-dependent customers operating under incomplete information.

The research is guided by the following questions:

RQ1(a) *How transparent and comparable are the sustainability reports and associated public environmental disclosures of Amazon, Microsoft, and Google, regarding data centre energy efficiency and carbon performance?*

RQ1(b) *What provider-level and selected region-level-information on data centre energy efficiency and carbon performance is available for AWS, Azure, and GCP, and how actionable is this information for customer decision-making?*

RQ2 *To what extent are sustainability-informed cloud deployment decisions possible for Neo4j within existing organisational, technical, and contractual constraints?*

1.5 Scope and Delimitations

The scope of this study is limited to selected energy efficiency related indicators and carbon intensity measures associated with the public cloud services of the three indicated hyperscale CSPs: Amazon, Microsoft, and Google. The analysis focuses on publicly disclosed metrics such as renewable energy sourcing claims, usage effectiveness criteria, and regional grid carbon intensity, and any other pertinent data available to cloud customers.

This study does not attempt to calculate actual emissions reduction, nor does it assess embodied emissions or full lifecycle impacts. Financial cost optimisation, customer data sovereignty, and on-premises deployment strategies are considered purely as situational factors given the constraint they place on cloud deployment decisions.

The empirical focus is restricted to Neo4j's internal and development environments, where there is greater flexibility in workload placement compared to customer-facing production.

1.6 Target Audience

This thesis is primarily addressed to three interrelated audiences:

(1) Mid-sides SaaS companies and cloud-dependent organisations

Particular focus is placed on companies and organisations seeking to align operational practices with sustainability objects; while operating under indirect control over cloud infrastructure, the study provides practical insights into how sustainability information from cloud providers can be interpreted and integrated into deployment decision-making.

(2) Academic research in environmental management and digital sustainability

The research addresses an identified gap in literature by linking provider-level sustainable disclosure with customer side decision feasibility.

(3) Policymakers and standard-setting bodies

Primary direction is towards bodies concerned with improving transparency, comparability, and accountability in data centre sustainability reporting; by highlighting the limitation of current disclosure practices from a customer perspective, the atheist providers insight into how reporting frameworks can better align with real-world decision-making needs.

1.7 Disposition

Chapter 1 introduces the research problem outlining the rapid expansion of cloud computing and its environmental implications, highlighting governance challenges around the concentration of cloud infrastructure among a small number of providers. This establishes the core research gap and presents Neo4j case as the empirical setting for this research.

Chapter 2 reviews the literature on data centre sustainability, carbon-aware computing, provider transparent, and the role of cloud customers, identifying a gap between potentials for technical optimisation and organisational feasibility. A conceptual framework is developed combining a reporting lens for analysing provider disclosures, and a bounded rationality lens for understanding constrained decision-making.

Chapter 3 outlines the methodological approach, adopting a qualitative, interpretative design with an instrumental case study, explaining the two-phased research design, data sources and collection, and analytical strategy.

Chapter 4 & 5 present the Phases 1 & 2 of the empirical analysis; Phase 1 assesses the CSPs' sustainability disclosures using a coding framework, and Phase 2 examines Neo4j's internal decision-making context.

Chapter 6 integrates findings from both phases with the literature and conceptual framework critically reflecting on the limits and constraints emerged.

Chapter 7 summarises the research, providers targeted recommendations for cloud customers, providers and policymakers, identifying practical pathways for improvement of the current narratives and directions for further research; it ultimately answers the RQs.

2 Literature Foundations and Conceptual Framework

2.1 Literature Review

The literature review covers digital sustainability, cloud infrastructure, AI energy demand, green computing, and information systems governance, organised into four themes: *data centre sustainability metrics*, *energy-efficient cloud computing*, *provider transparency in reporting*, and *customer role in sustainability*. While these areas are emerging and debated, peer-reviewed literature on region-level cloud deployment and organisational feasibility is scarce. Therefore, relevant industry reports and policy documents, i.e. grey literature, are used to supplement academic sources.

2.1.1 Data Centre Sustainability and Energy Efficiency Metrics

According to Long et al. (2022), energy efficiency metrics are the measurement standards for energy efficiency evaluation that operators use to understand inefficiencies and improve their performance through analysing of core parameters. These mainly include cooling, performance, sustainability, air management, security, and financial impact, e.g. carbon and oil price changes. Energy efficiency has been a central concern in data centre research for a long while, with evaluating standards developed primarily by industry, and not academics.

Data centre energy consumption comprises the electricity required to operate all the IT equipment (e.g. servers, storage devices and network equipment) and supporting power delivery and cooling system components. This consumption is measured using sustainability metrics like Power Usage Effectiveness (PUE), Carbon Usage Effectiveness (CUE), and Water Usage Effectiveness (WUE), introduced by the Green Grid Association (Nlyte Software, 2023). Additional indicators such as Data Centre Energy Productivity (DCeP), Green Energy Coefficient (GEC), and Energy Reuse Factor (ERF), attempt to link efficiency to broader suitability outcomes (Viegand Maagøe, 2022). While these metrics simplify complex systems into manageable indicators, their usefulness for customer decision-making remains limited.

Metrics differ in granularity: coarse metrics, e.g. PUE, capture overall facility efficiency, while finer metrics aim to isolate specific components, though often without directly reflecting total performance (Long et al., 2022); see the glossary in Appendix A for further details. Despite known limitations, PUE remains the dominant industry standard.

PUE measures the ratio of total facility energy to IT equipment energy, where values can range from 1.0 to infinite. Values closer to 1.0 indicate higher efficiency and values supervisor to 2.0 are considered inefficient. A perfect PUE means that almost all electricity that enters the data centre is used for computing rather than the overhead infrastructure. Indeed, lower PUE signifies a more efficient use of resources that support IT space, thereby reducing energy wastages and operating cost (Nlyte Software, 2023). Hyperscale CSPs typically achieve lower PUE values due to economies of scale and optimised infrastructures (Hankendi et al., 2025; Long et al., 2022; Safari et al., 2025). However, the literature consistently criticises PUE as an incomplete indicator, for it does not account for carbon intensity, workload utilisation, or lifecycle impacts; and it is often reported as aggregated averages, masking regional variation (Hankendi et al., 2025; Long et al., 2022; Safari et al., 2025). Inconsistent calculation methods further reduce comparability, with observed variations of up to ≈20% across facilities (Hankendi et al., 2025). While virtualisation improves data centre energy efficiency, scaling infrastructure beyond load reduction can conversely decrease efficiency and worsen PUE scores. This reveals a critical flaw: PUE is load-dependent and only measures physical infrastructure efficiency, ignoring IT computing power. Consequently, it fails to reflect the benefits of reduced energy consumption from more efficient IT devices, making it an inadequate energy efficiency metric (Long et al., 2022). Most importantly, PUE was not originally designed for cross-data centre comparison, yet it is consistently used as such in practice (Long et al., 2022). This critique is further reinforced by the argument that researchers are often forced nonetheless to rely on average PUE values precisely because operators do not disclose more granular operational data. Without facilities-specific power use or utilisation data, internal assessments remain approximate rather than decision ready (Hankendi et al., 2025). As such, PUE remains the dominant benchmark.

Similar limitations apply to WUE, which measures annual total site water usage relative to IT equipment energy use, i.e. the overhead or efficiency of a facility's usage of water (Nlyte Software, 2023). While highlighting the environmental relevance of water use and concerns on water scarcity, WUE is typically reported in aggregated form and lacks sufficient contextual detail to support meaningful comparison across locations or workloads (de Vries-Gao, 2026). However, recent work on data centre carbon and water footprints shows that even among major operators, disclosure is inconsistent and incomplete, with few firms providing enough detail to connect total water consumption to individual facilities or AI-related workloads. The consequence is that WUE can signal that water matters but rarely enables meaningful customer comparison across regions (de Vries-Gao, 2026).

CUE, introduced as complementary to PUE, helps determine the amount of GHG emissions produced per unit of IT energy consumed within a data centre, providing an effective way to measure the carbon footprint of operations. It is calculated based on the source of energy and the quantity consumed, expressed as the amount of CO_2 emitted per unit of energy produced (Nlyte Software, 2023). This is particularly relevant to initiatives signalling broader shifts from traditional energy efficiency metrics, such as operators committing to CO_2-neutral targets and 100% renewable energy sources, supported by clean energy procurement strategies like regional Power Purchase Agreements (PPAs). Safari et al. (2025) presents the case of Microsoft's long-term PPA with the Three Mile Island nuclear facility to secure a reliable energy supply. Moreover, thermal re-use systems like district heating with server exhaust and aquifer thermal storage, and water-efficient cooling technologies, are all becoming measures to mitigate further environmental impacts. These innovative solutions represent an integration of sustainability and reliance planning into the way data centres are designed and managed.

More advanced approaches like multi-metric frameworks attempt to address these limitations. For instance, Data Centre Energy Efficiency (DCEE) combines multiple indicators into a weighted assessment, namely summing IT Equipment Utilisation (ITEU), CUE, WUE, Data Centre Cooling System Efficiency (DCCSE), Energy Reuse Effectiveness (ERE), and PUE (Long et al., 2022). This approach such as many other innovative ideas, like Server Power Usage Effectiveness (SPUE), face significant barriers (Safari et al., 2025). Research may even advocate for these unified systematic frameworks to guide section and application in line with evolving infrastructure models; however, these remain still under development, and their implementation faces challenges regarding standardisation and integration, limited data availability, and variation in data centre design (Long et al., 2022; Safari et al., 2025). Nonetheless not all researchers seek a single solution: metrics should be problem-specific, so the optimal metrics tend to vary across data centre scenarios; the main issue is misuse. Long et al. (2022) agree that while a universal standard is lacking, limiting energy efficiency evaluation, the establishment of an all-encompassing metric would be impossible.

Overall, the literature highlights a fragmented landscape of metrics, each capturing only partial aspects of performance, typically excluding software-related inefficiencies, dynamic workloads, and grid-level CO_2 intensity. This fragmentation limits their ability to support consistent evaluation and comparison (Bandhavi Sakhamuri, 2025; de Vries-Gao, 2026; Hankendi et al., 2025; Lange et al., 2020; Long et al., 2022; Safari et al., 2025). Importantly, efficiency improvements do not translate into sustainability gains. Digitalisation can increase total energy demand through rebound effects and growth dynamics, meaning that more efficient infrastructure may still result in higher absolute emissions (Aslan et al., 2025; Lange et al., 2020). This is particularly evident with the rise of AI workloads, which significantly increase computational demand (Bandhavi Sakhamuri, 2025).

Consequently, while energy efficiency metrics remain essential for compliance and monitoring, they provide only indicative insights. Their limitations in granularity, comparability, and contextual relevance constrain their usefulness for cloud customers seeking to make sustainability-informed deployment decisions (Bandhavi Sakhamuri, 2025; de Vries-Gao, 2026; Hankendi et al., 2025; Lange et al., 2020; Long et al., 2022; Safari et al., 2025).

More broadly, the literature suggests that metric-based efficiency should not be conflated with sustainability performance. Data centres may become more efficient while total energy demand still rises because workloads, traffic, and computational intensity grow faster than efficiency gains (Lange et al., 2020). Bandhavi Sakhamuri (2025) shows that data centre electricity demand had remained relatively stable despite workload growth up to around 2022 because of efficiency improvement, but that the rapid rise of AI workloads is changing this pattern. Lange et al. (2020) argue that digitalisation should be understood through four interacting effects: direct energy use of information and communication technology (ICT); efficiency improvements; economic growth effects; and sectoral change. Their core conclusion is that efficiency gains can be offset by rebound effects and growth dynamics, so that digitalisation has so far tended to increase rather than decrease total energy demand (Aslan et al., 2025; Lange et al., 2020). Indeed, lower PUE or better cooling does not necessarily mean lower total environmental impact, especially in the context of AI-driven demand growth (Lange et al., 2020).

For this thesis, such metrics are treated as indicative rather than determinative. Their functions, strengths and limitations are summarised in Table 1.

Table 1. Key Data Centre Sustainability Metrics

Metric	***Measures***	***Formula***	***Strengths***	***Limitations***
PUE	Overall energy efficiency of data centre	Tot facility energy / IT equipment energy	Simple, widely used, comparable within same system	Ignores carbon intensity; aggregated; inconsistent methods; not designed for cross-site comparisons
CUE	Carbon emissions per unit of IT energy	CO_2 emissions / IT energy	Links energy use to emissions	Dependent on data availability; not widely standardised
WUE	Water consumption relative to IT energy	Water usage (litres) / IT energy (kWh)	Highlights environmental impact of cooling	Aggregated reporting; lacks regional and contextual comparability
DCeP	Useful work output relative to energy use	Output / energy consumption	Attempts to link efficiency to productivity	Difficult to define useful work; not widely adopted
GEC	Share of renewable energy use	Renewable energy / tot energy	Connects energy sourcing to sustainability	Not standardised across providers
ERF	Amount of reused energy	Reused energy / tot energy	Captures circular energy use	Limited applicability across sites
DCiE	Inverse of PUE, efficiency of infrastructure	IT energy / tot energy	Alternative expression of efficiency	Same limitations as PUE
SPUE	Energy efficiency at server level	Server-level metric	More granulate than PUE	Not standardised; limited adoption

DCEE	Combined efficiency across multiple indicators	Weighted sum of ITEU, CUE, WUE, DCCSE, ERE, PUE	Most comprehensive evaluation	Complex; lacks standardisation and data availability

Source: Author's elaboration, based on literature from Bandhavi Sakhamuri (2025); de Vries-Gao (2026); Hankendi et al. (2025); Lange et al. (2020); Long et al. (2022); Safari et al. (2025).

2.1.2 Energy-Efficient and Carbon Aware-Cloud Deployment

A second body of literature moves from status infrastructure metrics toward dynamic deployment decisions. This work argues that environmental performance depends not only on how efficient a data centre is, but also on where, when, and how workloads are executed. The core concept is carbon-aware cloud computing, which links workload placement and scheduling to temporal geographic variation in carbon intensity.

Bandhavi Sakhamuri (2025) outlines that modern cloud infrastructures, operating 24/7, often suffer from low average server utilisation and suboptimal resource use. Peak demand is the only time utilisation is high. The typical energy distribution in modern data centres: servers (40-45%), cooling (35-40%), storage (10-15%), and network (5-10%). This underutilisation creates unavoidable inefficiencies.

Identical workloads can have substantially different emissions outcomes, i.e. carbon intensity, depending on local energy mixed and time-based grid fluctuations (Beena et al., 2025). This variability is highlighted by Bandhavi Sakhamuri (2025), who provides examples such as the contracts between the peak generation times for solar midday and wind overnight. Indeed, data centres especially with high renewables grids present opportunities for carbon-aware workload scheduling. The article also argues that traditional Kubernetes scheduling privileges performance and resource efficiency, while generally neglecting environmental impact. This omission creates scope for carbon-aware scheduling approaches that use machine learning and carbon-intensity signals to inform workload placement, autoscaling, and temporal shifting.

Scheduling refers to the decision-making process of allocating resources to tasks over given time periods, traditionally modelled as an optimisation problem for the reduction of makespan and manufacturing costs, subject to constraints disruption (Mencaroni et al., 2025). Traditional Kubernetes schedulers are responsible for placing new pods (e.g. applications) onto worker nodes (e.g. machines), typically finding the best fit based on assessment of resources needed, which could be central processing units (CPUs), memory, constraints, or policies. However, these traditional kube-schedulers lack environmental awareness because they prioritise performance and resource efficiency, despite the significant geographical and temporal variations in carbon intensity. To address this, Bandhavi Sakhamuri (2025) introduces a novel framework for carbon-aware Kubernetes scheduling. This framework uses AI to optimise workload placement based on environmental factors, successfully reducing an organisation's carbon footprint without compromising performance (i.e. how fast or efficiently the system executes tasks) or reliability (i.e. how consistently the system performs without failing).

Carbon-aware scheduling introduces three time-dependent elements: the grid's carbon intensity, the availability of on-site renewable electricity, and the power requirements of individual jobs. From a managerial perspective, this is a cost-effective measure, for it acts on timing, and does not rely on additional machine capacity or infrastructure. By leveraging information that should be easily accessible within the industry, such as grid carbon intensity, renewable generation forecast, and job-specific power requirements, it can be integrated into existing planning systems with limited disruption (Mencaroni et al., 2025). The literature confirms how new arrangements for task scheduling and resource allocation are essential to improve the energy efficiency and sustainability of cloud computing, with particular relevance noting a major influence in task-sorting techniques i.e. methods and rules used to order, prioritise, and arrange a queue of tasks before they are assigned to resources (e.g. central processing units, virtual machines, or workers) for execution (Beena et al., 2025). Tripathi et al. (2023) provide a particularly useful contribution by making the trade-off structure visible. Their workload-shifting framework explicitly incorporates power constraints, Service Level Agreement (SLA) constraints, carbon-emission constraints, and routing constraints when evaluating whether workloads can be moved to lower-carbon periods or locations. The key insight is not merely that carbon-aware shifting can reduce emissions, but that such shifting is only feasible when workload type and operational requirements permit it. Batch and storage-intensive workloads, i.e. computing tasks designed to process vast amounts of data without human intervention. Because these tasks prioritise total throughput, i.e. completing a large volume of work, over low-latency responsiveness, these workloads appear more movable than interactive or latency-sensitive ones.

Presented by Beena et al. (2025), is green cloud computing as an emerging response to these challenges around traditional Kubernetes deployments and serves as a major step towards implementing carbon-aware scheduling. It is aimed at minimising carbon footprint in cloud services. It integrates energy-efficient technologies, intelligent resource management, and sustainable best practices to optimise cloud operations, but also lower operational costs as it reduces power consumption and maintains performance standards. Strategies like dynamic resource allocation, efficient workload scheduling, and the selection of renewable energy sources all play a vital role in attempting to make cloud computing more sustainable. The researcher's contribution proposes an adaptive and real-time task scheduling framework based on green cloud computing principles, with which CSPs and enterprises can dynamically adjust task scheduling using real-time carbon intensity data. Many existing proposed solutions lack real-time carbon intensity data, limiting the ability to adapt scheduling decisions accordingly. Moreover, geographical variation in energy consumption and fluctuating carbon content is overlooked. By incorporating live carbon intensity data from the electricity maps Application Programming Interface (API), Beena et al. (2025) addresses these gaps.

At the same time, the literature also warns against assuming that technical optimisation alone is sufficient. Carbon-aware scheduling is often modelled under idealised conditions with access to live carbon data, controllable workloads, and relatively clean migration pathways (Bandhavi Sakhamuri, 2025; Beena et al., 2025). Many proposed solutions remain theoretically appealing but operationally distant from the everyday realities of mid-sized SaaS firms facing latency requirements, compliance and governance issues, and contractual commitments. This suggests that the relevant question is not whether carbon-aware deployment exists, but rather under what condition does it become feasible without unacceptable trade-offs like power cost or SLAs (Tripathi et al., 2023). The wider digital sustainability literature deepens on this critique. Lange et al. (2020) shows that efficiency improvements often trigger rebound effects, increased demand and complementary growth rather than absolute reduction in energy use.

Carbon-aware cloud deployment must be understood as a bounded intervention within a system still shaped by growth in digital demand. This does not make workload shifting irrelevant, but it means that deployment optimisation should be assessed as a conditional emission-reduction strategy, not a complete solution.

2.1.3 Provider Transparency and Sustainability Reporting

A third body of literature focuses on what knowledge cloud customers can practically obtain. Here the central issue is not technical optimisation but transparency. Hankendi et al. (2025)

argue that the sustainability challenges of AI and data centres cannot be meaningfully addressed without greater data transparency. They note that operators often treat operational data as proprietary, thereby restricting access to the information needed to evaluate energy use, infrastructure efficiency, hardware impacts, and real-world emissions patterns.

This matters because provider sustainability reports, while increasingly common, do not necessarily provide decision-relevant information. Hankendi et al. (2025) also pointed out that aggregate disclosure on Scope 1-3 emission, renewable energy procurements, and data centre efficiency metrics is often too coarse for rigorous academic analysis or customer-side comparison. Important details such as per-inference energy use, real-time power draw, embodied carbon, or standardised location-level reporting are frequently absent. This problem is reinforced in de Vries-Gao's (2026) analysis, which finds that no major operator reports AI-specific environmental metrics and that even data centre level disclosure is incomplete and inconsistent across firms. As a result, estimates of AI carbon footprints must be presumed through wider data centre metrics. According to Hankendi et al. (2025) achieving carbon-neutral AI and in general data centre running is impossible without first addressing these transparency bottlenecks. Stakeholders cannot accurately assess where efficiency gains are possible or where regulatory interventions are necessary if there is no clarity on how the resources are consumed.

The literature also makes clear that transparency is not only about volume of disclosure but about granularity, comparability, and methodological clarity. Hankendi et al. (2025) illustrated how industry leaders like Amazon, Google, and Meta, now release high-level sustainability reports, disclosing aggregate figures of Scope 1-3 emissions, renewable energy procurement, and data centre efficiency metrics. However, these reports are insufficient for rigorous academic analyses. De Vries-Gao (2026) shows that many firms disclose different combinations of electricity consumption, Scope 2 emission, and water metrics, often organisation-wide and not at data centre or regional level. This makes cross-provider comparison difficult and weakens the ability of customers to relate corporate sustainability narratives to actual regional deployment choices. In this sense, the issue is not merely missing information but the presence of information that is structurally misaligned with customer decision needs (de Vries-Gao, 2026; Safari et al., 2025). By lacking granularity, e.g. per-inference energy use or real-time power draw, they omit embodied carbon metrics for AI hardware. At the same time, these major tech companies are committed to reduce their carbon emission to net-zero by 2030, which seems out of reach considering the increasing burden of AI technology (Hankendi et al., 2025).

Hankendi et al. (2025) go further by identifying distinct categories of data gaps: infrastructure data, operational data, hardware-vendor data, and big-tech disclosures. They argue that the absence of such data slows innovation and policy development because researchers are forced to rely on industry averages rather than real-world observations. The study shows that several reasons are used to justify nondisclosure, including cybersecurity risks, commercial secrecy, and lack of standardisation. These factors help explain why provider transparency remains partial even as environmental claims become more prominent.

Global initiatives promote transparency in computing infrastructure for efficiency and sustainability. In Europe, organisations have established comprehensive frameworks to ensure data transparency and ethical data usage in high-performance computing (HPC) and research initiatives. The EuroHPC Joint Undertaking is a leading framework in emphasising accessibility, transparency, and ethics in supercomputing. Germany's Gauss Alliance, the National Grid Initiative for Germany, and the National HPC centres reinforce this with transparent data management and rigorous accountability, supporting cutting-edge research. In Japan, RIKEN ensures data transparency and security in research (e.g. HOKUSAI-SR), guided by its Information Security Policy. North America sees industry-academia efforts: Meta's Open Compute Project (OCP) provides open server specifications, and Microsoft's Project Natick demonstrates superior efficiency in transparently operated underwater data centres (PUE 1.07

vs. 1.25). Singapore's government-led Green Data Centre Innovation Program mandates disclosure of energy metrics, yet over half of Asian economies reportedly lack reliable power data (Hankendi et al., 2025).

Implementing data transparency practices is a mutually beneficial deal: it sees researchers access detailed data to improve quality and impact of sustainability research, and operators receive academic output to improve their efficiency and achieve their goals. Additionally, it allows for policymakers to implement data-driven public policies, and customers to keep companies accountable with their environmental standards (Hankendi et al., 2025).

For this thesis, the implication is direct: cloud provider reporting should not be treated as neutral evidence; rather as a form of selective disclosure shaped by corporate incentives, methodological choices, and reputational considerations. This is why disclosure should be addressed through dimensions of transparency, comparability, and decision-relevance for customers.

2.1.4 The Role of Cloud Customers in Sustainability Transition

The final theme concerns the position of cloud customers themselves. Existing literature shows that cloud customers are neither fully powerless, nor fully sovereign. They do not control data centre design, energy procurements, or disclosure practices, but they may influence outcomes through provider choice, region selection, workload timing, and internal governance.

Limitations in transparency and comparability directly shape the decision space available to cloud customers, forming a critical basis for the analytical framework developed in this thesis (Hankendi et al., 2025).

This customer role is especially visible in the multi-cloud literature. Jaensson & Wiklund (2022) show that organisations using multiple CSPs face challenges not only of portability and interoperability, but also of sustainability, economy, talent management, and legal conditions. Their case material is useful here because it shows that even when a firm has reasons to use multiple providers, practical flexibility is limited by vendor lock-in, migration costs, uneven competencies, and the need to maintain variants across platforms. In other words, customer agency exists but are costly and unevenly distributed.

Alonso et al. (2023) similarly argue that multi-cloud application development is still characterised by conceptual ambiguity and unresolved practical challenges. Their review suggests that multi-cloud approaches can improve flexibility, fault tolerance, and optimisation potential, but also increase complexity for developers and operators across application lifecycle. They explicitly highlight unresolved issues around portability, runtime operations, and security.

This leads to a broader interpretative point. The literature can be grouped into three positions that do not simply disagree but rather illuminate different parts of the same problem. These are the following:

- An *optimistic position* assumes that technical optimisation, carbon-aware scheduling, and cleaner regions can materially improve outcomes if the right data and tools are available
- A *conditional position* argues that such gains depend on governance conditions, reporting quality, workload flexibility, and organisational capacity
- A *critical position* stresses that structural limitations dominate, because growth in digital demand, rebound effects, incomplete disclosure, and indirect control constraint what customers can achieve

This thesis takes the *conditional position*, because it best reflects the empirical reality of cloud-dependent firms such as Neo4j.

2.2 Conceptual Framework: Decision-Making Under Constraint in Cloud Sustainability

Technical studies often focus on optimisation, while governance studies highlight information asymmetry and limited customer agency. A framework linking technical possibility to organisational feasibility is clearly needed.

2.2.1 Reporting Framework: JRC Best Practices Guidelines as an Analytical Tool

To guide the analysis of RQ1(a) and RQ1(b), this thesis adopts the EC's Joint Research Centre's (JRC) *2025 Best Practices Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency* as its primary conceptual and analytical reference point. The JRC document is appropriate for this purpose because it is not merely a technical operations guide; it explicitly provides a common terminology and framework reference for describing energy-efficiency in practice in data centres and is interested in assisting both operators and customers in procurement and sustainability assessment. This document further states that customers or suppliers of IT services may use the CoC practices to support procurements of services that meet environmental or sustainability standards. In this sense, the JRC framework is directly relevant to a study concerned not only with infrastructure performance, but with what cloud customers can know and act upon from provider disclosure.

The JRC is especially useful because it bridges the gap between the literature reviewed above, and the temporal tasks required for addressing the RQ1 set. The literature shows that cloud sustainability metrics are often incomplete, inconsistently reported, and poorly aligned with customer decision needs. The JRC responds to this problem by specifying how environmental and energy-performance reporting should be structured, what should be measured, and how such measurements should be communicated. In methodological terms, the JRC functions in this thesis not as a normative benchmark of best performance, but as a reporting-oriented analytical lens through which the adequacy of CSP disclosure can be evaluated.

Three features of the JRC framework make it particularly suitable for this role:

Transparency and Comparability

The JRC places strong emphasis on standardised measurements and reporting, pulling work from the International Organisation for Standardisation (ISO)/International Electrotechnical Commission (IEC), and from the European Norms (EN). In its section on written reporting, it states that all reports and submissions should reference the reporting category and that the required methods of data collection and calculation should be based on the standardised ISO/IEC 30134 series or equivalent EN 50600-4-x standards where KPIs are reported. This provides the foundation for the Transparency and Comparability dimensions, used to guide the findings of the RQ1 set, further explained in Section 3.4.1. Where providers disclose metrics without explicitly anchoring them to these standardised methods, they become more difficult to compare strictly across providers. Accordingly, this thesis treats the presence of a metric as analytically distinct from the presence of a standardised and methodologically clear metric.

Granularity

The JRC explicitly recognises how reporting quality depends on granularity and reporting categories. Different categories of PUE represent increasing levels of reporting granularity and reporting should reference the category being reported. By emphasising improved visibility and granularity of infrastructure overheads it exemplifies practices such as distribution-board metering, cabinet-level metering, IT device-level metering, and row or cabinet-level temperature monetary. The framework goes beyond static reporting by distinguishing period manual readings, automated daily reading, and automated hourly readings, thereby embedding temporal resolution into the logic of what counts as strong monitoring and reporting practice. This is central to the thesis in terms of where providers disclose environmental information at a level that is useful to customers for making deployment directions. The JRC thus provides a clear basis for treating granularity and temporal resolution as core analytical criteria, operationalised when analysing the findings of the RQ1 set.

Boundaries

The JRC accounts for the fact that many actors do not control the entire data centre, and different areas of responsibility must therefore be distinguished. It identifies a structured set of responsibility areas: physical building, mechanical and electrical plant, data floor, cabinets, IT equipment, operating system and virtualisation, software, and business practice. Hyperscale providers frequently report cross mixed boundaries; some metrics relate to the physical facility, others to IT hardware, some to software or utilisation, and others to broader corporate or supply-chain activities. The JRC framework helps clarify why boundary clarity matters analytically. If reporting objects are not clearly distinguished, comparisons become unstable because providers may appear to be disclosing the same information, while in fact reporting different layers of the stack. The problem is often misaligned data: corporate-level sustainability narratives, data centre operational indicators, and customer-facing service claims and frequently disclosed together, without boundary structure.

Beyond these features, the JRC also strengthens the analytical treatment of specific sustainability metrics, placing the development of an energy monitoring and reporting strategy as central to efficient data centre operation and identifying income energy metering, IT energy metering, room-level and requirement-level temperature monitoring, and structured reporting consoles as foundational practices.

The JRC also acknowledges workload mobility and business relevance in ways that align with the study's concern for cloud customer decision-making, further developed through RQ2. It advises operators to consider the impact of mobile or shifting workloads during the design and operation of data centres, noting that modern application and software increasingly allow workloads to fluctuate and move between sites, thereby changing the ways installed power and cooling infrastructure need to be managed. Finally, it identified the potential value of a business relevant dashboard that reflects and helps manage the energy required to deliver IT services defined by specific business requirements. These elements support the thesis argument that a sustainability disclosure is more useful when it can be connected to customer relevant decision units such as workload, service, or location.

The role in this thesis is twofold. Conceptually, it provides a structural understanding of what robust environmental reporting in data centres should contain. Analytically, it provides the benchmark through which CSP sustainability reports and associated public disclosure can be assessed under RQ1(a) and RQ1(b). The JRC translates the literature reviewed above into a concrete defensible analytical instrument.

2.2.2 Bounded Rationality and Decision-Making Under Uncertainty

To answer RQ2, this thesis adopts a conceptual framework grounded in bounded rationality. In this context, bounded rationality means the deployment decisions are not made under conditions of full information, perfect comparability, or unconstrained choices. Instead, decision-makers operate within incomplete provider disclosure, uncertain environmental outcomes, multiple competing objectives, and limited room for reconfigure cloud infrastructure. They therefore do not optimise in a strict technical sense; they satisfice, selecting options that are good enough within the boundaries of what is known and what is feasible. Three literatures support this interpretation:

- The transparency literature shows that provider disclosures are partial, aggregated, and often not decision-ready (de Vries-Gao, 2026; Hankendi et al., 2025)
- The carbon-aware computing literature shows that lower-carbon deployment options exist, but only for certain workload types and only within operational and SLA constraints (Tripathi et al., 2023)
- The digital sustainability literature shows that efficiency improvements do not automatically reduce total energy demand because rebound and growth dynamics can offset gains (Lange et al., 2020)

On this basis, the analytical tool for RQ2 can be formulated around four dimensions of feasibility:

(1) *Informational feasibility*: is the relevant sustainability information available, granular enough, and comparable enough to support a deployment choice?
(2) *Technical feasibility*: is the workload movable in practice, given latency sensitivity, architecture, dependencies, and operational risk?
(3) *Organisational feasibility*: can a change be made without unacceptable disruption to engineering productivity, internal processes, or team capacity?
(4) *Governance feasibility*: do contractual commitments, provider dependencies, multi-cloud complexity, and limited customer control permit action?

These dimensions turn the literature into an analytical instrument. Rather than asking only whether a region appears greener, RQ2 asks whether Neo4j could usefully act on that information. A sustainability-informed option becomes feasible only when all found dimensions reach a sufficient threshold.

2.2.3 Implications for empirical analysis

The combined conceptual framework developed in Sections 2.2.1 and 2.2.2 have direct implication on how empirical material in this thesis is analysed. These sections establish that cloud sustainability decision-making must be approached through two linked but analytically distinct lenses, the first concerns the quality and structure of provider disclosure, which the JRC Guidelines provide the main benchmark; the second concerns for which bounded rationality provides the interpretive lens. The empirical analysis is therefore designed to move from the external information environment to the intent decision environments.

Chapter 4 should assess provider disclosure not simply as information sources, but as governance artifacts. The relevant question is not only whether CSPs disclose environmental indicators such as PUE, WUE, renewable energy claims, or carbon indicators, but whether these disclosures are sufficiently transparent, comparable, and regionally specific to support customer-specific decision-making. This is why the coding dimension of Transparency, Comparability, and Decision Relevance are analytically appropriate under the JRC Guidelines (Acton et al., 2025).

Chapter 5 should interpret Neo4j's deployment optimisation through feasibility rather than optimisation. The literature does not support the assumption that firms can simply relocate workloads to the greenest region. Workload-shifting studies explicitly show that power, SLA, routing, and workload characteristics constrain movement (Tripathi et al., 2023). Multi-cloud studies further show that portability, vendor lock-in, competence gaps, and architectural complexity shape the real cost of acting across providers or regions (Alonso et al., 2023; Jaensson & Wiklund, 2022). This means the analysis should identify where sustainability considerations can become realistic and relevant for making decisions, and also where they remain secondary or where they are blocked by structural limits. Ultimately, case findings should be read as describing a feasible decision space, not an ideal technical optimum.

The framework implies that empirical findings should be interpreted as trade-offs among competing rationalities. The JRC points towards structure and standardised approach to environmental reporting, while bounded rationality points toward the practical limits of acting on such information. Together they suggest that sustainability related cloud decisions emerge from the interaction between environmental ambition, reporting quality, engineering practicality, contract stability, and organisational capability.

2.3 Synthesis and Remaining Gaps

The literature demonstrates that technical potential for emissions reduction exists at infrastructure and scheduling level, and that carbon-aware deployment modes can theoretically optimise workloads. However, limited research addresses how such possibilities translate into actionable decisions for mid-sized SaaS firms operating under incomplete information and structural constraints. Data centre metrics, carbon-aware workload placement, temporal shifting, and resource optimisation all indicate that cloud computing can, conditionally, reduce emission intensity. At the same time, the literature equally shows that these opportunities are constrained by aggregation in provider reporting, inconsistent methodological disclosure, rebound dynamics in digitalisation, and limited control cloud customer exercise (Bandhavi Sakhamuri, 2025; Beena et al., 2025; de Vries-Gao, 2026; Hankendi et al., 2025; Lange et al., 2020; Safari et al., 2025; Tripathi et al., 2023).

The main research gap, therefore, is not whether greener cloud operation is technically imaginable, but how environmental information can be translated into feasible customer-side decisions under real organisational constraints. Mid-sized SaaS firms are particularly underexplored in this respect. They depend heavily on hyperscale infrastructure, often lack bargaining power over providers, and must balance sustainability aims against productivity, contractual arrangements, and platform dependencies. The developed conceptual framework attempts to bridge this gap and address these constraints, by linking two issues that are often treated separately: provider-level transparency and customer-side feasibility, examining whether the public environmental disclosures of major CSPs provide a sufficiently comparable basis for regional sustainability assessment, and assessing how far such information can inform internal deployment choices at Neo4j.

3 Research Methodology

3.1 Research Paradigm and Methodological Orientation

This thesis adopts a qualitative and interpretive research orientation, with an instrumental case study design. The study investigates suitability-oriented cloud deployment decision-making in context characterised by indirect control over infrastructure, incomplete information, and competing organisational objectives.

An interpretative approach is appropriate because the research questions do not aim to measure emissions quantitatively, but aim to understand:

- How sustainability information is produced and disclosed: ***RQ1(a)*** and ***RQ1(b)***
- How such information becomes actionable under organisational constraints: ***RQ2***

This aligns with the conceptual framework developed in Section 2.2, where bounded rationality suggests that decision-making occurs under conditions of uncertainty and constraint, rather than full optimisation.

3.1.1 Overall Research Design

The research design is structured in two sequential phases linked analytically, reflecting the dual nature of the research problem. The methodologies of each are summarised below in Table 2.

Phase 1 focuses on the external information environment, i.e. the CSP sustainability disclosures. This phase involves a comparative analysis of publicly available documents to determine what environmental information is available, how it is present, and to what extent it is usable for customer decision-making.

Phase 2 focuses on the internal information environment, i.e. the Neo4j case study on decision-making. By examining how such information is interpreted and acted upon in practice, this phase investigates organisational constraint, decision process, and the feasibility of suitability-formed deployment choices.

Alternative approaches such as quantitative emissions modelling or lifecycle assessment were considered. However, limited access to facility-level data and the study's focus on decision feasibility rather than optimisation render such approaches unsuitable. Instead, the current structure allows the study to move systematically from information availability to decision feasibility, ensuring that findings are both analytically grounded and practically relevant.

Table 2. Research Methods and RQs Alignment

Analytical lens	**Research question**	**Data Sources**	**Focus**
Phase 1, External Information Environment			
Comparative: JRC reporting framework	**RQ1(a)** *How transparent and comparable are the sustainability reports and associated public environmental disclosures of Amazon, Microsoft, and Google, regarding data centre energy efficiency and carbon performance?*	Sustainability reports & technical documents (2024-2025)	Transparency, Comparability, & Reporting practices
Comparative: JRC	**RQ1(b)** *What provider-level and selected region-level-information on data centre energy efficiency and carbon*	Sustainability reports &	Decision Relevance &

reporting framework	*performance is available for AWS, Azure, and GCP, and how actionable is this information for customer decision-making?*	technical documents (2024-2025)	regional differences
Phase 2, Internal Information Environment			
Interpretive: bounded rationality	**RQ2** *To what extent are sustainability-informed cloud deployment decisions possible for Neo4j within existing organisational, technical, and contractual constraints?*	Internal Neo4j consultations	Feasibility & organisational constraints

Source: Author's elaboration.

3.2 Case Study Rationale

The case study employs an instrumental case study design, an approach that allows for in-depth examination of organisational decision-making processes and contrast that are often underrepresented in large-scale technical or policy focused studies. As such, Neo4j is examined as a means to explore broader questions concerning sustainability transition and governance, and feasibility of action in cloud-dependent organisations.

Neo4j is selected as the case study organisation due to its position as a mid-sized SaaS firm with high dependence on public cloud infrastructure and limited influence over cloud provider operations. Neo4j operates in a multi-cloud environment, balancing performance requirements, customer preferences, commercial agreements, and investor expectations. Moreover, Neo4j faces increasing sustainability expectations from regulators, customers, and investors, particularly in relation to CSRD (Neo4j, Sustainability Specialist, personal communication, January 28, 2026). These characteristics make Neo4j a critical case for examining the gap between provider-level sustainability performance and customer-level decision capacity.

The selection of Neo4j is further grounded in prior embedded project work conducted within the organisation. Before the thesis was initiated, a preliminary project was undertaken to map Neo4j's data centre usage and establish a foundational GHG accounting framework. This work included a detailed breakdown of internal, canary, and external cloud usage, and aimed to support Neo4j in meeting regulatory requirements and responding to client demands.

This initial project revealed several critical insights that directly motivated the thesis. Cloud data centre usage emerged as a key GHG hotspot within Scope 3 emissions. The analysis of this usage highlighted substantial inconsistencies in CSPs' emissions data, including differences in methodological approaches, reporting boundaries, and data granularity. The increasing demand from clients and investors for emission transparency exposed a growing gap between sustainability expectations and available decision-relevant information within Neo4j. The case is therefore empirically grounded in a real organisational challenge: how to operationalise sustainability decision-making in a context where environmental impacts are mostly externally generated and partially observable (Neo4j, Sustainability Specialist, personal communication, January 28, 2026).

Importantly, this case reflects challenges common across most of the SaaS sector, including reliance on hyperscale providers, long-term contractual lock-in, latency and SLA constraints, growing external pressure to align operations with climate targets, and limited internal resource allocation for sustainability. At Neo4j, sustainability is integrated across teams, with responsibility largely residing within finance, due to Swedish regulatory reporting requirements. This organisational structure further reinforces the suitability of the case for examining bounded rationality and constrained decision-making.

While findings are not statistically generalisable, they support analytic generalisation, where Neo4j represents a critical and representative case. Insights and findings can be conceptually transferable to similar cloud-dependent organisations operating under comparable conditions.

3.3 Data Corpus and Inclusion Criteria

3.3.1 Phase 1, External Information Environment

This phase assesses the annual environmental and sustainability disclosure practices of the three largest hyperscale CSPs chosen as research group: AWS, Azure, and GCP. Focus is taken with respect to data centre energy efficiency and carbon performance. The inclusion criteria reflect the following data corpus:

(1) A report corpus: CSPs' annual sustainability or environmental reports
(2) An expanded public disclosure corpus: associated environmental disclosures, including enterprise digital hubs, technical documents, region-specific public datasets, etc.

The corpus of public environmental disclosures prioritises 2025 publications where available and otherwise uses the latest public disclosure issues within the 2024-2025 period. Amazon's principal source is the 2024 Amazon Sustainability Report, which covers the period from 1 January to 31 December 2024. Microsoft's principal source is the 2025 Environmental Sustainability Report, which explicitly states that reports on the 2024 fiscal year, and that detailed environmental data are also published separately in an Environmental Data Fact Sheet. Google's principal source is the 2025 Environmental Report, which reports 2024 progress and includes a substantial environmental data appendix. The analysis treats the report corpus as the primary object of Phase 1. Accordingly, the analysis does not aim to determine environmental superiority of all cloud regions, but rather to assess what a customer can reasonably infer from publicly available evidence. This distinction is crucial because the three providers disclose different kinds of information at different levels of granularity.

The expanded public disclosure corpus was however necessary because data relevant to the RQs were not consistently contained within the sustainability reports alone. This involved supplementing the reports by exploring hyperlinks and other resources on the enterprises' digital hubs. The integration of these resources supports further triangulation regarding the issue of data availability and further reflects the challenge on retrieving information and seeking out CSP transparency as a cloud customer.

Appendix B provides a comprehensive list of resources analysed, including descriptions and the corresponding hyperlinks to facilitate direct access to the original source.

3.3.2 Phase 2, Internal Information Environment

This phase draws on internal organisation data to examine how sustainability-related information is interpreted and acted upon within Neo4j. The objective is to understand the decision processes and constraints in practice. The data corpus consists of three main components:

(1) Internal analytical work and documentation

The study builds on previous internal work conducted to establish Neo4j's GHG baseline. This concluded the development of a metric to break down emissions per environment (internal, canary, external) and standardise emission across CSPs. The analysis relied on data from AWS, Azure and GCP customer interfaces, combined with consumption data. However, due to inconsistencies in provider reporting, multiple assumptions were required, including estimation of annual emission from partial dataset and interpretation of unspecified units and calculation methods. The process revealed significant methodological uncertainties: lack of clarity on location- and market-based emissions; absence of Scope 3 allocations in AWS; incomplete temporal fiscal year coverage in Azure and GCP; and approximation of GCP emission values based on graph representations. These limitations are used as empirical evidence of transparency challenges analysed in RQ1 and RQ2. Practically for the focus of cloud services, this challenge concerned data from 2024 reporting, as the research commenced in March 2025, at which point GCP and Azure platforms only retained customer-facing emissions data for the preceding twelve months. Obtaining a complete dataset for the full 2025 financial year reporting period therefore required this prior knowledge of how the individual provider interfaces functioned, including the need to request or extract data immediately following the end of the financial year, i.e. extracting at the end of January 2025 to obtain the dataset for January 2024-January 2025. Further investigation identified that historical datasets can be obtained but not be exported directly through the standard user interface. Instead, downloading complete emissions records required the use of specialised engineering environments and technical support, making this extraction process beyond the capabilities of a typical sustainability practitioner. This finding illustrates that data accessibility is constrained not only by disclosure practices but also by the technical design of customer-facing systems. Consequently, the collection of sustainability data depended on cross-functional collaboration with engineering personnel, reinforcing the central observation of this thesis regarding sustainability data's existence with usability constraints.

(2) Access to cloud provider tools and data interfaces

Customer-facing emission tools provided by AWS, Azure, and GCP were examined as part of the empirical analysis. These included the AWS Customer Carbon Footprint Tool (CCFT), Azure Carbon Optimisation Dashboard, and Google Cloud Carbon Footprint Dashboard. These tools represent the primary interface through which cloud customers of this type access emissions data in practice. However, as mentioned, these tools revealed limitations consistent with those reported in the broader literature, including insufficient methodological transparency (AWS, Azure) and ambiguity in allocation and double-counting mechanisms (GCP).

This interaction with provider tools forms an important empirical bridge between Phase 1 on disclosures, and Phase 2 on decision-making feasibility.

(3) Semi-structures consultations

Two semi-structured informal consultations were conducted with key stakeholders: (1) Aura engineering team, for technical perspective; (2) Financial Planning & Analysis (FP&A) team, for financial and governance perspective. Participants were identified through internal referrals, reflecting the distributed and non-centralised nature of sustainability knowledge within the organisation. The process of identifying relevant participants required iterative engagement, reflecting the absence of a dedicated suitability function. The consultations were guided by a semi-structured protocol, allowing for both consistency and flexibility. Topics included decision criteria (cost, latency, performance); constraints (technical, contractual, organisational); perceptions of sustainability relevance; and availability and usability of environmental data. Participants are anonymised by role and consultation date in accordance with ethical requirements. Appendix C contains further details on the consultation guidelines and consent form.

3.4 Data Collection Process

The data collection process was designed to ensure both analytical rigor and alignment with the RQs by capturing the interaction between externally available sustainability information and internally situated decision-making practice. The data was treated as a separated corpus; and collection followed a two-context logic, distinguishing between the information environment available to cloud customers and the organisational environment in which decisions are made.

For the external information environment (Phase 1), data collection focused on publicly accessible environmental disclosures from the three selected CSPs: Amazon, Microsoft, and Google. The process involved systematic identification and review of primarily sustainability reports for the 2024-2025 period, complemented by targeted retrieval of supplementary material such as technical documents, infrastructure sustainability pages, and provider-authored resources. This approach ensures that the datasets reflect the actual informational boundary available to customers, rather than internal or proprietary provider knowledge. The selection process was iterative, with additional material incorporated where suitability-relevant information was not fully contained within primary reports, thereby supporting triangulation and completeness.

For the internal information environments (Phase 2), data collection was shaped by the researcher's embedded engagement within Neo4j, enabling access to organisational knowledge that is typically not formalised or centrally documented. This included interaction with internal systems, explorer to customer-facing emission tools provided by cloud vendors, and engagement with stakeholders involved in deployment or contractual decisions. The process required navigating a fragmented information landscape, where relevant knowledge was distributed across teams rather than institutionally consolidated. This exploratory and iterative process itself reflected the organisation's reality under investigation. To complement this the semi-structured informal consultations were conducted with participants identified based on their involvements in cloud deployment and infrastructure use, and cost governance. The format enables consistent coverage of core themes concerning decision criteria and feasibility of integrating sustainability, while allowing for the emergence of context-specific insights.

The data collection process maintained a strict separation between external and internal sources, ensuring that CSP disclosures are not conflated with internal interpretations. This separation is critical for preserving analytical validity, allowing the study to assess what information is publicly available (RQ1a,b), and to examine how such information is viewed and used in practice (RQ2).

3.5 Analytical Strategy

The conceptual framework developed in Section 2.2 directly informs the empirical design of the thesis. Phase 1 applies the JRC reporting framework to evaluate the extent to which CSP disclosures enable external assessment and comparison. Phase 2 then examines how such information can be interpreted and acted upon within an organisational context operating under bounded rationality, using Neo4j as a case study. Data analysis occurs in two-fold interconnected phases, explained below.

3.5.1 Phase 1, The Comparative Analysis

The data corpus of CSPs' disclosures is assessed in terms of transparency, comparability, and decision relevance of reported energy efficiency and carbon-related metrics across regions. The

CSPs' disclosures are systematically coded according to predefined analytical dimensions: Transparency (T), Comparability (C), and Decision Relevance (D).

These dimensions are operationalised from specific coding criteria further elaborated in Appendix D. Below are broader summaries of the code definitions derived from the JRC reporting framework:

- *Transparency codes (T)* answer whether CSP disclosures are sufficiently explicit, granular, and methodologically clear
- *Comparability codes (C)* answer whether providers report equivalent objects with sufficiently similar definitions, scopes, and boundaries to allow meaningful comparison
- *Decision Relevance (D)* codes answer what disclosure can support customer-facing region selection or workload placement decisions

A conservative evidentiary rule is applied. Claims are only coded positively where they were publicly available and sufficiently explicit in the selected corpus. Absence of public evidence is treated as absence of publicly accessible disclosure, not as evidence that the provider does not internally measure. This evidentiary rule is especially important where providers refer to separate fact sheets, customer dashboards, or supplementary technical materials rather than publishing a consolidated, directly comparable regional dataset. This choice may understate some element of operational sophistication, but it increases the external validity of the analysis for the specific purpose of this thesis: assigning the adequacy of CSP disclosure for customer-facing transparency, comparability, and decision-making.

Coding proceeds deductively, while allowing for constructive emergence of additional themes. Each disclosure is assessed at both provider and region-level where possible. Thematic patterns are then interpreted in relation to the conceptual framework in Section 2.2. The coding system is not arbitrary; it is directly operationalising the RQs. The JRC Guidance supports this orientation by stressing common terminology, reporting category, and procurement usefulness.

3.5.2 Phase 2, The Interpretative Analysis

Insights gathered from Neo4j's current internal deployment practices are assessed to evaluate which sustainability-oriented deployment options are feasible within organisational constraints, such as latency sensitivity, contracts, and engineering workflows. The analysis recognises the organisation may pursue *satisficing* rather than optimising strategies when navigating sustainability trade-offs. Data is analysed along four dimensions identified in the conceptual framework for bounded rationality, Section 2.2.2:

(1) Informational feasibility
(2) Technical feasibility
(3) Organisational feasibility
(4) Governance feasibility

The analysis is conducted through thematic coding and pattern identification, guided by the concept of bounded rationality. The purpose of this focus is to assess whether sustainability-informed deployment decisions are realistically feasible under existing constraints.

3.6 Methodological Limitations

This study is subject to several limitations that are considered when interpreting the findings.

Firstly, the analysis relies on self-reported CSP data, which has been established by the literature in Section 2.1.3 to be often aggregated, methodologically heterogeneous, and selectively presented. This introduces uncertainties regarding the accuracy and completeness, but also the comparability amongst reported metrics. The study does not aim to validate provider performance, but nonetheless these limitations directly affect the robustness of cross-provider comparison and reinforce the indeed to interpret findings as indicative rather than definitive.

Secondly, the study is constrained by limited access to facility-level operational data. The absence of granular, region-specific, and standardised data restricts the ability to perform precise environmental assessments or emissions modelling. This limitation is analytically consistent with the research objective: investigating what information is available to customers. As such, the limitation reflects the empirical reality of the study, and not a methodological shortcoming.

Thirdly, the case study design introduced limitations related to scope and generalisability. The findings are based on a single organisation, Neo4j, and are not statistically generalised. However, the study follows an analytical generalisation logic, where insights are transferable to similar mid-sized SaaS firms operating under comparable conditions of cloud dependency, limited control, increased external pressures, and decentralised coordination on internal sustainability practices.

Fourthly, the research is conducted within a dual academic-industry setting, where the research's embedded position may introduce interpretive bias. To mitigate this reflexive awareness is maintained through the process by explicitly distinguishing empirical and normative observations, critical integrating internal assumptions and narratives, and structuring the analysis through predefined conceptual framework, developed in Section 2.2.

Fifthly, the study draws on many grey literature and industry sources due to the highly technical nature of the research field. These sources may vary in methodological rigor and carry institutional biases; their inclusion is therefore deliberate but critically assessed.

Finally, the study focuses on economic feasibility under real-world constraints, which limits the ability to derive precise environmental impact outcomes, however it strengthens the reliance on findings of organisational practice.

These limitations define the scope: providing a contextualised, decision-oriented analysis of cloud sustainability under conditions of incomplete information and constrained agency.

3.7 Ethical Considerations

This research is conducted in collaboration with Neo4j, the sponsoring company, which provides support through access to internal intelligence, data, expertise, and other relevant resources. While this partnership naturally introduces a company-centric perspective, analytical independence is maintained. This is achieved by critically assessing cloud provider data and clearly separating empirical observations from normative recommendations. Conclusions are grounded in the conceptual framework and RQs.

All internally collected data followed the principles of informed consent, confidentiality, and responsible data handling. Data is securely stored and accessed only for research purposes, and research findings may be used to inform Neo4j's internal strategy and external communication. Some data were provided under confidentiality agreements and have been anonymised or abstracted. Nonetheless, the underlying information and insights derived from this complete, internal data set were fully utilised. These insights informed the assessment, synthesis of findings, and formulation of the actionable recommendations presented, ensuring the thesis remains academically sound and analytically robust.

4 Phase 1 Comparative Results: CSP Disclosure Assessment

4.1 Alignment Between the T-C-D Framework and the RQs

This phase addresses the first RQ set:

RQ1(a) *How transparent and comparable are the sustainability reports and associated public environmental disclosures of Amazon, Microsoft, and Google, regarding data centre energy efficiency and carbon performance?*

RQ1(b) *What provider-level and selected region-level-information on data centre energy efficiency and carbon performance is available for AWS, Azure, and GCP, and how actionable is this information for customer decision-making?*

RQ1(a) is addressed primarily through the T & C dimensions of the T-C-D framework; RQ1(b) is addressed primarily through the D dimension, depending on a degree of T & C established in the first part of analysis.

4.2 Results from the Report Corpus: Transparency, Comparability, and Decision Relevance (T-C-D)

4.2.1 Dimension T: Transparency By Provider

Table 3 summarises the transparency coding results. Scores range from 0 to 2, where 0 indicates absence of insufficiencies, 1 indicates partial or limited disclosure, and 2 indicates comprehensive or sufficient disclosure.

Table 3. Transparency Coding Scores by Provider

Transparency code	***AWS***	***Azure***	***GCP***
T1 - KPI availability	1	1	2
T2 - Standardised KPI disclosure	1	1	1
T3 - Granularity of disclosure	0	0	2
T4 - Methodological clarity	1	1	2
T5 - Temporal resolution	1	1	1
T6 - Boundary and responsibility clarity	1	1	2

Source: *Author's elaboration, based on findings from coding framework (Appendix D).*

The table reads in terms of disclosure-quality: a higher score indicates provision of clearer and more decision-relevant information. It should not be interpreted as evidence of supervisor environmental performance, especially in relation to the JRC guidance, which treats reporting adequacy and operational efficiency as related but analytically distinct issues.

Amazon

AWS discloses a limited but meaningful set of operational indicators in public form. Amazon reports that data within the report reflects calendar year 2024, and that the report includes many business units and subsidiaries, including AWS, Devices, Amazon Fresh, and Whole Foods Markets. Within this broader corporate boundary, Amazon reports a global PUE of 1.15 for AWS data centres and a data-centre WUE of 0.15 L/kWh. It also reports that 100% of the electricity consumed by Amazon was matched with renewable energy sources in 2024 and that 100% of electricity consumed in all data centre regions continued to be matched with renewable energy sources. These disclosures justify partial scores for KPI availability and boundary clarity, but they do not support a high transparency score overall because the core data remain highly aggregated and are not region-specific.

What remains opaque in the AWS case is not whether sustainability activity exists, but whether the public disclosure preserves the link between the activity and customer choice. The report does not publish region-level operational KPI tables, and the reporting object is mixed: provider-specific data are embedded in a broader Amazon sustainability narrative that spans retail, devices, grocery, and logistics. AWS also does not explicitly anchor its reported PUE or WUE to EN 50600-4-x or ISO/IEC 30134 in the main report. Under the JRC benchmark, this weakens both standardisation and comparability.

Microsoft

Microsoft similarly discloses key sustainability information, but primarily at a broad portfolio level. The report is explicit that it is published annually, that it reflects FY2024, and that detailed environmental data are provided separately in an Environmental Data Fact Sheet. Microsoft also frames its operational sustainability footprint across campuses, data centres, devices, software, and value chains, which immediately signals that the reporting boundary is broad rather than infrastructure specific. In methodological terms, the report is comparatively strong: Microsoft describes the use of life-cycle analysis (LCAs), activity-based accounting for Azure hardware, bill of material analysis across the global data centre fleet, and location-sensitive comparisons of carbon and water trade-offs in cooling technologies. This supports the score of 1 for methodological clarity, though it still falls short of full public auditable standardisation.

Microsoft's transparency profile is therefore internally sophisticated but externally uneven. The report itself does not provide the same kind of site- or region-level operational tables that Google published. At the same time, Microsoft describes a strong customer-facing optimisation layer: Carbon Optimisation in Azure provides more granular emissions data, identifies underutilised resources for deletion or rightsizing, and is accessible directly within the Azure Portal. This improves Microsoft's standing as a provider of decision support, but not its standing in public report-based transparency. The distinction is critical for the coding logic of this chapter.

Google

Google is the strongest performer on transparency because its report combines narrative reporting, methodological explanations, and location-level environmental data in a single public source. Google reports that its average annual fleet-wide PUE fell to 1.09 in 2024, that is maintaining a 100% renewable energy match on a global basis, and that its global average 24/7 carbon-free energy reached 66% in 2024. The report also makes clear that the Carbon Free Energy (CFE) metric is calculated at hourly granularity and Google's carbon intelligent computing platform shifts tasks across locations and times of day based on local grid carbon intensity. In addition, Google discloses regional CFE data in its environmental tables, including strong regional variation, with Asia Pacific materially lagging other regions.

Google's strength is not only that it discloses more, but that is disclosed at a level that is analytically usable. It provides public infrastructure-level signals that can be linked to specific geographies and operational choices. At the same time, a conservative interpretation remains necessary. Although the report gives a robust methodological account of hourly CFE and provides PUE values and water-risk explanation, it does not explicitly state in the main disclosed that PUE, WUE, or related indicators are calculated under EN 50600-4-x or ISO/IEC 30134 formulas. This justifies a score of 1 rather than 2 for standardised KPI disclosure.

A further refinement concerns temporal resolution. Google is stronger than AWS and Microsoft because the CFE metric is explicitly hourly, and the JRC recognised automated hourly readings as higher-order monitoring practice (Acton, et al., 2025). However, Google's disclosure is not uniformly dynamic across all environmental KPIs; PUE and water are still primarily annualised in the public report. For this reason, the more defensible score for T5 is 1 rather than 2 when assessing the total disclosure system rather than a single metrics family.

4.2.2 Dimension C: Comparability Cross-Provider

Table 4 identifies the principal comparability blockers in the corpus. The evidence suggests that the main comparability problem is structural rather than incidental. The providers are disclosing different quantities of information, alongside disclosures of different objects, at different scales, using different metric constructs, and without consistently trying those constructs to the same formal standards.

Table 4. Comparability Blockers Across Providers

Blocker	***Amazon***	***Microsoft***	***Google***	***Consequence***
PUE standard explicitly cited	No	No	No	Prevents strict method equivalence under the JRC requirements that reported KPIs use ISO/EN formulas.
Region-level energy data	No	No	Yes	Prevents region-to-region comparison across all three providers
Boundary clarity	Partial	Partial	Stronger	Broad, mixed reporting weakens comparability of footprint claims
Water context	Limited	Limited	Better but incomplete	Cross-provider water comparison remains weak without water-source energy intensity context
Customer decision linkage	Partial	Partial / strong via tools	Strong	Only Google links public infrastructure data directly to location choice in a public and auditable manner

Source: Author's elaboration, based on findings from coding framework (Appendix D).

The first and most basic comparability failure concerns ***standardisation***. The JRC guidance states that if KPIs are reported, the standardised ISO/IEC 30134 series or EN 50600-4-x equivalences should be used in written reporting. It explicitly lists PUE, REF, CUE, and WUE among the relevant standardised KPIs, and further specifies that PUE should be reported according to EN 50600-4-2 or ISO/IEC 30134-2, REF according to EN 50600-4-3 or ISO/IEC 30134-3, and WUE according to ISO/IEC 30134-9. None of the three providers explicitly anchors its public KPI claims to these standard formulas in the report text itself. This means that even where all three use familiar metric labels, strict method equivalence can be assumed from the public corpus alone.

The second comparability blocker is ***granularity***. The JRC notes that different categories of PUE represent increasing levels of reporting granularity and encourages grouped-by-location reporting as a surefire practice. This is precisely where comparability breaks down in present material. AWS and Microsoft largely report global or fleet-level values, whereas Google publishes public regional and location-level data. This asymmetry means that a comparison between an AWS region and a Google region is structurally unbalanced before any substantive environmental analysis even begins. One side of the comparison is public and region-specific; the other is global and aggregated.

A third blocker is ***boundary mismatch***. Amazon states that the report includes many business units and subsidiaries, including AWS, Devices, Amazon Fresh, Whole Foods Market, Amazon Private Brands, Stores and Twitch. Microsoft frames its sustainability footprint across campuses, data centres, devices, software, and value chain. By contrast, Google more clearly separates location-specific data centre indicators from broader environmental reporting, even though some of its environmental metrics still cover both data centres and offices. The effect is that provider disclosure does not mean the same thing across all three cases. Corporate sustainability narratives, global fleet metrics, and location-specific infrastructure data are being juxtaposed under one umbrella, creating a mismatch in levels of analysis that constrains provider comparison.

A fourth blocker concerns ***renewable electricity claims***. The providers do not report the same construct. Amazon emphasise 100% matching of electricity consumed by its global operations with renewable energy and 100% matching in all data centre regions. Microsoft frames its ambition as matching 100% of electricity consumption with zero-carbon electricity purchases 100% of the time and reports large-scale contracted carbon-free electricity growth. Google reports both annual renewable matching and hourly/local 24/7 CFE. The JRC benchmark, however, points to REF as the standardised renewable-energy factor when renewable share is reported. These are all relevant disclosures, but they are not interchangeable metrics and therefore should not be treated as directly comparable without qualification.

A fifth blocker concerns ***water comparability***. The JRC explicitly warns that water consumption cannot be directly compared with energy efficiency unless the energy intensity of the water source is understood, and that comparing water consumption between buildings is therefore not useful without such context (p.50). None of the three providers supplies a fully standardised contextual disclosure that would make cross-provider WUE or volumetric water comparisons robust. Google offers more context than the others through its climate-conscious cooling narrative and local water-risk framework, but even this does not fully resolve the problem. Cross-provider water comparison should therefore be treated as highly provisional.

4.2.3 Dimension D: Decision Relevance for Customers

The central practical question for this chapter is whether a mid-sized SaaS firm can choose one cloud region over another on defensible environmental grounds using only publicly available CSP disclosure. Table 4.3 summarises the answers.

Table 5. Decision Relevance for Region-Selection Decisions

Decision Relevance question	***AWS***	***Azure***	***GCP***
Can a customer compare regions publicly?	No	No	Yes partially
Is the information linked to customer decision units such as region or workload placement?	Weakly	Indirectly; stronger portal tools	Yes
Are business-facing dashboards or optimisation tools disclosed?	Limited in public report	Yes	Yes
Is a regional selection publicly defensible?	No	No	Yes, with caveats

Source: Author's elaboration, based on findings from coding framework (Appendix D).

For GCP, the answer is partially yes. Google's report publicly links environmental information to geographic deployment choices through regional CFE values and location-level data centre indicators, while also explaining that its carbon-intelligent computing platform shifts computing tasks across locations and times of day based on local grid carbon intensity. This means that a customer could reasonably justify, for example, preferring a region associated with much higher CFE to one with much lower CFE when carbon-aware siting is the decision criterion. The disclosure is not sufficient for a complete sustainability optimisation across all trade-offs, but it is sufficient for a defensible carbon-aware region selection argument.

For AWS and Azure, the answer is no if the decision must be based on public report evidence alone. AWS reports global PUE, WUE, and renewable-energy matching, but not public region-specific operational KPI tables. Microsoft disclose strong internal measurement practices and customer-facing optimisation capability, but its public report likewise does not publish region-level operational tables, in consequence, a customer cannot determine from the report corpus whether one AWS region is environmentally preferable to another AWS region, whether one Azure region is preferable to another Azure region, or whether a specific AWS region is preferable to a specific Azure region. The necessary region-specific infrastructure evidence is absent.

The distinction between public disclosure and post-deployment tooling is specifically important here. Microsoft states that Carbon Optimisation in Azure provides more granular emission data and actionable recommendations within the Azure Portal. This is clearly decision-supportive for customers already using Azure. Yet it does not fully solve the public procurement problem addressed in this thesis, because the tool's logic and outputs are not presented in a public, standardised, and independently auditable reporting format. In other worlds, Microsoft is stronger on business facing optimisation than on public ex ante comparability. The same principle applies, though less explicitly in the present corpus, to AWS's broader sustainability tooling ecosystem.

Three analytical levels can therefore be distinguished, first, what can be decided: within Google's ecosystem, customers can make a reasonably defensible region-based carbon-aware decisions by using public provider disclosures. Second, what cannot be decided: region selection within AWS or Azure, or between AWS and Azure, cannot be defended based on public report evidence. Third, what requires inference beyond public disclosure: any attempt to infer regional environmental performance for AWS or Azure would require proprietary portal outputs, provider-specific local energy announcements, or third-party grid datasets. That shifts the analytical burden from the provider onto the customer, which is precisely what actionable disclosure should avoid.

4.3 Integrated Provider Findings from the Expanded Public Disclosure Corpus

AWS; strong fleet-level efficient disclosure, weak public regional comparability

AWS reports several clear public indicators of data centre efficiency, but these remain predominantly global or fleet-level rather than region-specific. AWS Sustainability (n.d.) states on its Data Centers page that it reported a global PUE of 1.15 in 2023, and that this is better than both the public cloud industry average and estimated on-premises enterprise averages. The page also states that customers running compute-heavy workloads can reduce emission by moving from on-premises to AWS, with further reductions possible through AWS-designed silicon, and that Graviton-based EC2 instances can use substantially less energy for the same performance. These claims support a provider-level narrative of efficiency and lower carbon intensity relative to on-premises infrastructure (AWS Sustainability, n.d.).

At the same time, AWS does not provide a public region-by-region environmental dataset comparable to Google's regional CFE disclosure. The publicly available AWS materials identify innovation in cooling, server placement, silicon efficiency, circularity, and water positive provider, but the key metrics remain aggregated at global fleet level. This limits both cross-provider comparability at regional levels and customer actionability when choosing one AWS region over another on environmental grounds. A customer may reasonably infer that AWS infrastructure is designed for efficiency overall, but not which AWS regions are relatively more carbon- or energy-advantageous from publicly available evidence alone (AWS Sustainability, n.d.).

Methodologically, AWS therefore performs better on transparency at the provider level than at the regional level. It discloses meaningful indicators, but the chapter should distinguish between provider-level sustainability signalling and customer-usable regional disclosure. The former is present; the latter is largely absent from the public corpus reviewed here.

Azure; methodologically rich but fragmented disclosure

Microsoft presents a different disclosure profile. Its 2025 Environmental Sustainability Report is explicit that environmental progress is reported annually and that environment data are also published and reported in an Environmental Data Fact Sheet. This is important because it signals a relatively sophisticated reporting architecture but also means that the core annual report is not the sole location of relevant evidence. Microsoft further describes its sustainability work across campuses, data centres, devices, software, and value chain, which reflects broad internal accounting sophistication but also widens the reporting boundary beyond a narrowly defined cloud-infrastructure dataset (Taylor, 2022).

Azure's strongest contribution to methodological transparency appears in supplementary technical communication. In its Azure sustainability guidance and data centre metrics blog, Microsoft directly acknowledges that customers consider energy and water consumption in commercial decisions, introduced public PUE and WUE metrics, and explains that it tracks these globally by operating geographies. The same blog defines PUE as an industry metric and notes that local environment and infrastructure affect how it is calculated, with slight variations across providers. This is an important admission for the comparability analysis, because it publicly recognised that PUE is not automatically comparable without attention to methodological differences (Taylor, 2022; Walsh-Elwell, 2022).

Microsoft also discloses selected region- or market-specific sustainability information, but not in a unified cross-region public table equivalent to Google's region carbon page. For example, Microsoft announced 24/7 hourly matching and renewable-energy monitoring for its new data centre regions in Sweden, explicitly framing this as a first among hyperscale cloud providers for commercial products. This offers evidence that selected local and regional sustainability characteristics can be probably disclosed, yet these examples remain case-based rather than systematically comparable across Azure regions (Walsh-Elwell, 2020).

In embodied carbon accounting, Microsoft is ahead in methodological sophistication. Its Azure hardware disclosure explains that high-granularity, process-based data now cover over 97% of cloud server rack emissions and nearly 80% of semiconductor emissions, improving Scope 3 accounting and making disclosure more representative of real hardware configurations in Azure data centres. This is valuable for transparency, but it still does not translate directly into a simple region-by-region customer comparison tool (Leoaspauza, 2026).

As a result, Azure should be interpreted as comparatively strong on methodology explanation and internal accounting sophistication, but weaker than Google on systematic public regional actionability. This disclosure is not absent; rather it is fragmented access annual reporting, infrastructure pages, and technical blogs.

GCP; the strongest public basis for carbon-aware regional decision-making

Among the three providers, Google offers the strongest public basis for linking sustainability disclosure to customer relevant region selection.

Its public Google Cloud region-carbon page explicitly framed cloud-region choice as partially a carbon decision, stating that electricity used to power an application may have different carbon intensity depending on region and that customers can use the published CFE% indicated when locating workloads. Google Cloud (n.d.) determined that for a location to be considered low-carbon, it must belong to a region with a Google CFE% of at least 75%, or if CFE% information is unavailable, a grid carbon intensity of maximum 200gCO_2-eq/kWh.

Google Cloud (n.d.) defines this indicator as the average percentage of carbon-free energy consumed in a particular location on an hourly basis and states that for customers it represents the average percentage of time an application will run on carbon-free energy. This is unusually direction-oriented disclosure because it converts provider sustainability reporting into directly interpretable region-level variables.

Google's supplementary guidance also makes this information actionable. In the Google Cloud Blog (2021) announcing regional CFE publication, Google states that customers can use locations with higher CFE percentage, can pick a lower-carbon region for new applications, can run batch jobs in lower-carbon regions, and can restrict organisational policy to preferred lower-carbon regions. This is a rare example of a provider not only disclosing a region-related indicator but also explicitly telling customers how to use it in infrastructure design and workload placement (Google Cloud Blog 2021; Google Cloud n.d.).

On energy efficiency, Google Data Centers (n.d.) also discloses stronger technical standardisation than the other providers. It is public data centres efficiency page providers ongoing PUE reporting at fleet and campus level, and its 2024 annual reporting corresponds to a fleet-wide trailing-twelve-month PUE of 1.09. This gives Google a stronger position on metric continuity and public visibility than providers that disclose only a global annual number or only selected regional examples.

Ultimately, Google scores highly on transparency because it publishes named indicators, methodological explanation, and region-level carbon information; highly on comparability within the provider because regions are present through the same metric logic; and highly on decision-relevance because explicitly links disclosure to customer region choice. However, even Google's disclosure is stronger for carbon performance than for regional energy-efficient comparison, because public PUE remained less directly available to customers at the cloud region decision interface and the regional carbon indicators (Google Cloud, n.d.; Google Cloud Blog, 2021; Google Data Centers, n.d.).

4.4 Synthesis: What Phase 1 Implies for Actionable Region Comparisons

The coding is intentionally conservative. Where a provider referred to separate data sheets, customer portals, or non-public optimisation environments, those capabilities were not treated as equivalent to publicly auditable report disclosure. This may understate operational sophistication, particularly for Microsoft and AWS, but it strengthens the validity of the analysis for external customer-facing transparency assessments. Several conclusions follow from the coding analysis:

(1) comparisons that are currently defensible across all three CSPs are mostly macro-level comparisons. The providers can be compared at the level of broad decarbonisation orientation, high-level renewable energy strategy, or global fleet efficiency signals. AWS reports 100% renewable energy matching across global operations and public AWS PUE and WUE values. Microsoft reports broad operational decarbonisation progress, larger-scale contracted carbon-free electricity growth, and extensive internal measurement and supplier engagement programs. Google reports similar macro-level ambition while also going further in public location-level disclosure. These macro-level comparisons are real and informative, but they remain too aggregated to resolve concrete customer deployment choices in most cases.

(2) region-to-region comparisons are not equally possible across providers. Google is the only provider in the corpus that comes close to what may be termed decision-composable disclosure: public information that can be combined into coherent agreement about the environmental implications of choosing one deployment geographies over another. AWS and Microsoft do not disclose enough public regional infrastructure data for the same purpose. This is not because they disclose nothing, but because the disclosure remains too aggregated and too detached from the level at which customers actually choose cloud locations.

(3) cross-provider water comparison remains methodologically weak. The JRC states that water consumption cannot be directly compared with energy efficiency unless the energy intensity of the water source is known, and that comparison between buildings is not useful without such context. Because none of the three CSPs fully discloses this contextual layer in standardised form, direct comparison of volumetric water consumption or WUE across providers remains scientifically fragile. Google improves interpretability through a local water risk and cooling strategy narrative, but the underlying cross-provider comparability limitation remains.

(4) the most important divide is between corporate sustainability disclosure and actionable infrastructure disclosure. The former signals ambition and progress; the latter enables cloud customers to make location-sensitive decisions, in this these, actionable disclosure refers to publicly available, standardised, methodologically clear, and region-relevant information sufficient to support a defensible deployment choice without requiring proprietary internal data or post hoc reconstruction, on that definition, only Google's public disclosure currently approaches the threshold for actionable region-level decision relevance. AWS and Microsoft both show important strengths: AWS in global KPI and renewable matching disclosure, Microsoft in methodological sophistication and customer-facing tooling. However, neither published enough public region-specific infrastructure information to meet the same threshold.

These findings imply that further progress in cloud sustainability disclosure depends on less generic increases in reporting volume, more on harmonisation of disclosure objects and methods. The JRC benchmark points towards what harmonised reporting would entail: using the standardised KPI formulas, explicitly reporting categories, grouped-by-location reporting where relevant, and contextualisation of water and energy metrics, without this kind of convergence, CSP disclosure will remain rich on claims but uneven in customer usefulness.

4.5 Answering the Research Questions: RQ1(a) and RQ1(b)

RQ1(a)

Public environmental reporting by the three CSPs is only partially transparent and only unevenly selectively comparable at regional level with respect to data centre energy efficiency and carbon performance.

All three providers disclose meaningful environmental information, but they do so at different levels of granularity and through different combinations of reports, webpages, and technical materials. Google provides the strongest public disclosure architecture for regional carbon-related comparison by combining methodological explanation of 24/7 CFE with region- and grid-level disclosure. Amazon and Microsoft both disclose important environmental progress and performance signals, but these remain predominantly global, fleet-wide, or otherwise insufficiently tied to public regional decisions. AWS providers clear credible provider-level efficiency disclosure, but little publicly usable region-level environmental comparisons. Microsoft provides the richest methodological explanation of measurement and accounting, but its public disclosure is more fragmented and less systematically region-comparable than Google's.

RQ1(b)

Publicly available information is stronger at the provider level, and only selectively available at region-level. Within Google, some regions appear relatively more carbon advantageous than others because the report differentiates performance by grid region and explicitly links carbon-aware computing to local grid conditions.

Google provides the most suitable public information because it publishes average hourly CFE% in Google Cloud regions and explicitly encourages customers to use these differences when choosing regions and running batch workloads. AWS publicly discloses a global PUE of 1.15 and several infrastructure-efficient innovations support provider-level assessment but not robust regional differentiation. Microsoft publicly discloses measure logic, customer guidance, and selected regional initiatives such as Sweden's hourly renewable-energy matching but not standardised public regional comparison datasets across Azure. Accordingly, provider-level actionability is present for all three providers, but region-level customer actionability is strongest for Google, partial for Azure, and weak for AWS in the reviewed corpus.

5 Phase 2 Case Results: Neo4j Feasibility and Constraints

This phase addresses the second RQ:

RQ2 *To what extent are sustainability-informed cloud deployment decisions possible for Neo4j within existing organisational, technical, and contractual constraints?*

5.1 Case Baseline: Where Deployment Decisions Happen

Neo4j's cloud deployment decision-making is distributed across multiple functions rather than concentrated in a single unit. Provider-level commercial relationships and spend commitments are negotiated at executive and cross-functional levels, while decisions regarding deployment configurations, internal development environmental, and day-to-day operational use are shaped more directly by engineering, platform, product, and IT teams (Neo4j Consultation #2 - FP&A, 15 Apr. 2026; Neo4j Consultation #1 - Aura, 26 Mar. 2026).

This division produced a layered decision environment. At the highest level, top management, the agreed upon financial commitments generate contractual lock-in, which limit flexibility. Once these commitments are in place, more localised deployment decisions are shaped by customer location, latency requirements, data security and legal constraints, service availability, and pricing (Neo4j Consultation #2 - FP&A, 15 Apr. 2026). In this sense, provider selection and regional deployment are linked, but not determined by the same decision logic.

The consultation also suggests that Neo4j's internal cloud relationship is historically path dependent, reflecting both legacy choices and the practical difficulties of migrating working systems without disrupting ongoing operations (Neo4j Consultation #2 - FP&A, 15 Apr. 2026).

These findings show how cloud deployment decisions at Neo4j are shaped by a combination of commercial commitments, customer-facing requirements, and organisational structures. This is the context in which sustainability-oriented deployment decisions must be understood.

5.2 Results from Internal Context and Organisational Structure

The consultation findings indicate that Neo4j's internal cloud usage is differentiated across environments and functions. Broadly, two organisational areas can be distinguished: first, the teams building and operating the SaaS offering, particularly Aura; and second, the IT-managed environments that support internal systems, tooling, demos, and proof-of-concept work (Neo4j Consultation #1 - Aura, 26 Mar. 2026). This distinction matters because the incentives governing cloud optimisation differs across these areas.

On the Aura side, cloud costs are more visible and more tightly managed because the product is sold through usage-based SaaS economics. In this context, resource optimisation already exists, but it is primarily driven by cost rather than by environmental performance. By contrast, internal IT environment's function more as cost centres and appear to be less intensively optimised (Neo4j Consultation #1 - Aura, 26 Mar. 2026). As a result, opportunities to integrate sustainability considerations may differ across the organisation depending on where cost discipline is already embedded.

The consultations further show that Neo4j's position as a cloud customer shapes its assumption about sustainability. Both respondents suggested that major hyperscale cloud providers are generally premised to meet baseline sustainability threshold, largely because energy efficiency is already in the provider's financial interest. From this perspective, cloud providers are assumed to optimise infrastructure design, cooling systems, and electricity usage more effectively than most firms could achieve in their own facilities, due to economies of scale and greater engineering capacity (Neo4j Consultation #1 - Aura, 26 Mar. 2026; Neo4j Consultation #2 - FP&A, 15 Apr. 2026). This assumption reduces the extent to which sustainability becomes an explicit variable in internal deployment decisions.

At the same time, the consultations reveal that this assumption remains largely untested in operational practice. Sustainability is acknowledged as relevant, particularly in relation to reporting obligations, customer expectations, and broader strategic positioning, but it is not yet embedded as a routine decision criterion at the level of deployment planning (Neo4j Consultation #2 - FP&A, 15 Apr. 2026).

5.3 Consultation Results: Thematic Analysis

The consultation findings are organised into four themes reflecting decision-making logic, constraints, flexibility, and information gaps; these themes provide insights into how sustainability considerations interact with organisational realities in cloud deployment. To ensure analytical alignment with the conceptual frameworks (Section 2.2.2), these themes are explicitly interpreted through the four dimensions of feasibility: informational technical, organizational, and governance. Table 6 summarises these themes and main characteristics that define each category, and the results are presented per theme below.

This approach ensures that empirical findings are not only descriptive, but analytically structured to directly address RQ2.

Table 6. Thematic Structure of Consultation Findings and Feasibility Dimensions

Theme	***Characteristics***	***Feasibility Dimensions***
A: Decision-Making Hierarchy & Priorities	Cost, latency, and performance dominate decision-making; sustainability remains secondary and largely implicitly	Organisational (decision criteria); Informational (lack of visibility)
B: Structural & Operational Constraints	Latency, compliance, contractual commitments, and workforce capacity constrain deployment choices	Technical (workload constraints); Governance (contracts, lock-ins); Organisational
C: Feasible Levers for Sustainability	Flexibility exists in internal workloads, new deployments, and architecture optimisation	Technical (workload mobility); Organisational (implementation capacity)
D: Information Gaps & Decision Relevance	Data is unavailable, unclear, not actionable; cost acts as proxy	Informational (primary contractions); Governance (provider dependency)

Source: Author's elaboration based on conceptual framework (Section 2.2.2).

Theme A: Decision-Making Hierarchy & Priorities

Across both consultations, cloud deployment decisions were described as being primarily driven by cost, service quality, and latency, with sustainability playing at most an assumed and largely implicit role. This reflects a condition of bounded organisational feasibility, where only measurable and operationally integrated criteria influence decision-making. While sustainability criteria in cloud deployment are acknowledged as a strategic add-on, it is not embedded by choice; cost and performance remain the dominant criteria, largely due to their visibility and measurability. As one respondent noted "cost is much easier to track than emissions, which makes sustainability harder to act on" (Neo4j Consultation #1 - Aura, 26 Mar. 2026). In the Aura context, cost is especially visible because teams are closely tied to usage-based SaaS economics and therefore have stronger incentives to optimise cloud resources.

The FP&A consultation reinforces this picture from a financial governance perspective. Regional and provider decisions were described as fundamentally demand-driven: the organisation aims to align hosting with customer demand, project margins, and maintain service availability. In this logic, sustainability is not absent, but is often treated as an underlying assumption, namely, that hyperscale providers are already operating efficient facilities because doing so is economically beneficial from them (Neo4j Consultation #2 - FP&A, 15 Apr. 2026).

From a feasibility perspective, this highlights a combined organisational and informational constraint. Sustainability is absent due to the lack of integration into decision systems, without measurable, comparable, and operationally embedded indicators, sustainability cannot compete with cost and performance as a decision variable, especially in an organisational context like Neo4j without centralised coordination of sustainability practices. Despite these hurdles, there remains the potential for Neo4j to refine its data centre landscape and adopt a new narrative driven by a sustainability-informed perspective with the fragmented information available.

These findings suggest that sustainability is currently not embedded as an explicit and measurable decision factor within the operational cloud deployment process. Instead, it remains largely subordinate to cost, customer demand, and service performance.

Theme B: Structural & Operational Constraints

The findings reveal that sustainability decisions are bound by technical and governance feasibility constraints; both consultations identified multiple limits to the feasibility of sustainability-oriented deployment changes.

Latency and service quality requirements emerge as limits to technical feasibility, considered non-negotiable for customer-facing workloads, where relocation would compromise performance or customer preferences. Hosting data closure to customers remains essential in many cases to avoid degraded services levels or potential conflicts with service expectations (Neo4j Consultation #2 - FP&A, 15 Apr. 2026). Similarly, data security requirements and legal frameworks can constrain where customer data may be hosted.

Contractual and financial arrangements appear as governance constraints, for these create lock-in effects at a provider level. Spend commitments such as savings plans are typically negotiated across the organisation as a whole and are not region-specific. While these commitments may not prevent all regional flexibility, they can make cross-provider movement difficult and reduce incentives to reconfigure environments purely for environmental reasons (Neo4j Consultation #1 - Aura, 26 Mar. 2026; Neo4j Consultation #2 - FP&A, 15 Apr. 2026).

Engineering resources limit organisational feasibility; even where more sustainable alternatives may be conceivable, implementing regional migration or architectural redesign requires time and team capacity, which can often compete with core business priorities. One consultation noted that if a more sustainable option were clearly available and migration were technically easy, it would not necessarily be controversial; however, in practice, organisations of this type face resource constraints (Neo4j Consultation #2 - FP&A, 15 Apr. 2026).

These constraints indicate that sustainability cannot be treated as an independent optimisation variable. Rather, it must be incorporated, if at all, within existing technical, contractual, and organisational boundaries.

Theme C: Feasible Levers for Sustainability

The conclusions identified several areas where sustainability-informed action may still present conditional feasibility, under these aforementioned constraints.

The most flexible are internal development and testing environments, including canary developments. These internal environments exhibit higher technical feasibility, as they are less constrained by customer latency requirements and SLAs, and may therefore offer greater scope for experimentation with more sustainable regions or configurations (Neo4j Consultation #1 - Aura, 26 Mar. 2026). New deployments offer higher organisational feasibility, seen as more promising than migration of existing environments, since sustainability consideration could be integrated at the point of rollout without requiring cost transitions (Neo4j Consultation #1 - Aura, 26 Mar. 2026; Neo4j Consultation #2 - FP&A, 15 Apr. 2026).

Respondents identified architectural optimisation choices as another lever. For example, designing system to smoothen demand rather, than scaling aggressively during short-lived peaks, reduce computing intensity where technically feasible (Neo4j Consultation #1 - Aura, 26 Mar. 2026). Similarly, machine-level optimisation and workload design choices may influence recourse use independently of provider or region selection. This was raised in relations to CPUs versus Graphics Processing Units (GPUs): GPUs were described as substantially more energy-intensive than CPUs. Where workloads can be run efficiently on CPUs, they are often preferable from both cost and resource-use perspectives (Neo4j Consultation #2 - FP&A, 15 Apr. 2026). This reflects low-friction feasibility where sustainability aligns with cost and efficiency.

However, these levers are currently considered opportunistic rather than strategic, driven mainly by cost and efficiency logic, and not by explicit sustainability objectives. Environmental gains may occur indirectly through financial optimisation but are rarely the primary rations for action.

Theme D: Information Gaps & Decision Relevance

A central finding across both consultations is the perceived lack of actionable sustainability information available to decision-makers, directly correlating informational feasibility as the dominant constraint.

Cost data is highly visible and integrated into operational thinking. By contrast, emissions and sustainability metrics were described as much harder to interpret and much less available in a form that engineers or decision-makers can readily use (Neo4j Consultation #1 - Aura, 26 Mar. 2026). Reports from CSPs may offer broad sustainability information, but they do not provide the same granularity or operational relevance as cost data.

The first consultation suggested that cost may function as partial proxy for sustainability, since more efficient regions may also tend to be more cost-efficient. However, this relationship was described as imperfect and insufficient for confident decision-making. Expensive regions may reflect infrastructure underdevelopment and higher emission but cost alone does not provide a reliable environmental signal (Neo4j Consultation #1 - Aura, 26 Mar. 2026).

The second consultation similarly indicated that sustainability metrics such as PUE and WUE are recognised as relevant in principle, but are not currently considered in internal deployment decision-making, environmental thresholds are assumed to be met by the major hyperscalers, and deeper evaluation of provider- or region-level sustainability is not currently part of routine procurement or deployment logic (Neo4j Consultation #2 - FP&A, 15 Apr. 2026).

Both consultations converge on the point that sustainability would only become genuinely actionable if:

(1) Clear region-level environmental data were available
(2) Such data were comparable across providers
(3) Metrics could be integrated into existing operational or financial decision processes

Without these conditions, sustainability remains difficult to operationalise beyond broad strategic commitments within the context of Neo4j.

5.4 Synthesis: Feasibility for Sustainability-Informed Deployment

The empirical findings from Phase 2, interpreted through the bounded rationality framework developed in Section 2.2.2, indicate that sustainability-informed cloud deployment at Neo4j operates within a bounded and conditional decision space. Rather than represent a binary condition of possibility, feasibility emerges as the outcome of the interaction between four dimensions: informational, technical, organisational, and governance feasibility. These dimensions operate collectively, as they determine whether sustainability considerations can be meaningfully integrated into deployment decisions.

Sustainability does not function as a primarily optimisation objective; decision-making follows a *satisficing* logic: sustainability considerations are incorporated only when they align with existing operational priorities and do not conflict with core constraint. The analysis therefore distinguishes between actions that are currently feasible, those that require improved enabling conditions, and those that remain structurally constraints.

5.4.1 Feasible Actions Under Current Conditions

Under present organisational and infrastructural conditions, sustainability-informed actions are feasible primarily in contexts where all four feasibility dimensions reach a sufficient threshold. This is most evident in internal and non-production environments, where constraints are comparatively weaker.

Informational feasibility provides perspective on how Neo4j retains partial visibility through internal tooling interfaces provided by the CSPs, even in conditions where CSPs' disclosures remain limited in granularity and comparability (as established in Phase 1). These interfaces offer emissions estimates and usage breakdown at account or service level. Although these data are methodologically inconsistent, or at least methodologically inconsistent in their ways of reporting, and subject to identified uncertainties, they do provide sufficient directional insights to support internal discussions and exploratory analysis. This suggests that informational feasibility, while incomplete, is not entirely absent.

Technical feasibility is perspective closely ties to workload characteristics. Internal workloads, such as development, testing, and canary environments, are not subject to strict latency or service-level constraints and can therefore be considered relatively movable. This contrasts sharply with customer-facing production workloads, where relocation would introduce unacceptable performance risks.

Organisational feasibility further conditions what actions are realistically pursued. The findings indicate that sustainability considerations are most likely to be adopted when they align with existing operational logics, particularly cost efficiency. For example, practices such as avoiding overprovisioning, smoothing in demand peaks, or favouring less resource-intensive compute configuration (e.g. CPU over GPU where feasible) are already embedded within operational workflows. These actions are not driven explicitly by sustainability objectives, yet they produce indirect environmental benefits.

Governance feasibility is likewise perceived as a determining factor. Actions that do not challenge contractual commitments, provider relationship, or customer expectations, remain permissible within existing governance structure. Consequently, sustainability initiatives that are incremental and low risk are more likely to be implemented.

These conditions suggest that sustainability-informed deployment is currently feasible in limited but meaningful subset of cases: prioritisation of internal and flexible workloads for experimentations; integration of sustainability considerations into new deployments and not retrofitting existing systems; and leverage on existing efficiency optimisation practices. Such actions remain incremental and do not fundamentally alter organisational deployment strategies.

Neo4j's Practical Assessment Approach

An important empirical insight emerging from this study concerns the role of internal data access shaping the feasibility of sustainability-informed decisions-making. Through specialised customer interface tooling, Neo4j maintains precise visibility into its cloud infrastructure usage, including detailed percentage breakdowns of data centre regions utilisation across its various providers. Although the specific underlying of these insights cannot be disclosed due to a non-disclosure agreement, their existence constitutes a critical empirical anchor for this research.

This internal access enabled a direct, albeit confidential, alignment between the external sustainability signals identified in Phase 1, such as GCP's regional level CFE disclosures, and Neo4j's actual operational footprint. In practical terms, this allowed for the construction of an internally grounded, comparative view of deployment patterns across regions and providers, combining usage data with publicly available sustainability indicators. Although methodologically inconsistencies and data gaps prevent fully standardised comparisons, this process demonstrates that a composite, evidence-based understanding of infrastructure sustainability is achievable in practice, even under conditions of incomplete information.

This finding is analytically significant because it challenges the implicit assumption that cloud customers are entirely constrained by the lack of data. While public disclosure remains insufficient, as demonstrated in Phase 1, internal tooling providers access to partially disaggregated environmental information that can support exploratory analysis. It also illustrates that sustainability-informed reasoning does not require perfect data; rather it can emerge through the structured combination of heterogeneous data sources, including provider disclosure, grid-level proxies, and internal suage metrics.

Building on this insight, this study introduces a practical decision-support approach, developed during the empirical work with Neo4j. this approach consist of constructing side-by-side comparisons, implemented internally, through spreadsheet-based analysis, linking:

- Organisational cloud distribution across cloud regions and CSPs
- Publicly available environmental indicators

While such comparisons inevitable contain gaps and uncertainties, they allow organisation to assemble what may be described as a fragmented but actionable sustainability-informed deployment landscape. This approach supports informed judgment under uncertainty, consistent with the bounded rationality perspective adopted in this thesis.

Due to confidentiality constraints, the full application of this approach within Neo4j cannot be presented. However, its logic is demonstrated using the GCP as a proxy case, Neo4j's predominant cloud usage. The most actionable sources were identified in Google's data centre efficiency disclosures and cloud region carbon intensity indicators (Google Cloud, n.d.; Google Data Centers, n.d.); these illustrate how region-level signals can be operationalised.

Importantly, this approach is not proposed as a prescriptive decision-making framework, but as a transferable analytical practice. It demonstrates that mid-sized SaaS firms can, within existing constraints, construct a sufficiently robust understanding of their sustainability context to inform incremental deployment decisions. This reinforces the central argument f this thesis: sustainability-informed decision-making in cloud environments does not depend on perfect information, but on the ability to navigate and operationalise imperfect, heterogenous data within a bounded decision space.

5.4.2 Conditions Required for Enhanced Feasibility

While limited forms of sustainability-informed actions are currently feasible, more ambitious deployment strategies, particularly those involving systematic region-level optimisation, remain contingent upon improvements that access all feasibility dimensions.

Informational feasibility emerges as the primary bottleneck. The absence of standardised, region-level environmental metrics across cloud providers significantly constrains the ability to compare deployment options in a methodologically robust manner. As demonstrated in Phase 1, CSPs' disclosures differ in granularity, definitions, systems boundaries, and accounting methods. This lack of comparability undermines the reliability of sustainability signals and prevents their consistent integration into decision-making processes.

Technical feasibility is present in certain workload categories, though it requires further development to support broader optimisation. This includes clever identification of workload mobility, integration of environmental indicators into deployment tools, and in some cases, architectural adjustments to enable flexibility.

Organisational feasibility depends on the extent to which sustainability can be institutionalised within existing processes. The findings indicate that sustainability is currently treated as an implicit or background consideration rather than an explicitly decision criterion. Enchanting feasibility would therefore require the development of internal reporting mechanisms, the translation of environmental metrics into operationally meaningful indicators, and the incorporation of sustainability into performance evaluation and planning frameworks.

Governance feasibility similarly requires transformation. At present, sustainability is not formally embedded within procurement or deployment policies; its integration would necessitate the alignment with external pressures, such as regulatory requirements, investor expectations, and customer demand, and the development of internal accountability structures.

Sustainability becomes actionable when it is rendered both visible and governable. That entails not merely the availability of data, but its standardisation, comparability, and integration into organisational decision-making structures.

5.4.3 Structural Constraints and Limits of Action

Despite the potential for incremental improvements, certain aspects of cloud deployment remain structurally constrained across all feasibility dimensions. These constraints define the outer boundaries of the decision space within which Neo4j operates.

Technical feasibility sees the relocation of customer-facing production workloads, based on prioritised environmental considerations, as severely limited in feasibility due to these workloads being inherently constrained by latency sensitivity, reliability, and SLAs.

Governance feasibility further restricts flexibility in terms of long-term contractual commitments with cloud providers, including spend-based agreements. This creates a form of lock-in that reduce the organisation's ability to switch providers or significantly reconfigure deployments. Additionally, multi-cloud environments introduce complexity that increased the cost of coordination and change.

Organisational feasibility also plays a role in terms of limited engineering capacity and competing task priorities that sustainability initiatives must compete with core business objectives. Consequently, actions that do not deliver immediate operational or financial benefits are less likely to be prioritised.

Informational feasibility persists, even with access to internal tooling. The lack of standardised, externally verifiable environmental data limits confidence in sustainability assessments and reduces the likelihood of decisive action.

These combined constraints indicate that sustainability cannot be treated as an independent optimisation variable. Instead, it is subordinate to performance, cost, and customer requirements, and can only be incorporated where these priorities permit.

5.4.4 Answering the Research Question: RQ2

The synthesis and overall interpretation of these findings suggest that Neo4j's sustainability decision-making is best understood as operating within a bounded feasibility space, rather than an optimization landscape. Within this space, sustainability consideration is neither absent nor dominant, but conditionally integrated based on the interaction of informational, technical, organisational, and governance factors.

This interpretation directly reflects the bounded rationality framework adopted in this thesis. Decision-makers do not optimise for sustainability in a strict sense; rather, they select options that are sufficiently aligned with existing constraints and objectives. Sustainability thus emerges as a context-dependent and contingent factor, influencing decisions at the margins rather than determining them.

RQ2

Sustainability-informed cloud deployment decisions at Neo4j are partially feasible, but only under specific conditions where sufficient informational clarity, workload flexibility, organisational capacity, and governance permissibility coexist. These decisions are inherently incremental and constrained, reflecting a broader pattern of *satisficing* under structural limitations rather than comprehensive optimisation.

6 Discussion: Linking Energy Systems, Digital Infrastructure, and Cloud Decision-Making Constraints

The findings of this thesis empirically support the conditional and critical strands of the literature reviewed in Chapter 2. While the technical potential for lower-carbon cloud computing clearly exists through carbon-aware scheduling, workload shifting, and renewable-energy optimisation, the empirical findings demonstrate that these possibilities remain constrained by incomplete transparency, organisational bounded rationality, and infrastructural dependencies, and governance asymmetries. The study therefore extends existing digital sustainability literature by showing that sustainability in cloud computing should be understood not as a purely technical optimisation problem, but as a socio-technical governance challenge emerging from the interaction between provider disclosure practices, energy systems, organisational feasibility, and hyperscale market structures.

These findings can be further contextualised within broader development in energy systems, digital infrastructure, and sustainability governance. Across the energy transition and digitalisation dimension, a consistent pattern emerged: challenges are no longer purely based on technological expansion, but also coordination, transparency, and system integration.

6.1 Answering the RQs

RQ1(a) *How transparent and comparable are the sustainability reports and associated public environmental disclosures of Amazon, Microsoft, and Google, regarding data centre energy efficiency and carbon performance?*

The analysis found that public CSP sustainability disclosure is only partially transparent and unevenly comparable. All three providers disclose meaningful sustainability information, including global efficiency indicators, renewable claims, and board decarbonisation strategies. However, the disclosure differs substantially in granularity, methodological clarity, reporting boundaries, and regional specificity.

GCP has the strongest public basis for regional carbon-related interpretation through its carbon-free energy data and location-linked disclosure (see Appendix B, Table 9). AWS and Azure disclose important sustainability information, but much of it remains at global, fleet, corporate, or tool-specific level rather than in a public region-comparable format. As a result, cross-provider comparison remains fragile. The issue is not simply that information is missing, but that providers disclose different objects, at different scales, and through different methods.

RQ1(b) *What provider-level and selected region-level-information on data centre energy efficiency and carbon performance is available for AWS, Azure, and GCP, and how actionable is this information for customer decision-making?*

Provider-level information is available for all three CSPs, but region-level actionability is limited. The strongest actionable public evidence is found in GCP's regional carbon-related disclosure, which can support some carbon-aware region-selection reasoning. AWS and Azure provide useful provider-level signals, and Microsoft appears comparatively strong in custom-facing tools and methodological sophistication, but neither offers the same publicly auditable region-level comparability in the reviewed corpus.

Nonetheless, given that carbon-aware region selection is based on the local energy grid, it is tempting to assume that two providers operating in the same geographic area are equally sustainable. However, this is a flawed assumption.

RQ2 *To what extent are sustainability-informed cloud deployment decisions possible for Neo4j within existing organisational, technical, and contractual constraints?*

For Neo4j, sustainability-informed deployment decisions are possible, but only within a bounded and conditional decisions space. The most feasible opportunities are internal, development testing, and canary workload, where customer-facing latency, contractual obligations, and SLA and expectation are less restrictive. New deployments also offer more realistic opportunities, rather than employing efforts into migrating existing workloads from one region to another.

While not within the scope of this research, customer-facing production environments have been identified as strongly constrained by customer requirements, latency, compliance and governance restrictions, provider commitments, and other commercial realities. Sustainability is acknowledged as relevant, but it is not yet embedded in operational decision-making in the same way cost and performance are. The findings therefore support a bounded-rationality interpretation: Neo4j cannot optimise cloud sustainability in an unconstrained way, but can make incremental, *satisficing* decisions where information on optimal regions, workload environment flexibility, and organisational capacity align.

6.2 From Disclosure to Decision: A Structural Coordination Problem

Phase 1 of the analysis demonstrated that sustainability information from cloud providers is available in unstructured ways disabling comparison or action. Phase 2 showed that even where sustainability is acknowledged, it cannot be operationalised due the informational, organisational, technical and governance constraints. The translation of information into action is constrained by systemic fragmentation; it is a coordination and governance problem. This systemic perspective directly reinforces the findings from RQ1(a,b) and RQ2, and the central contribution of this thesis: lower-carbon deployment options are technically possible, yet operationally dependent on coordination across infrastructure, energy systems, and organisational capabilities (Bandhavi Sakhamuri, 2025; Beena et al., 2025; Tripathi et al., 2023).

Recent energy system analyses indicate that we are witnessing a critical turning point in the global energy transition: a shift from a problem of supply expansion to one of system coordination (Net Zero Insights, 2026a). Historically, the focus was on building infrastructure to support supply expansion accommodating wind turbines and solar panels to replace fossil fuels. Now evidence shows focus on updating grids and storing energy for intermittent power produced to be used reliably. Renewable capacity continues to grow rapidly, with projections suggesting a massive influx of retiring systems, namely 209 Gigawatts (GW) capacities replacing 104 GW by 2030. Nonetheless, successful building capacity faces hurdles as grid congestion and infrastructural bottlenecks intensify. This highlights that the energy system is "not short of electrons, but short of flexibility" (Net Zero Insights, 2026a). *Flexibility*, defined by the ability to shift, store, or adapt energy demand and supply in real time, has emerged as a critical constraint (Net Zero Insights, 2026a). Importantly, this constraint mirrors the findings of this thesis. Just as energy systems struggle to coordinate distributed resources, cloud-dependent organisations struggle to translate sustainability information into actionable decisions. In both cases, the limiting factor is not availability of resources, but ability to align and operationalise under complex constraints. This parallel is particularly relevant for the argument made by Hankendi et al. (2025) that sustainability challenge of AI and cloud infrastructure is fundamentally constrained by transparency and coordination failures.

6.3 Reassessing Efficiency: From Metrics to System-Level Impact

The limitations identified in RQ1(a,b) and RQ2 further clarify the literature on data centre metrics and digital sustainability in Section 2.1.1. The environmental footprint of cloud computing is difficult to assess due to limited transparency and inconsistent reporting. Providers use different metrics, definitions of renewable energy, and reporting boundaries, making it challenging for organisations to evaluate their actual impacts (Jungblut, 2026a). This directly supports the findings of Phase 1, where provider-level metrics such as PUE, WUE, CUE, and renewable energy matching, were shown to be insufficient for region-level decision-making. As highlighted in Section 2.1.1, this directly supports de Vries-Gao (2026), Safari et al. (2025), and Long et al. (2022), who argue that these metrics provide simplified indicators of infrastructure efficiency, however, only capture partial aspects of performance: they are aggregate, methodologically inconsistent, and often fail to capture key contextual factors such as grid carbon intensity and workload characteristics.

The broader digital sustainability literature reinforces that cloud computing, and AI systems cannot be evaluated solely through efficiency improvements (Aslan et al., 2025; Bandhavi Sakhamuri, 2025; Lange et al., 2020). The literature reviewed in Sections 2.1.1 and 2.1.2 repeatedly warns that digitalisation and AI may generate rebound effects whereby efficiency gains are offset by expanding computational demand and correlating environmental impacts (Aslan et al., 2025; Lange et al., 2020). The assumption that cloud services are more environmentally beneficial due to shared infrastructure and economies of scale is increasingly contested; cloud infrastructure remains highly material and resource intensive (Bandhavi Sakhamuri, 2025; Hankendi et al., 2025). Data centres have substantially growing electricity, water, and material demands, particularly under AI-driven workloads altering energy demand patterns (Jungblut, 2026a). The findings strongly align with this perspective. For RQ1(b), this implies that even where providers disclose improvements in efficiency, these signals remain indicative rather than determinative, limiting their usefulness for customer decision-making. CSPs' disclosures consistently framed AI and cloud expansion through narratives of efficiency and renewable procurement, yet simultaneously revealed rapidly increasing infrastructure growth, energy demand, and water usage. This supports Lange et al.'s (2020) argument that digitalisation historically rends towards increased total energy consumption despite efficiency gains. Disclosures cannot be directly translated into actionable deployment decisions until they capture simultaneously regionally energy system variation, workload-specific impacts, and temporal fluctuations in carbon intensity. This reinforces the conclusion of Phase 1: that provider disclosure does not currently enable robust, region-level sustainability comparisons, and metrics are structurally misaligned with customer decision needs.

The findings challenge optimistic assumptions within parts of the green cloud computing literature, while studies such as Beena et al. (2025) and Bandhavi Sakhamuri (2025) demonstrate the technical potential of carbon-aware scheduling and AI-driven optimisation, the empirical findings indicate that such optimisation remains structurally disconnected form forms of disclosure available to customers. Neo4j could not operationalise many theoretically available optimisation opportunities precisely because the underlying environmental data was incomplete, aggregated, or non-comparable across providers.

Therefore, the literature is extended by demonstrating that the practical value of sustainability metrics is not determined by their technical sophistication alone, but also by their availability to support decision-making under real organisational conditions.

6.4 Transparency as a Structural Limitation

Across both the literature and the empirical findings, transparency emerges as the central constraint linking RQ1(a,b) and RQ2. Phase 1 showed that CSPs' disclosures are fragmented, aggregated, and inconsistently defined. Phase 2 demonstrates that this lack of transparency directly translates into organisational inaction, as sustainability cannot be integrated into decision-making processes without usable data.

This aligns with border findings in sustainability in cloud infrastructure and AI systems, and directly confirms arguments developed by Hankendi et al. (2025) and de Vries-Gao (2026). Environmental assessment of digital systems faces barriers by missing data, insufficient granularity, inconsistent methodologies, and proprietary reporting practices (Jungblut, 2026a). Particularly, this reinforces de Vries-Gao's (2026) argument that no major provider currently discloses AI-specific environmental metrics, forcing both researchers and customers to infer sustainability impacts indirectly through broader data centre metrics. Indeed, current estimates of emission often rely on partial disclosures of model-based approximations, and key lifecycle impacts such as hardware production and e-waste remain underreported (Jungblut, 2025). As a result, even the most widely used AI systems cannot be comprehensively evaluated from a sustainability perspective. Neo4j's inability to connect provider-level disclosure to workload-specific deployment decisions empirically demonstrates this exact problem: the absence of information aligned with customer-side decision units.

Similarly, the findings strongly support Hankendi et al.'s (2025) argument that transparency bottlenecks so both innovation and sustainability governance because actors are forced to rely on averages and corporate narratives, confirming how cloud customers remain dependent on provider-controlled sustainability information ecosystems, creating a structurally asymmetry in which providers control disclosure while customers are increasingly expected to operationalise sustainability objectives. This also extends the JRC analytical framework adopted in Section 2.2.1: the JRC Guidelines emphasise granularity, comparability, and boundary clarity as prerequisites for meaningful environmental reporting (Acton et al., 2025). The findings demonstrate that these dimensions are based on technical reporting preferences, with rooted foundational conditions for sustainability governance itself. Where providers lacked methodological clarity, regional specificity, or workload relevance, sustainability information became difficult to operationalise despite appearing extensive at corporate level.

This reinforces a central conclusion of this thesis: sustainability in cloud computing is not limited by technical potential, but by informational constraints and governance structures. Providers control of the production, definition, and disclosure of sustainability data, and customers are expected to act on it. This creates a governance asymmetry, where information is controlled upstream, and responsibility is pushed downstream. This misalignment explains the empirical findings of Phase 2 and strengthens the thesis' core contribution: sustainability decision-making in cloud environments is constrained by structural information asymmetry, not simply organisational priorities.

6.5 Bounded Reality in Practice: Conditional Feasibility

The bounded rationality framework presented in Section 2.2.2 is a robust lens through which to interpret Phase 2 findings. Rather than optimising sustainability outcomes, Neo4j operates within a feasibility envelope, where decisions are shaped by the interaction of informational, technical, organisational, and governance constraints. This directly confirms the bounded rationality assumption embedded in the conceptual framework: decisions emerge through satisficing under uncertainty, not through optimisation under perfect information.

This is consistent with the literature on carbon-aware computing identified in Section 2.1.2; while theoretical models demonstrated that workload shifting and carbon-aware scheduling can reduce emissions, these models assume access to real-time carbon data, controllable workloads, and low migrations costs (Bandhavi Sakhamuri, 2025; Beena et al., 2025; Tripathi et al., 2023). Phase 2 shows that these assumptions rarely hold in practice, and *feasibility* is conditional to:

- Informational feasibility, limited by incomplete and on-comparable CSP data
- Technical feasibility, constrained by workload characteristics like latency sensitivity
- Organisational feasibility, shaped by cost priorities and engineer capacity
- Governance feasibility, restricted by contractual commitments and provider lock-in

This confirms that sustainability-informed deployment decisions are selectively feasible, emerging only in contexts where the four feasibility dimensions align. This explains why sustainability-informed deployment potentials can be conceived primarily in internal or canary environments, new deployments, rather than migrations.

Similarly, the findings reinforce Jaensson & Wiklund's (2022) and Alonso et al.'s (2023) arguments regarding the practical limits of multi-cloud flexibility. Although multi-cloud architectures theoretically increase customer agency and deployment flexibility, the Neo4j case demonstrates that portability, interoperability, engineering complexity, and vendor-specific dependencies significantly constrain actual strategic freedom. The findings therefore support the literature's argument that customer agency exists but is unevenly distributed and organisationally expensive.

These findings refine the literature on carbon-aware computing, identifying the condition under which such strategies become operationally viable.

6.6 Flexibility as the Missing Link: Energy Systems and Digital Workloads

The findings support the literature reviewed in Section 2.1.2 concerning the growing importance of flexibility in sustainable computing systems. Studies on carbon-aware scheduling consistently argue that workload mobility, temporal shifting, and intelligent resource allocation may reduce emission without requiring additional physical infrastructure (Bandhavi Sakhamuri, 2025; Tripathi et al., 2023). The empirical findings partially confirm this potential: internal and canary development were identified as relatively feasible spaces for experimentation with sustainability-informed optimisation, due to their greater workload flexibility and lower operational risk.

However, the findings also challenge the literature by demonstrating that flexibility itself is unevenly distrusted across organisational contexts. While hyperscalers may increasingly operationalise large-scale demand flexibility through AI-driven orchestration and grid-aware computing, Neo4j's ability to exploit these possibilities remained constrained by workload dependencies, SLA obligations, organisational properties, and informational limitations. This extends the literature by showing that flexibility is not simply a technical property of workloads, but a socio-technical capability.

The concept of flexibility provides a critical bridge between energy systems and cloud computing. In energy systems flexibility refers to the ability to shift, store, or adapt demand in response to system conditions; in digital systems this translates into workload mobility, scheduling and resource allocation (Net Zero Insights, 2026a). The increasing role of demand-side flexibility in both energy and digital systems. In energy systems, demand flexibility, e.g. demand response and virtual power plants can scale more rapidly than infrastructure-based solutions because it avoids physical constraints like permitting and interconnection (Net Zero Insights, 2026a). Emerging development illustrates this concern. For instance, grid-aware systems and carbon-aware computing projects dynamically adjust workload execution based on real-time carbon intensity, adapted to align with energy system conditions and reduce energy consumption during high-carbon periods (Jungblut, 2026b). Hyperscalers such as Google are increasingly integrating demand flexibility into their operation, shifting workloads and participating in large-scale demand response mechanisms to align computing with energy availability (Net Zero Insights, 2026b).

These developments suggest that flexibility in digital systems is becoming a key mechanism for reducing environmental impact. However, they also highlight a critical limitation: such approaches require access to real-time data, advanced coordination systems, and organisational capacity to act on this information. As shown in Phase 2, such flexibility is unevenly destructed, and while hyperscalers can operationalise flexibility at scale, these conditions are often not met in mid-sized SaaS organisations. This creates a divergence between system-level optimisation potential and organisational-level feasibility. Flexibility is thus a function of organisational position within the infrastructure hierarchy.

6.7 Power, Market Structure, Governance, and the Limits of Customer Responsibility

The findings also reinforce the critical perspective identified in Section 2.1.4 that structural asymmetries within hyperscale cloud markets fundamentally constrain customer agency. The dominance of a small number of CSPs creates concentrated control not only over infrastructure, but also over sustainability narratives, environmental metrics, and disclosure boundaries. This strongly confirms the arguments advance by de Vries-Gao's (2026), Hankendi et al.'s (2025), and the broader transparency literature.

This perspective must be understood in the context of broader market and governance structure. The discussion highlights the overall issue of power concentration in digital infrastructure; the dominance of a small number of hyperscale providers creates structural asymmetries in control, transparency, and accountability. These firms control not only infrastructure, but also the data and tools used to evaluate sustainability performance and have strong influence over the reporting standards. Customers operate under dependency (O'Sullivan, 2025). This concentration of power has several implications. First, it limits the ability of customers to independently assess environmental impact. Second, it creates barriers to alternative, more transparent or sustainable solutions. Third, it requires coordinated action across industry, policy, and civil society, including mandatory reporting standards, lifecycle assessments frameworks, and transparency requirements (Jungblut, 2025; Lucks, 2026).

The Neo4j case demonstrates that even highly motivated organisations remain dependent on provider-controlled infrastructure of disclosure and operational possibility. Sustainability therefore cannot be realistically decentralised entirely to cloud customers without corresponding changes in provider disclosure and regulatory frameworks. The findings suggest that customer responsibility narratives may overestimate the degree of agency available to firms operating within hyperscale cloud ecosystems characterised by vendor lock-in, interoperability challenges, and asymmetric information structures.

This finding extends the literature by connecting digital sustainability more explicitly to infrastructure governance and political economy. Sustainability in cloud computing emerges as a question of power relations within digital infrastructure systems.

7 Conclusion

7.1 Research Summary and Synthesis of the Findings

This thesis set out to examine how sustainability-informed cloud deployment decisions can be understood from the perspective of mid-sized SaaS companies operating under indirect control over cloud infrastructure. The study investigated the relationship between provider-side disclosure and customer-side decision feasibility. The central purpose was to assess whether publicly available environmental information from Amazon, Microsoft, and Google specific to their cloud services AWS, Azure, and GCP, can meaningfully support deployment decisions, and how far such information can be acted upon within Neo4j's organisation constraints.

The findings show that the main challenge is not the complete absence of sustainability information, nor the absence of technical opportunities for lower-carbon cloud use. Instead, the central challenge in translation: transform partial, aggregated, and uneven provider disclosures into decisions that are feasible for cloud customers. This conclusion is important because it shifts the discussion away from a simple expectation that firms should choose greener regions and towards a more realistic understanding of sustainability decision-making under constraint. The findings suggest that cloud sustainability sits at the intersection of three tensions.

(1) *Transparency tension*; CSPs increasingly disclose environmental information but much of it is not structured around the decision customers need to make when choosing providers and regions
(2) *Agency tension*; cloud customers are increasingly expected to account for and reduce emission, yet they do not control the underlying infrastructure, electricity procurements, operational energy use, and disclosure practices. Their agency exists, but is parietal and uneven
(3) *Feasibility tension*; carbon-aware deployment is technically imaginable, but it becomes practically relevant only where workloads are moveable, environmental signals are trustworthy, and trade-offs with cost, latency, compliance, and productivity are acceptable

The thesis therefore concludes that sustainable cloud decision-making is not primarily a matter of identifying one "greenest" provider or region; it is a governance problem: how to make environment information visible, comparable, usable enough, and upon these conditions also bounding, to enter mid-sized SaaS organisational realities.

Ultimately, sustainability in cloud computing is a relational problem, emerging at the intersection of provider disclosure practices and supply chain transparency, organisational constraints within cloud-dependent firms, and system-level dynamics of governance and coordination.

7.2 Contribution to Knowledge

This thesis contributes several insights to the emerging literature on digital sustainability, cloud infrastructure governance, AI-driven energy demands, and organisational decision-making.

The first contribution regards the link between provider transparency and customer-side feasibility. Much existing literature focuses either on data centre metrics and carbon-aware computing, or on sustainability reporting and disclosure, but rarely these reflect thoroughly upon how these blockers impede a smooth and informed sustainability transition for cloud deployment practices. This thesis connects these debates by showing that disclosure only matters for customers if it becomes decision relevant.

The second contribution reflects the application of bounded rationality to cloud sustainability decision-making. This analytical lens supports the understanding of why potential technical sustainability action may not be organisationally feasible. In doing so, the thesis moves being idealised optimisation models and contributes a more realistic account for decision-making under incomplete information and competing priorities.

The third contribution provides empirical insights not the position of mid-sized SaaS companies. As many discussions of cloud sustainability focus on either hyperscale providers or technical systems, Neo4j illustrates the intermediate position of firms that are environmentally accountable on cloud but constrained in means of influence.

7.3 Recommendations

The *practical contributions* of this thesis are provided through a grounded understanding of where sustainability-informed cloud decisions are realistically possible, how organisations can integrate sustainability into existing decision processes, and what limits such integration in practice. Consequently, the research yields actionable insights for SaaS enterprises navigating environmental sustainability within the constraints of infrastructure dependency. These results are paired with formal calls for heightened transparency from CSPs and the imposition of more robust regulatory frameworks by policymakers.

For **cloud customers** and **mid-sized SaaS firms**, the main recommendation is to begin where feasibility is highest. Sustainability should first be integrated into internal, development, testing, canary, and batch-like workloads, rather than customer-facing production. Firms should build internal visibility by tracking cloud usage by region, workload type, provide, and business function, and this information should be widespread to the directly interested parties involved with the corresponding business unit. The alignment between sustainability and cost optimisation should be used strategically, while recognising that cost is only an imperfect proxy for environmental performance. Introducing a centralised coordination of sustainability management allows for traceability of these sustainability reduction strategies, and with the appropriate resource allocations, sustainability-informed cloud deployment choices can become embedded into the current frameworks of cloud provider and region selection.

For **cloud providers**, the main recommendation is to make sustainability disclosure more standardised, enforcing decision-relevant elements. This requires region-level indicators, clever methodological boundaries, standardised reporting formula, and customer-facing tools whose assumptions are transparent. Providers should distinguish more clearly between corporate sustainability claims, data centre performance metrics, and other deployment signals to assist customers in making informed decisions. By doing so, providers will be further encouraged to better their performance in "less green" regions, ultimately reducing the risk of carbon leakage in data centre regions with lower carbon prices or weak climate policies and mitigating the net increase in global emissions.

For **policymakers** and **standard-setting bodies**, the findings support the needs for stronger harmonisation of data centre reporting accessible to the public. Existing frameworks such as the EU CoC and JRC Best Practices Guidelines provide a strong foundation, but customer-side usefulness should become a more explicitly reporting objective, and while respecting the desire for businesses to maintain proprietary information, data should be more openly shared to enhance sustainability-informed decisions. Regulatory efforts should encourage disclosure that is comparable, auditable, and linked to the level at which cloud decisions are made.

The *academic contributions* of this thesis are illustrated through the contributions to literature as linking provider transparency to customer feasibility, introducing the concept of decision-composable disclosure, and empirically grounding the concept of bounded agency in cloud sustainability. This fosters awareness of broadening the perspectives for sustainability governance in digital infrastructure systems.

7.4 Limitations and Future Research

The study has several limitations. The analysis relies on public CSP disclosure, which are self-reported and unevenly structured. The Neo4j case is based on a limited number of consultations and internal contextual insights, meaning the findings should be treated as analytically transferable rather than statistically generalisable. Moreover, as the thesis does not calculate actual emissions reductions or conduct any full LCAs, its contribution lies in assessing decision feasibility, not in quantifying environmental impacts.

Further research cloud extends this work into a larger comparative study of SaaS firms, testing whether feasibility patterns identified here apply across different organisational sizes and sectors, and how sustainability information plays decisive roles in operational priorities and decision-making. Further research could also evaluate customer-facing cloud carbon tools in more detail, comparing their methods, outputs, and usefulness for deployment decisions. Another important direction would be to examine AI-specific cloud workloads, especially where GPU use, inference demand, and data storage create new sustainability pressures. Finally, future work could develop and test practical governance models for integrating sustainability into cloud architectures, procurement, and platform engineering decisions.

7.5 Final Reflection

The thesis begins from a practical problem: mid-sized Saas firms are expected to reduce environmental impacts and their emissions; while relying primarily on cloud infrastructure they do not fully control or have significant influence over.

The research shows that this problem cannot be solved simply by better intentions or by assuming that CSPs' efficiency gains and lower cost options automatically translate into customer actions based on assumed direct correlation with better environmental performance. Sustainability-informed cloud deployment requires better disclosure from CSPs, clearer internal organisational decision processes within SaaS firms, and a realistic understanding of where customer agency begins and ends both for when enforcing environmental targets, and for enforcing further transparency upon CSPs.

The most important conclusion is therefore not that cloud customers' power is conditional; they can act where information is usable, workloads are flexible, and governance structure allows sustainability to become part of decision-making. Until provider reporting becomes more granular and comparable, sustainability in cloud deployment will remain possible mainly at the margins. Those margins still matter, but they must be understood as incremental steps within a constrained system and not as full control over the environmental footprint of the cloud. We must redesign the system so that those who must act are finally given the power, and the information, to do so.

If the cloud remains, as is commonly said, 'just someone else's computer', sustainability in the cloud is ultimately someone else's decision.

Bibliography

Acton, M., Booth, J. & Paci, D. (2025). *2025 Best Practice Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency*. (JRC141521). Publications Office of the European Union. https://data.europa.eu/doi/10.2760/9449356.

Al-Sharafi, M. A., Iranmanesh, M., Al-Emran, M., Alzahrani, A. I., Herzallah, F., & Jamil, N. (2023). Determinants of cloud computing integration and its impact on sustainable performance in SMEs: An empirical investigation using the SEM-ANN approach. *Heliyon, 9(5)*, e16299. https://doi.org/10.1016/j.heliyon.2023.e16299.

Alexandre. (2025). *Cloud Market share 2026: Top cloud providers and trends*. Holori. https://holori.com/cloud-market-share-2026-top-cloud-vendors-in-2026/.

Alonso, J., Orue-Echevarria, L., Casola, V., Torre, A. I., Huarte, M., Osaba, E., & Lobo, J. L. (2023). Understanding the challenges and novel architectural models of multi-cloud native applications: A systematic literature review. *Journal of Cloud Computing, 12*, Article 6. https://doi.org/10.1186/s13677-022-00367-6.

Amazon. (2024). *Amazon Sustainability Report 2024*. https://sustainability.aboutamazon.com/reports.

Arshi, P., & Miller, J. (2025). *Our approach to carbon-aware data centers: Central data center fleet management*. Google Cloud Blog. https://cloud.google.com/blog/topics/sustainability/googles-approach-to-carbon-aware-data-center.

Aslan, T., Holzapfel, P., Stobbe, L., Grimm, A., Nissen, N. F., & Finkbeiner, M. (2025). Toward climate neutral data centers: Greenhouse gas inventory, scenarios, and strategies. *iScience 28*, 111637. https://doi.org/10.1016/j.isci.2024.111637.

AWS. (2024). Amazon Sustainability Report: AWS Summary. https://sustainability.aboutamazon.com/reports.

AWS Sustainability. (n.d.). *AWS Sustainable Infrastructure*. AWS Sustainability Data Centers. https://aws.amazon.com/sustainability/data-centers/.

Bandhavi Sakhamuri, N. S. (2025). Sustainable cloud infrastructure: AI-driven carbon-aware kubernetes scheduling and resource management. *World Journal of Advanced Research and Reviews*, *26(2)*, 2138-2145. https://doi.org/10.30574/wjarr.2025.26.2.1854.

Beena, B. M., Ranga, P. C., Manideep, T. S. S., Saragadam, S., & Karthik, G. (2025). A Green Cloud-Based Framework for Energy-Efficient Task Scheduling Using Carbon Intensity Data for Heterogeneous Cloud Servers. *IEEE Access*, *13*, 73916-73938. doi:10.1109/ACCESS.2025.3562882.

Borra, P. (2024). An Overview of Cloud Computing and Leading Cloud Service Providers. *International Journal of Computer Engineering and Technology* (IJCET). *15(3)*, 122-133. DOI: 10.17605/OSF.IO/5HQ4M.

de Vries-Gao, A. (2026). The carbon and water footprints of data centers and what this could mean for artificial intelligence. *Patterns, 7*, 101430. https://doi.org/10.1016/j.patter.2025.101430

Directorate-General for Energy. (2025). *In focus: Data centres – an energy-hungry challenge. Energy, Climate change, and Environment*. European Commission. https://energy.ec.europa.eu/news/focus-data-centres-energy-hungry-challenge-2025-11-17_en.

European Commission: Directorate-General for Energy; Ernest & Young (EY); Austrian Institute of Technology (AIT); & Borderstep. (2025). *Assessment of the energy performance and sustainability of data centres in EU: First technical report*. Publications Office of the European Union. https://data.europa.eu/doi/10.2833/3168794.

European Commission (EC). (n.d.). *Energy performance of data centres*. https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-directive/energy-performance-data-centres_en.

Ewim, D., Ninduwezuor-Ehiobu, N., Orikpete, O., Egbokhaebho, B., Fawole, A., & Onunka, C. (2023). Impact of Data Centers on Climate Change: A Review of Energy Efficient Strategies. *The Journal of Engineering and Exact Sciences*. *9*, 16397-01e. 10.18540/jcecvl9iss6pp16397-01e.

Google. (2025). *Google Environmental Report 2025*. https://sustainability.google/reports/google-2025-environmental-report/.

Google Cloud. (n.d.). *Carbon free energy for Google Cloud regions*. https://cloud.google.com/sustainability/region-carbon.

Google Cloud Blog. (2021). *How carbon-free is your cloud? New data lets you know*. https://cloud.google.com/blog/topics/sustainability/sharing-carbon-free-energy-percentage-for-google-cloud-regions.

Google Data Centers. (n.d.-a). *Driving sustainable innovation for society*. https://datacenters.google/operating-sustainably/.

Google Data Centers. (n.d.-b). *Growing the internet while reducing energy consumption*. https://datacenters.google/efficiency/.

Google Sustainability. (n.d.). *Innovating across our operations and supply chains*. https://sustainability.google/operations/.

Hankendi, C., Coskun, A. K., & Sovacool, B. K. (2025). Why transparency matters for sustainable data centres and carbon-neutral artificial intelligence (AI). *iScience 28*, 113705. https://doi.org/10.1016/j.isci.2025.113705.

Hasbe, S. (2026). *2025: Year of AI and scalability*. Neo4j Blog. https://neo4j.com/blog/news/2025-ai-scalability/.

International Energy Agency (IEA). (2025). *Energy and AI*. IEA Publications. World Energy Outlook Special Report. https://www.iea.org/reports/energy-and-ai.

Jaensson, P., & Wiklund, O. (2022). *Handling multiple cloud service providers - Common challenges and best practices*. [Master's thesis, Lund University].

Jungblut, S.-I. (2025). *Sustainable AI means looking beyond data centres*. RESET. https://en.reset.org/sustainable-ai-means-looking-beyond-data-centers/.

Jungblut, S.-I. (2026a). *Why do our clouds look so gloomy, and how can we make the switch to a green cloud?* RESET. https://en.reset.org/why-do-our-clouds-look-so-gloomy-and-how-can-we-make-the-switch-to-a-green-cloud/.

Jungblut, S.-I. (2026b). *Grid-aware websites respond to the power grid*. RESET. https://en.reset.org/grid-aware-websites-respond-to-the-power-grid/.

Kandasamy, A. (2025). *More data centers, more environmental problems?* Conservation. National Wildlife Federation. https://www.nwf.org/Magazines/National-Wildlife/2025/Fall/Conservation/AI-Data-Centers.

Kumar, A. (2026). *AWS vs Azure vs GCP: Which Cloud is Best for Businesses in 2026?*. teleGlobal, elevating business to technology. https://teleglobals.com/blog/aws-vs-azure-vs-gcp

Lal, A. & You, F. (2025). Advances and challenges in energy and climate alignment of AI infrastructure expansion. *Advances in Applied Energy*. *Volume 20*, 100243. https://doi.org/10.1016/j.adapen.2025.100243.

Lange, S., Pohl, J., & Santarius, T. (2020). Digitalization and energy consumption. Does ICT reduce energy demand?. *Ecological Economics*. *Volume 176*, 106760. https://doi.org/10.1016/j.ecolecon.2020.106760.

Leoaspauza. (2026). *Advancing embodied carbon measurement at scale for Microsoft Azure hardware*. Azure Infrastructure Blog. https://techcommunity.microsoft.com/blog/azureinfrastructureblog/advancing-embodied-carbon-measurement-at-scale-for-microsoft-azure-hardware/4485784.

Long, S., Li, Y., Huang, J., Li, Z., & Li, Y. (2022). A review of energy efficiency evaluation technologies in cloud data centers. *Energy and Buildings*. *Volume 260*, 111848. https://doi.org/10.1016/j.enbuild.2022.111848.

Lucks, B. (2026). ECO:DIGIT shows how the environmental impact of software can be measured. RESET. https://en.reset.org/ecodigit-shows-how-the-environmental-impact-of-software-can-be-measured/.

Mell, P. & Grance, T. (2011). *The NIST Definition of Cloud Computing*. NIST Special Publication 800-145. National Institute of Standards and Technology.

Mencaroni, A., Leyman, P., Raa, B., De Vuyst, S., & Claeys, D. (2025). Towards net-zero manufacturing: Carbon-aware scheduling for GHG emissions reduction. *Journal of Cleaner Production, 529,* 146787. https://doi.org/10.1016/j.jclepro.2025.146787.

Microsoft. (2025). *2025 Environmental Sustainability Report. Accelerating progress to 2030.* https://www.microsoft.com/en-us/corporate-responsibility/sustainability/report/.

Microsoft. (n.d.-a). *Microsoft for Sustainability*. https://www.microsoft.com/en-us/sustainability/cloud.

Microsoft. (n.d.-b). *Powering sustainable transformation. Microsoft Datacenters*. https://datacenters.microsoft.com/globe/powering-sustainable-transformation/.

Microsoft Reporter. (2020). *Microsoft announces investments to accelerate Sweden's digital transformation and plans to open its sustainable datacenter region in 2021*. https://news.microsoft.com/europe/2020/11/24/microsoft-announces-investments-to-accelerate-swedens-digital-transformation-and-plans-to-open-its-sustainable-datacenter-region-in-2021/.

Neo4j. (n.d.-a). *Neo4j: Graph Intelligence Platform*. Company https://neo4j.com/company/.

Neo4j. (n.d.-b). *Neo4j Aura*. Fully Managed Cloud Service. https://neo4j.com/cloud/platform/.

Neo4j. (n.d.-c). *Aura with cloud provider marketplace*. https://neo4j.com/docs/aura/cloud-providers/.

Neo4j. (n.d.-d). *Neo4j Aura DB - Frequently Asked Question*s. https://neo4j.com/cloud/platform/aura-graph-database/faq/.

Neo4j. (n.d.-e). *Why Neo4j for Knowledge Graphs*. Use Cases. https://neo4j.com/use-cases/knowledge-graph/.

Neo4j informal consultations (2026). Internal consultations conducted with engineering and platform stakeholders, March-April 2026:

- Consultation #1, Neo4j - Aura Software Engineer, 26 March 2026
- Consultation #2, Neo4j - FP&A Director, 15 April 2026

Net Zero Insights. (2025). *The Environmental Cost of Data Centres*. Climate Tech 101. https://netzeroinsights.com/resources/data-centers-environmental-cost/.

Net Zero Insights. (2026a). *Where the grid gets flexible: What the investment data tells us about the energy transition's next frontier*. https://netzeroinsights.com/resources/where-the-grid-gets-flexible/.

Net Zero Insights. (2026b). *How Google is building a flexible grid to power AI*. https://netzeroinsights.com/resources/how-google-building-flexible-grid-power-ai/.

Nlyte Software. (2023). *White Paper: Fundamental Measures of Data Center Sustainability. Carrier Global Corporation.*

O'Sullivan, L. (2025). *EU alternatives to Big Tech: Why local is the future*. RESET. https://en.reset.org/eu-alternatives-to-big-tech-why-local-is-the-future/.

Pardalis, G., Palm, J., & Reindl, K. (2026). The cloud has wires: Toward circular governance of smart grids. *Journal of Cleaner Production*. *Volume 556*. 148238. https://doi.org/10.1016/j.jclepro.2026.148238.

Pavlinich, E. J. (2026). *The Dangers of Data Centers.* Environmental Health Project. https://www.environmentalhealthproject.org/post/the-dangers-of-data-centers.

Safari, A., Sorouri, H., Rahimi, A., & Oshnoei, A. (2025). A Systematic Review of Energy Efficiency Metrics for Optimizing Cloud Data Center Operations and Management. *Electronics*, *14 (11)*, 2214. https://doi.org/10.3390/electronics14112214.

Sun, H. (2022). Resource Deployment and Task Scheduling Based on Cloud Computing. *2022 2nd International Conference on Computer Systems (ICCS)*. IEEE. DOI: 10.1109/ICCS56273.2022.9988014.

Taylor, A. (2022). *Improve your energy and carbon efficiency with Azure sustainability guidance*. Microsoft Azure Blog. https://azure.microsoft.com/en-us/blog/improve-your-energy-and-carbon-efficiency-with-azure-sustainability-guidance/.

Tripathi, S., Kumar, P., Gupta, P., Misra, R., & Singh, T. N. (2023). Workload Shifting Based on Low Carbon Intensity Periods: A Framework for Reducing Carbon Emissions in Cloud Computing. *2023 IEEE International Conference on Big Data, 2023*. 3387-3395. COI: 10.1109/BigData59044.2023.10386447.

Viegand Maagøe. (2022). *Energy efficiency metrics for data centres: Final report*. IEA 4E EDNA. https://www.iea-4e.org/wp-content/uploads/2022/10/EDNA-Studies-Metrics-for-data-centre-efficiency-Final.pdf.

Walsh-Elwell, N. (2020). *Achieving 100 percent renewable energy with 24/7 monitoring in Microsoft Sweden*. https://azure.microsoft.com/en-us/blog/achieving-100-percent-renewable-energy-with-247-monitoring-in-microsoft-sweden/.

Walsh-Elwell, N. (2022). *How Microsoft measures datacenter water and energy use to improve Azure Cloud sustainability*. Microsoft Azure Blog. https://azure.microsoft.com/en-us/blog/how-microsoft-measures-datacenter-water-and-energy-use-to-improve-azure-cloud-sustainability/.

Yañez-Barnuevo, M. (2025). *Data Centers and Water Consumption.* Environmental and Energy Study Institute. https://www.eesi.org/articles/view/data-centers-and-water-consumption.

Appendix A. Glossary

The Glossary presented in table 7 informs the significance of various technical terminologies and chosen thesis definitions utilised throughout this research; all terms described are represented as italicised in text.

Table 7. Glossary

Term	***Description***
Bounded agency	The limited capacity of cloud customers to influence the environmental performance of cloud infrastructure due to constraints such as incomplete information, provider dependency, contractual lock-in, technical limitations, and organisational priorities. In this thesis bounded agency describes how organisation can influence sustainability outcomes only within restricted acres of control (Author's thesis definition built on the *bounded rationality theory* of Herbert A. Simon).
Carbon intensity	Carbon intensity refers to the amount of CO_2 emissions produced per unit of energy consumed, fluctuating based on regional energy sources and time-based variations (Beena et al., 2025).
Cloud integration	Cloud integration is the synchronisation of data between local on-premises servers and remote SaaS applications and cloud services. Where companies are expected to employ efficient and innovative technologies to maintain adequate levels of productivity within their ecosystem, cloud computing provides a cost-effective and flexible solution to achieve these objectives (Al-Sharafi et al., 2023).
Comparability	Comparability is the extent to which the disclosed data can support like-for-like comparison across CSPs and across regions (Author's thesis definition, built on the JRC Best Practices Guidelines).
Decision relevance	Decision relevance is the extent to which disclosed CSP information can support a mid-sized SaaS firm's deployment choice, especially region selection (Author's thesis definition, built on the JRC Best Practices Guidelines).
Cloud deployment	The placement, configuration, and operation of workloads, applications, and services across cloud infrastructure environments and regions. (Neo4j Consultation #1 - Aura, 26 Mar. 2026). In this thesis, cloud deployment specifically refers to organisational decisions regarding workload placement across CSPs and regions.
Energy efficiency evaluation	At present, energy efficiency evaluation is completed by short-term on-site measurements and engineering calculations. These are called "Measurement-based evaluations" and are the most intuitive and accurate methods. "Simulation-based" and "Analytical Modeling-based" are alternative methods, but with much lower accuracy (Long et al., 2022).
Energy efficiency Metrics granularity	Energy efficiency metrics category separation, presented by Long et al. (2022): Coarse-granularity metrics are those proposed at the data centre level, generally rough estimates but easy to use as they employ only one metric and eliminate need for several details. PUE is a typical example, assessing the energy efficiency as a ratio of total data centre energy consumption to energy consumption of IT equipment. Other metrics include: DCiE, DCeP, CADE, TUE, and DCPE. However, details of sub-components are lacking, and as such these evaluation results cannot be used to advise data centre energy savings.

Term	Description
	Medium-granularity targets data centre components, giving information for specific equipment, infrastructure, and green energy in data centres, resulting in more accurate and extensive results due to their diversity and pertinence. Namely, the metrics include: for IT equipment ITEE, SCE, DWPE, EPI, ITEU, TGI, CNEE, NPUE, COP, RER; for cooling equipment DCCSE, MLC, CLF; and for green energy WUE, CUE, ERE, GEC. Fine-granularity provides detailed performance information for specific IT equipment components, responsible for normal operation of data centres, i.e. CPU, disk, memory, network, and storage. More specifically, the metrics include: for CPU utilisation, CPU workload, Balance Factors, and Waiting Time; for disk Usage Rate; for memory network Commands Per Second, Memory Usage, and Throughput; and for storage Data Transmission Speed and Storage Usage Average. These metrics alone cannot represent directly the data centre energy efficiency, but can be evaluated by looking at changes in performance and energy consumptions, measured as performance/watt. This is the most significant metrics group for cloud operators.
Hyperscale data centres	Hyperscale data centres perform highly scalable operations, and handle extensive workloads through an optimised networks architecture. These centres typically contain 5,000+ servers, with extensive networking infrastructure (Net Zero Insights, 2025).
Hyperscale providers	Generally, hyperscale data centres are owned by Tech giants, which can also be referred to as Hyperscale Providers (Net Zero Insights, 2025). These are companies capable of delivering computing resources at massive scale through geographically distributed facilities, interconnected via high-capacity networks.
Internal, canary, & external environments	The internal, canary, and external are the three stages in a software's lifecycle. Each environment balances the development with risk. Each environment ensures that a new code is tested by developers, validated by small user groups, and finally released to the public safely (Neo4j Consultation #1 - Aura, 26 Mar. 2026).
Low-latency	Systems or infrastructure designed to minimise delay between a user request and system response. In cloud deployment, low-latency requirements often constrain workload mobility because applications must remain geographically close to users (Author's thesis definition, informed by cloud computing literature).
Measurement-based evaluations	Measurement-based evaluations require specialised equipment, e.g. using the Power Distribution Unit (PDU) to measure the power consumption of IT devices, and specialised evaluation metrics. PUE is frequently used in the industry as an easy-to-use metrics; however, it may only evaluate a data centre in a rough sense. In the realm, there are no approved measurement methods (Long et al., 2022).
Post-deployment tooling	Post-deployment tooling refers to software, scripts, and platforms dedicated to validating, monitoring, and managing applications and infrastructure immediately after they are released into a production environment (Microsoft, 2025).
Redundancy	The duplication of critical infrastructures, systems, or components to ensure service continuity and fault tolerance in the event of failure. Cloud regions typically include redundant systems to maintain reliability an uptime (Author's thesis definition, informed by cloud computing literature).
Satisficing	A decision-making logic associated with bounded rationality in which organisations select options that are "good enough" under existing contractions, rather than fully optimising outcomes. In this thesis, cloud sustainability decisions are interpreted as satisficing behaviour under informational, technical, organisational, and governance

Term	Description
	limitations (Author's thesis definition built on the *bounded rationality theory* of Herbert A. Simon).
Service Level Agreement (SLA)	A formal contractual agreement defining service standards between a provider and customer, including uptime, latency, reliability, and performance guarantees; SLA requirements often constrain sustainability-informed workload relocation decisions (Neo4j, Sustainability Specialist, personal communication, January 28, 2026).
Transparency	Transparency is the extent to which Amazon, Microsoft, and Google disclose environmental performance information with sufficient availability, granularity, methodological standardisation, and boundary clarity to permit informed interpretation (Author's thesis definition, built on the JRC Best Practices Guidelines).

Source: Author's elaboration

Appendix B. Data Corpus Phase 1, List of Resources

The list of resources provides a structure overview of all the document corpus analysed in Phase 1 (Chapter 4). In line with Section 3.3.1, the corpus is divided into two categories:

(1) The *report corpus*, i.e. annual sustainability or environmental reporting forming the primary basis for the T-C-D coding analysis
(2) The *expanded public disclosure corpus*, i.e. the supplementary material used to triangulate findings and assess transparency, including provider sustainability webpages, technical blogs and methodology notes, infrastructure and data centre pages, and region-specific sustainability disclosures.

This appendix enhances traceability and reproductivity of the analysis by enabling direct access to all sources used.

Table 8. List of Resources: Report Corpus

Provider	***Document Name***	***Year***	***Coverage Period***	***Source***	***Citation***
Amazon	Amazon Sustainability Report	2024	2024 fiscal year (Jan.1 - Dec.31, 2024)	https://sustainability.aboutamazon.com/reports	Amazon. (2024).
Microsoft	Environmental Sustainability Report 2025	2025	2024 fiscal year (Jan.1 - Dec.31, 2024)	https://www.microsoft.com/en-us/corporate-responsibility/sustainability/report/	Microsoft. (2025).
Google	2025 Environmental Report	2025	2024 fiscal year (Jan.1 - Dec.31, 2024)	https://sustainability.google/reports/google-2025-environmental-report/	Google. (2025).

Source: Author's elaboration based on data corpus in Phase 1 (Section 3.1.1).

Table 9. List of Resources: Extended Public Disclosure Corpus

Amazon Web Services				
Source Type	***Source Name***	***Description***	***Link***	***Citation***
Report	2024 Amazon Sustainability Report: AWS Summary	Global efficiency metrics (PUE, WUE, etc.); design practices; reducing embodied carbon	https://sustainability.aboutamazon.com/products-services/aws-cloud	AWS. (2024).
Enterprise webpage	AWS Sustainability	General cloud sustainability strategy and claims	https://aws.amazon.com/sustainability/	AWS Sustainability. (n.d.-a).
Enterprise webpage	AWS Sustainable Infrastructure	Materials & lifecycle focus	https://aws.amazon.com/sustaina	AWS Sustainability. (n.d.-b).

			bility/data-centers/	

Source: Author's elaboration based on data corpus in Phase 1 (Section 3.1.1).

Microsoft Azure				
Source Type	***Source Name***	***Description***	***Link***	***Citation***
Enterprise webpage	Microsoft for Sustainability	Overview of Azure sustainability approach	https://www.microsoft.com/en-us/sustainability/cloud	Microsoft. (n.d.-a).
Enterprise webpage	Improve your energy and carbon efficiency with Azure sustainability guidance	Explains PUE, WUE, and other measurement logics	https://azure.microsoft.com/en-us/blog/improve-your-energy-and-carbon-efficiency-with-azure-sustainability-guidance/	Taylor, A. (2022).
Enterprise webpage	How Microsoft measures datacenter water and energy use to improve Azure Cloud sustainability	Methodological explanation of metrics	https://azure.microsoft.com/en-us/blog/how-microsoft-measures-datacenter-water-and-energy-use-to-improve-azure-cloud-sustainability/	Walsh-Elwell, N. (2022).
Enterprise webpage	Achieving 100 percent renewable energy with 24/7 monitoring in Microsoft Sweden	Hourly renewable matching concept	https://azure.microsoft.com/en-us/blog/achieving-100-percent-renewable-energy-with-247-monitoring-in-microsoft-sweden/	Walsh-Elwell, N. (2020).
Enterprise webpage	Microsoft Datacenters: Powering sustainable transformation	High-level infrastructure sustainability claims	https://datacenters.microsoft.com/globe/powering-sustainable-transformation/	Microsoft. (n.d.-b).
Enterprise webpage	Microsoft announces investment to accelerate Sweden's digital transformation and plants to open its sustainable datacenter region in 2021	Region-specific PUE/WUE example	https://news.microsoft.com/europe/2020/11/24/microsoft-announces-investments-to-accelerate-swedens-digital-transformation-and-plans-to-open-its-sustainable-datacenter-region-in-2021/	Microsoft Reporter. (2020).
Enterprise webpage	Advancing embodied carbon measurements at scale for Microsoft Azure hardware	Scope 3 and lifecycle measurement	https://techcommunity.microsoft.com/blog/azureinfrastructureblog/advancing-embodied-carbon-measurement-at-scale-for-microsoft-azure-hardware/4485784	Leoaspauza. (2026)

Source: Author's elaboration based on data corpus in Phase 1 (Section 3.1.1).

Google Cloud Platform				
Source Type	***Source Name***	***Description***	***Link***	***Citation***

Enterprise webpage	Growing the internet while reducing energy consumption	Granular PUE metrics and efficiency calculations	https://datacenters.google/efficiency/	Google Data Centers. (n.d.-b).
Enterprise webpage	Carbon free energy for Google Cloud regions	Regional-level CFE data	https://cloud.google.com/sustainability/region-carbon	Google Cloud. (n.d.).
Enterprise webpage	How carbon-free is your cloud? New data lets you know	Technical blog	https://cloud.google.com/blog/topics/sustainability/sharing-carbon-free-energy-percentage-for-google-cloud-regions	Google Cloud Blog. (2021).
Enterprise webpage	Innovating across our operations and supply chains	Overall environmental reporting	https://sustainability.google/operations/	Google Sustainability. (n.d.)
Enterprise webpage	Our approach to carbon-aware data centers: Central data center fleet management.	Workload shifting & carbon-aware logic	https://cloud.google.com/blog/topics/sustainability/googles-approach-to-carbon-aware-data-center	Arshi & Miller. (2025).
Enterprise webpage	Driving sustainable innovation of society	Operating sustainability: cooling, water, infrastructure	https://datacenters.google/operating-sustainably/	Google Data Centers. (n.d.-a).

Source: Author's elaboration based on data corpus in Phase 1 (Section 3.1.1).

Appendix C. Semi-Structured Informal Consultations: Guide & Consent Form

Semi-Structured Informal Consultations Guide

Aura Cloud Deployment & Sustainability

Estimated duration: 30 minutes

Format: Semi-structured

Participants: [...]

Questions:

Section A: Role & Context

(1) Could you briefly describe your role in relation to Aura development and cloud deployment decision?
(2) To what extent are you involved in decisions about region or provider selection?

Section B: Current Deployment Logic

(1) What are the main criteria currently used when selecting cloud regions or providers for Aura workloads?
(2) Are there internal guidelines or formal decisions frameworks guiding region selection?

Section C: Flexibility & Contractions

(1) What types of Aura workloads (e.g., development, testing, canary deployments) are most flexible in terms of region placement?
(2) Which workloads are least flexible and why?
(3) Are there contractual or financial constraints (e.g., long-term commitments, marketplace agreements) that limit region or provider changes?
(4) How would changing deployment regions potentially affect engineering productivity or system performance?

Section D: Sustainability Awareness & Decision Relevance

(1) To what extent is environmental performance currently considered in cloud deployment decisions?
(2) If region-level environmental information (e.g., carbon intensity, PUE, renewable resourcing) were clearer and comparable would it realistically influence deployment decisions?
(3) What information would be necessary to make sustainability considerations actionable?

Section E: Hypothetical Scenarios

(1) If two regions were technically equivalent in performance and cost, but one had demonstrably lower carbon intensity, would switching be feasible?
(2) Under what conditions would sustainability become a decisive factor in deployment decisions?
(3) What trade-offs would be unacceptable?

Section F: Closing

(1) From your perspective, what are the main barriers to making cloud deployment more environmentally sustainable?

Consent Form

Thesis title:

The Weight of the Cloud: Who Owns the Rain?

Cloud Sustainability Beyond Metrics: Provider Transparency and Customer-Side Feasibility in SaaS Deployment Decisions

Researcher:

Aoife Michelle Ní Uanacháin, IIIEE Lund University

Purpose of the Study

This study examined how mid-sixes SaaS companies can make sustainability-informed cloud deployment decisions under conditions of limited infrastructural control and incomplete information.

The consultations aim to understand organisational constraints and feasibility considerations related to Aura cloud deployment.

What Participation Involves

You have been invited because of your role in the Aura development, cloud deployment, or related operational decision-making processes. Participation involves a semi-structured consultation lasting approximately 30 minutes. The consultation will focus on deployment logic, constrained, and sustainability considerations. No personal performance evaluation is involved.

Voluntary Participation

Participation is voluntary. You may decline to answer any questions and may withdraw at any time without consequences.

Data Handling & Confidentiality

No sensitive personal data will be collected. Data will be anonymised in the thesis and stored securely in accordance with Lund University research standards

Use of Results

Findings will be used for academic purposes and may inform internal learning at Neo4j. No individual names or identifiable information will be disclosed.

If you are happy to proceed with a consultation and participate in this study, please complete the following consent form:

Consent to Participate in Research

Participants Name:

Signature:

Date:

Researcher Signature:

Date:

Appendix D. Coding Framework

This coding system is built on the 2025 Best Practice Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency. All references to the Guidelines are from this unique source.

The structure is two-tiered: (Level 1) the coding dimensions T, C, D; (Level 2) the standards-linked subcodes.

Each subcode is scored 0-2, then aggregated cautiously into dimension totals, not as a claim of objective performance scores, but as disclosure adequacy scores.

Table 10 breaks down the coding framework with code descriptions and rationale.

Table 10. Coding Framework, Detailed Breakdown of Individual Codes

Transparency (T) *Transparency is the extent to which Amazon, Microsoft, and Google disclose environmental performance information with sufficient availability, granularity, methodological standardisation, and boundary clarity to permit informed interpretation*		
Code	***Code Description & Scoring***	***Code Rationale***
T1	***T1 - KPI availability*** Does the provider disclose the core KPI families relevant to data centre environmental assessments? Check for: - PUE - REF / Renewable electricity share - CUE / Carbon/emissions intensity - WUE - ERF / heat reuse where relevant - Utilisation-related indicators Scoring: 0 = absent 1 = partial / selective disclosure 2 = broad disclosure across core KPI families	The 2025 Guidelines states that written reporting should use standardised KPIs and explicitly like: PUE, pPUE, REF, ITEESV, ITEUSV, ERF, CER, CUE, and WUE (p.54).
T2	***T2- Standardised KPI disclosure*** Are those metrics reported using recognised standards rather than provider-defined labels? Check whether the disclosure references: - EN 50600-4-x - ISO/IEC 30134 series - Equivalent recognised KPI methods Scoring: 0 = proprietary / undefined metrics 1 = partially aligned or inferable 2 = explicitly standardised	The Guidelines require that if these KPIs are reported, they should follow the standardised ISO/IEC 30134 series (or EN 50600-4-x equivalent) methods (p.54). For instance, reporting PUE under EN 50600-4-2 or ISO/IEC 30134-2 (p.54); REF under EN 50600-4-3 or ISO/IEC 30134-3 (p.12); and WUE under ISO/IEC 30134-9 and EN 50600-4-9 (p.50).

T3	***T3- Granularity*** At what unit is the information disclosed? Code the highest consistent level. Scoring: 0 = provider / global only 1 = partial region / service / location disclosure 2 = consistent region-level or site / data-centre-level disclosure	The Guidelines repeatedly link better monitoring to increasing visibility and granularity, including infrastructure granularity (p.51,54), and cabinet- & device-level metering (p.51-52).
T4	***T4 - Methodological clarity*** Does the provider explain: - What the metric means - How it is calculated - Used denominator/numerator - What averages, spot values, or ranges are shown - Whether exclusion or estimation methods apply Scoring: 0 = unclear 1 = partially explained 2 = clearly documents	This is directly supported by the requirement that reporting references the category being reported and uses the required method of data collection and calculation (p.54).
T5	***T5 - Temporal resolution*** How current and decision-timely is the disclosure? Scoring: 0 = one-off/static statement 1 = annual reporting only 2 = more frequent, dynamic, dashboard-based, or regularly updated disclosure	The Guidelines distinguish periodic written reporting (p.54), automated daily readings, automated hourly refining, and integrated reporting consoles, which provides a solid standard-based basis for temporal coding.
T6	***T6 - Boundary and responsibility clarity*** Is it clear what part of the stack the provider's disclosure covers? Using the 2025 "areas of responsibility" (p.9) logic: - Physical building - Mechanical / electrical plant - Data floor - Cabinets - IT equipment - Operating Systems / virtualisation - Software - Business practices Scoring: 0 = boundary unclear 1 = partially inferable 2 = clear scope/responsibility boundary	This gives a formal basis for saying why CSP disclosures are hard to interpret when they mix facility, service, and software layers.
Comparability (C) *Comparability is the extent to which the disclosed data can support like-for-like comparison across CSPs and across regions*		
Code	***Code Description & Scoring***	***Code Rationale***
C1	***C1 - Metric equivalence*** Are providers disclosing the same kind of metric? Scoring: 0 = different/non-equivalent indicators 1 = broadly similar but not fully equivalent	This avoids false comparison between, for example, generic renewable claims and standardised REF.

	2 = same metric family across providers	
C2	***C2 - Method equivalence*** Are the disclosed metrics computed under the same recognised method? Scoring: 0 = no shared method 1 = partially aligned 2 = clearly aligned to same EN/ISO KPI	The Guidelines insist standardised KPI methods be used if reported (p.54).
C3	***C3 - Boundary equivalence*** Do the providers disclose comparable system boundaries? Code whether the disclosure clearly pertains to: - Whole provider estate - A cloud region - A facility - Operational electricity only - Broader life cycle or corporate emissions Scoring: 0 = incomparable boundaries 1 = partially comparable 2 = clearly comparable	This code explicitly uses the EU responsibility categories in the conceptual anchor.
C4	***C4 - Granularity equivalence*** Do the providers report the same level of detail? Scoring: 0 = one provider global-only, another region/site-level 1 = mixed partial equivalence 2 = same granularity level across providers	
C5	***C5 - Context comparability*** Do the disclosures provide enough contextual information to avoid misleading comparisons? Check for: - Climate/location effects - Cooling configuration differences Water source differences - Energy source definitions - Heat resume context Scoring: 0 = no context 1 = limited context 2 = sufficient contextualisation	This code is supported by the Guidelines' warning that water consumption cannot be directly compared unless the energy intensity of water source is understood, and that comparison waste consumption between buildings is not useful without context (p.50).
C6	***C6 - Aggregation bias*** Do the providers aggregate in ways that mask relevant variation? Scoring: 0 = high aggregation obscurers comparison 1 = some disaggregation but important masking remains 2 = low aggregation/meaningful separation	The Guidelines emphasises increasing metering granularity (p.47), infrastructure visibility, and grouped-by-location IT reporting (p.52).
Decision-relevance (D) *Decision relevance is the extent to which disclosed CSP information can support a mid-sized SaaS firm's deployment choice, especially region selection*		
Code	***Code Description & Scoring***	***Code Rationale***

D1	***D1 - Region-actionability*** Can a customer identify whether Region A is environmentally preferable to Region B? Scoring: 0 = not possible 1 = weak/indirect inference only 2 = yes, directly supported	
D2	***D2 - Workload linkage*** Is this information linked to the actual decision unit customers control? Check whether the disclosure related to: - Cloud regions - Workload placement - Service choice - Architectural choice - Time-varying operational decisions Scoring: 0 = no link 1 = indirect link 2 = clear link	The Guidelines explicitly connects monitoring and dashboard to the energy used to deliver IT services and to workload-linked reporting.
D3	***D3 - Interpretability for non-specialists*** Can a sustainability manager or engineer lead a mid-sized SaaS firm understand and use the information without reconstructing hidden assumptions? Scoring: 0 = inaccessible/highly technical 1 = partially interpretable 2 = clear and user-usable	
D4	***D4 - Business-facing decision support*** Does the provider offer dashboards, calculators, or other interfaces that translate infrastructure data into a customer decision context? Scoring: 0 = none 1 = partial 2 = yes	The Guidelines explicitly mention integrated reporting consoles and relevant dashboards that reflect and help reduce the energy needed to deliver IT services (p.54).
D5	***D5 - Defensibility of inference*** Could a researcher reasonably defend a customer recommendation based on the disclosed data lone? Scoring: 0 = not defensible 1 = only with strong caveats 2 = defensible with limited caveats	

Source: Author's elaboration, based on Acton et al. (2025).

Appendix E. Declaration of AI Form

Declaration of Use of Generative AI (GAI) Tools

1) Permitted Uses in This Course

Students may use GAI tools for:

- Brainstorming and idea generation
- Planning and outlining
- Language support (grammar, clarity, tone)
- Translation of non-personal, non-sensitive content
- Drafting outlines (not full assignments)

Requirements:

- You must critically review, fact-check, and edit all AI outputs.
- Do not submit unedited AI-generated work.
- Do not input personal, sensitive, or confidential data into GAI tools.

2) Student Declaration (to be completed)

Project/Assignment title:

"The Weight of the Cloud: Who Owns the Rain?
Cloud Sustainability Beyond Metrics: Provider Transparency and Customer-Side Feasibility in SaaS Deployment Decisions.
Neo4j as a Case Study"

Did you use GAI tools for this work? **Yes**
If Yes, complete the items below.

a) Tools used

ChatGPT & Capgemini.

b) What you used them for (check all that apply)

- ☑ Brainstorming / idea generation
- ☑ Planning / structuring
- ☑ Language editing (grammar/clarity)
- ☑ Outline drafting

c) How you used the tools (be specific)

Describe prompts, steps taken, and how you revised the outputs:

I used ChatGPT for the initial brainstorming to generate 6 different topic angles that merged academic soundness and organisational relevance for the case study; I selected 2, discarded the rest, and rewrote them in my own words to propose to the company chosen for the case study.
I used ChatGPT to provide assistance in laying out the skeleton for the thesis to accurately balance the academic and organisation focuses. This was used as a baseline but later modified as the research progressed.
I used mainly Capgemini (provided by the case company Neo4j) for additional language polishing, clarity, and occasionally shortening long paragraphs.

d) Your own contributions

What parts are entirely your original work?

All of the work is my original work, aside from the brainstorming assistance and the language polishing. All assistance sought out from by AI was reworded and critically reviewed before incorporating it into the thesis.

e) Verification and editing

How did you verify accuracy and edit AI output?

I always proofread the provided assistance and rewrote everything suggested. I did additional research on top of brainstorming suggestions to validate reliability of suggestions.
All prompts were ultimately used primarily for guidance.

Student signature: *Aoife Michelle Ní Uanacháin*

Date: May 17, 2026