Environmental Strategies of the Cloud Computing Industry

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Gabe Silverstein

Spring, 2025

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Bryn E. Seabrook, Department of Engineering and Society

Environmental Strategies of the Cloud Computing Industry

Introduction: A Balancing Act

Emerging technologies like artificial intelligence (AI), machine learning (ML), and the Internet of Things (IoT) are transforming industries at an unprecedented pace—but behind every chatbot response, facial recognition scan, and recommendation algorithm lies an invisible powerhouse: the cloud. Cloud computing provides the fundamental building blocks needed to process large swaths of data in a scalable and flexible manner. Despite how important cloud infrastructure is to the current technological landscape, its usage comes with socioeconomic and environmental implications. For instance, a single cloud data center consumes the same amount of electricity as fifty thousand homes in just one year of operation. Additionally, yearly carbon emissions from data centers are so large that they eclipse the entire airline industry (Monserrate, 2022). These concerning statistics, among others, leave cloud service providers in a tough spot as they struggle to balance consumer demand and environmental ramifications amidst a lack of federal regulation (Marwah et al., 2010).

To shed some light on that sociotechnical issue, this paper explores how and why major cloud providers adopt their environmental strategies, specifically through the lens of the Social Construction of Technology (SCOT) theory. Utilitarian and deontological ethics are also used to investigate why certain policy decisions are considered standard. By understanding the factors that play into the sustainability policies of cloud providers, it becomes easier to pinpoint where changes need to be made to strike a balance between technological innovation and preserving the environment.

Methods: Analyzing Industry Leaders

This paper primarily utilizes document, policy, and ethical analyses to examine the main aspects that major cloud service providers consider when developing their environmental strategies. In particular, the research focuses on the three largest and most influential cloud providers: Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). These providers were selected because they collectively dominate the global cloud market, set industry standards, and have the resources and visibility to significantly influence sustainability practices across the technology sector. Keywords like cloud computing, sustainability, and environmental policy were used when researching those providers' strategies. Document analysis is used on environmental reports and policy statements to determine which metrics the companies deem to be important. Those primary source documents help illustrate the interaction between corporate social responsibility and consumer expectations that this paper tries to explore. Additionally, current data center statutes and proposed bills undergo a policy analysis to investigate how the lack of federal regulation influences the strategies that cloud industry leaders adopt. The document and policy analyses are supplemented by an ethical analysis of public statements from cloud executives. Evaluating those transcripts through ethical frameworks illustrates the extent to which ethics play a role in the highest levels of cloud management. Each research method produces different findings, but they all seek to interpret commonalities or gaps across the different policies that cloud providers support. Those results are grouped into how and why sections that summarize cloud providers' policies. The how focuses on the implementation details of each provider's strategy, while the why takes a broader view and examines the underlying motivations present in the industry. The insights from each section are then used to outline areas where cloud providers can improve their policies and environmental impact.

Background: Understanding The Cloud

To analyze the sustainability decisions in cloud policies, it is first necessary to understand the cloud itself. At its core, the cloud is one large distributed system. In other words, it is a network of geographically dispersed computers connected over the internet that work together to complete tasks. This grid of computers provides extensive compute resources almost anywhere there is internet connection, driving the appeal of cloud computing technologies. Cloud service providers provide ubiquitous and convenient access to computing power by building their systems around five core tenets. The National Institute of Standards and Technology (NIST) defines those principles as on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service (Mell & Grance, 2011). Those standards result in consumers being able to scale resources whenever they want, and through nearly any device they want (computer, phone, tablet, etc.), without having to know where data is stored or requiring human interaction with the service provider. On top of that flexibility, consumers only have to pay for their recorded resource usage.

The versatility that cloud computing provides has led to it becoming the de facto place to process large datasets. Coupled with the surge in demand for AI, an inherently data-intensive technology, the cloud computing market has more than doubled in the past few years. The cloud is growing at such a fast rate that Goldman Sachs predicts the market for cloud computing technologies to compound at a staggering annual growth rate of 22% between 2024 and 2030 (*Cloud Revenues Poised to Reach $2 Trillion by 2030 amid AI Rollout*, 2024). This growth in the cloud computing industry has been paralleled by the proliferation of data centers to meet demand (Yan et al., 2024). From a pure infrastructure standpoint, the United States is not currently equipped to handle such a rapid increase in energy demands. Specifically, the energy providers that power these data centers are already struggling to meet current demand, and the construction of more data centers will only exacerbate this issue (Li & Zhang, 2024).

Besides massive energy consumption and carbon emissions, the growth of cloud computing brings other significant environmental and socioeconomic challenges. One of the most noteworthy sociotechnical problems is the excessive usage of water. In 2021 alone, Google's data centers consumed 4.3 billion gallons of water. To put that number in context, 4.3 billion gallons of water can irrigate 29 golf courses yearly or 17 acres of lawn daily (Hölzle, 2022). This reliance on water stems from the need to irrigate the multiple server rooms in data centers and prevent the machines from overheating. The need for water is so significant that it impacts the surrounding communities of data centers. For example, residents in Bluffdale, Utah experience power outages and water shortages due to their proximity to the Utah Data Center (Hogan, 2015). This short description of a major socioeconomic and environmental issue is only one of the many challenges facing the cloud computing industry. Understanding the context surrounding cloud infrastructure makes it easier to recognize why it is important to investigate the factors that go into developing the environmental policies of major cloud providers.

STS Framework: Socially Constructing The Cloud

The Social Construction of Technology (SCOT) is the main theory used throughout the paper to analyze how and why major cloud providers adopt their environmental strategies. SCOT was chosen as the STS framework to apply to this research question because of how it explores the social context surrounding technology. Championed by STS researchers Wiebe Bijker and

Trevor Pinch, SCOT theory asserts that the development of technology is shaped by human interaction and influenced by social, cultural, economic, and political factors (2012). That type of viewpoint aligns with the paper's research question as it explores sustainability in cloud computing—a topic with significant sociotechnical context. SCOT also provides tools like interpretive flexibility and closure/stabilization, which help analyze the factors contributing to the current state of a particular technology. Interpretive flexibility refers to the idea that different social groups can design and interpret technology in multiple ways, while closure/stabilization describes how the different groups reach a consensus on a design.

While current literature lacks an analysis of cloud environmental policies through SCOT, it does contain applications of SCOT to the broader infrastructure surrounding cloud computing. For example, Professor Rhinesmith from the University of Illinois Urbana-Champaign wrote an article describing how different social groups influenced a nonprofit's implementation of a cloud computing project in St. Louis, Illinois (2015). He combined SCOT's principles of interpretive flexibility and stabilization with Star's ethnography of infrastructure (1999) to create a successful analysis that highlighted the tensions between external stakeholders and internal organizational needs. The paper also noted that the social and technical aspects of the project were intertwined to form one complex sociotechnical issue. That point supports the concepts developed by Bijker and Pinch, illustrating that "the social and technical are mutually constitutive and cannot be analyzed separately" (2015, p. 3). Similar to how Rhinesmith used interpretive flexibility and stabilization to describe cloud computing infrastructure, this research on cloud sustainability strategies seeks to use those same concepts to draw attention to the sociotechnical factors that drive policy adoption.

Results & Discussion: Considering Cloud Sustainability Strategies

Cloud service providers adopt their environmental strategies based on a variety of political, social, economic, and ethical considerations. SCOT helps characterize the common factors in those sustainability efforts and explain why all their policy implementations incorporate renewable energy, water stewardship, carbon offsetting, and energy efficiency initiatives. It also illustrates how different social groups—like consumers, corporations, and the government—shape the policies that are adopted. Utilitarianism and deontology viewpoints supplement the SCOT analysis and explore how moral obligations and consequences further shape cloud service providers' environmental strategies. By examining both the how and why behind the adoption of these sustainability policies, it becomes easier to identify areas for improvement, ensuring a balance between technological innovation and environmental preservation.

The How

From restarting the nuclear power plant at Three Mile Island (Mandler, 2024) to using hydrotreated vegetable oil to power backup generators (Venkatesan & Karibandi, 2024), cloud providers like Microsoft, Amazon, and Google are pursuing many innovative solutions in their environmental strategies. Despite differing implementation details, each company's policy addresses the same core sustainability challenges in renewable energy, water stewardship, carbon emissions, and energy efficiency. This convergence around similar objectives can be described by the SCOT principle of closure and stabilization. Although the actual technology solutions have not coalesced due to how rapidly cloud computing is growing, cloud providers realized that the aforementioned obstacles kept appearing as the industry expanded. Certain iterations of cloud infrastructure gained more traction over time, and best practices emerged for implementing a cloud sustainability strategy.

One of the core tenets of any cloud environmental policy is investment in renewable energy. This push for cleaner energy is the byproduct of different social groups—including consumers, corporations, and the government—agreeing on a best practice for being more sustainable. To meet that goal, cloud service providers have collectively focused on power purchase agreements (PPAs) for wind and solar energy. A PPA is an agreement with an energy provider to invest in a renewable energy project and buy the energy output. The main catch is that the renewable energy is not directly used to power data centers. A PPA "merely ensure[s that] an equivalent amount of a customer's agreed energy demand is being generated by renewable sources" (Swinhoe, 2023, para. 12). This loophole introduces an ethical dilemma under a deontological framework. Deontology would see a moral obligation to use the renewable energy to power the data centers, as one of the main purposes of the investment was to reduce the company's reliance on fossil fuels. In contrast, a utilitarian standpoint might view the status quo as acceptable since renewable energy is inherently intermittent, and improving sustainability at the cost of availability should not outweigh the cloud computing needs of consumers worldwide.

Another focal point of these policies is water stewardship, which is the process of using water equitably. While each company has stabilized around this principle, SCOT's interpretative flexibility also plays a role here, as water stewardship means something different to each of them. For Amazon (2024), it entails working toward being water positive—i.e., replenishing more water than is used—by 2030 and investing in clean water initiatives in India and Indonesia.

To Microsoft (2024), it means not only being water positive by 2030, but also improving their water use efficiency by 40% from their 2022 baseline. Google (2024) takes a different approach and wants to replenish 120% of the annual freshwater volume they consume by 2030, in addition to their 74 other water stewardship projects. Despite their differences and the lack of closure on a single technological implementation, each plan has common themes and sets goals for 2030 while investing in water replenishment initiatives.

Reduction in carbon emissions is probably the most well-known initiative by consumers because of how often cloud providers tout it. Similar to how there are varying interpretations of water stewardship, each cloud provider takes their own approach despite agreeing on the overall goal. Although AWS (2024) holds the largest market share, its emissions target is the least aggressive, trying to be carbon-neutral by 2040. Conversely, Microsoft's Azure (2024) takes the most ambitious approach and wants to be carbon-negative by 2030. GCP (2024) sits somewhere in the middle, intending to be carbon-neutral by 2030. In this instance, SCOT illustrates how different stakeholders and company cultures influence the commitments that each company undertakes. While these corporate pledges are commendable, they are not enforceable or necessarily feasible with how rapidly cloud infrastructure is growing (Monserrate, 2022). It is particularly important to keep track of how these plans evolve and stabilize as conditions change and it gets closer to 2030.

Energy efficiency improvements are also a key component of a cloud environmental policy. Most of the large cloud providers have decided to pursue machine learning (ML) optimizations for their power usage effectiveness (Oberhaus, 2019). It is interesting that they have converged around this idea because ML requires significant resources to train and run. Using an energy-intensive technology to solve an energy efficiency problem can seem counterintuitive and raises questions about why Google, Microsoft, and Amazon have socially constructed that approach as one worth investing in. Their energy efficiency strategies also pose ethical dilemmas regarding access and equity. Lucivero (2019) highlights that with data centers moving to colder climates—and even under the ocean—to save energy on cooling, who ensures that economic gaps between warm countries and cold countries are avoided? Furthermore, how do the benefits and consequences that local data center communities experience reconcile with utilitarian and deontological ethics? Is it okay to forego fairness if the resulting energy savings help the environment?

The Why

Beneath the surface-level implementations of these cloud sustainability strategies lies a complex combination of political, social, economic, and ethical factors. The underlying considerations that cloud providers have to grapple with are influenced by what SCOT calls relevant social groups. Governments, customers, investors, and more all play a role in shaping the policies that cloud providers adopt. These stakeholders influence not only the adoption of specific technologies but also the broader narratives surrounding sustainability within the cloud computing industry. While each group might assign a different meaning to what an environmental strategy entails, together, they turn sustainability from a technical challenge into a socially constructed necessity.

Government regulations, or the lack thereof, are one of the driving factors behind why cloud providers adopt specific objectives and disclose certain information in their policies. In the U.S., cloud computing is mostly unregulated (Monserrate, 2022). This lack of oversight requires providers to navigate the political landscape mostly on their own and decide to what extent they wish to be sustainable. It also results in cloud providers underreporting their emissions by around seven times the actual amount due to their differing interpretations of reporting standards (New York State Sustainable Data Centers Act, 2024). That status quo is being challenged by local governments in states like Virginia and New York, which are trying to create reporting mandates and more detailed sustainability standards. The policymakers, environmental advocacy groups, and industry leaders who propose that legislation are changing the way that cloud service providers enact their environmental strategies by pushing them to be more transparent and proactive. EU policies are also driving this shift with the Energy Efficiency Directive (EED) and Corporate Sustainability Reporting Directive (CSRD). Despite their slow rollout, the EED and CSRD are requiring cloud providers to become more sustainable and inspiring other governments to strengthen their regulations (Wong, 2024). The combination of interpretations from these different governments and companies has begun to find closure around transparency and effective resource utilization. That stabilization creates politically motivating factors for cloud providers to consider when creating their environmental strategies.

Social influences are also critical in the development of these policies as cloud providers balance environmental expectations with consumer demands. Consumers want near 24/7 availability from their cloud infrastructure, but also want it to be environmentally friendly. That combination is difficult to achieve as quality of service is inversely related to resource consumption (Panwar et al., 2022). The "green" technologies that cloud providers adopt in their strategies are a direct result of these disparate demands. To preserve brand reputation and competitive advantages, strategies must be sustainable enough to attract ESG (Environmental, Social, and Governance) investors yet realistic enough to maintain service agreements. The policies must also address a variety of components, from waste to water to emissions, to indicate to governments and consumers how comprehensive their plans are. This breadth makes it challenging to enact meaningful change in any one area. SCOT helps unpack the balancing act between those competing pressures as differing viewpoints from media narratives, corporate social responsibility initiatives, and consumer activism shape what objectives are considered sustainability best practices. Those socially constructed expectations around sustainability become self-reinforcing as companies integrate these objectives into their branding and corporate identity, further driving the industry toward greener practices.

In addition to the social aspects, cloud environmental strategies are impacted by financial considerations. For instance, cloud providers like Google, Amazon, and Microsoft recognize that improving energy efficiency and reducing waste directly contribute to cost savings, making sustainability not just an ethical choice but a financially sound one (Cao et al., 2023). Furthermore, state governments frequently offer tax incentives, subsidies, and grants to companies that are investing in green technologies, reinforcing the economic benefits of sustainability (Wong, 2024). Those combined factors result in significant investment in more efficient cooling systems and renewable energy to mitigate long-term financial risks associated with climate change. It also incentivizes cloud providers to adopt sustainable practices preemptively to avoid future financial penalties. SCOT reveals how economic pressures from different social groups, like internal stakeholders, customers, and regulatory bodies, influence environmental decisions. Cloud providers are not only protecting their bottom line for current stakeholders, but are also looking for ways to ensure success down the road by responding to industry trends. The push for energy-efficient technologies by those groups illustrates that sustainability is as much about financial judgment as it is about social responsibility.

Besides the political and socioeconomic factors, corporate social responsibility and ethical obligations play a significant role in why cloud sustainability policies are adopted in their current fashion. Major cloud providers know that it is no secret how pronounced their environmental impact is and believe that they, therefore, have a moral and fiduciary duty to mitigate any harm that comes as a result of their operations (Nakagawa & Smith, 2023). That viewpoint aligns with a deontological perspective that argues that corporations have a fundamental duty to act responsibly, independent of other motivations. It also supports a utilitarian standpoint as having a corporate culture that prioritizes ethical behavior benefits the planet and all who depend on it. SCOT's closure and stabilization provide insight into how those ethical expectations evolve within communities and corporate culture, shaping the adoption of greener technologies. As discussions over climate change and environmental stewardship become more prevalent, ethical decision-making becomes embedded in industry practices as a way to preserve customer trust.

Limitations and Future Research

While this paper identifies how different social groups and their interpretations shape how and why major cloud providers develop their environmental strategies, certain limitations remain. Namely, this analysis focused solely on the policies of Amazon, Microsoft, and Google. Smaller non-hyperscale providers might not experience the same political, social, economic, and ethical pressures as their larger counterparts. Therefore, they might also prioritize different objectives in their environmental policies that are not covered in this paper. Additional research is needed to analyze how the size of a cloud provider affects its environmental policies and evaluate the resulting impact on what policies are considered standard. This paper is also limited to a primarily SCOT analysis of these sustainability strategies. While some ethical frameworks are incorporated, the focus is on investigating how those policies are socially constructed and the varying factors and interpretations that play into their creation and implementation. Future research from STS scholars is needed to provide further context on how and why cloud providers adopt their environmental policies and perhaps analyze those decisions through different frameworks like Star's ethnography of infrastructure (1999) and Law's actor-network theory (1992).

Conclusion: Learning From The Status Quo

As cloud service providers navigate the intersection of sustainability and technological advancement, the influence of political, social, economic, and ethical factors remains central to their decision-making. SCOT highlights how various stakeholders—including consumers, corporations, and governments—reach consensus on and drive the adoption of green technologies and policies. Deontological and utilitarian perspectives further shape these strategies, offering different justifications for corporate sustainability efforts. Despite these insights, current sustainability policies still have room for improvement. Greater transparency in energy reporting, stricter enforcement of renewable energy commitments, and enhanced water stewardship strategies could make these efforts more impactful. Additionally, government incentives could be more effectively structured to ensure that carbon offset programs lead to tangible environmental benefits rather than serving as corporate loopholes. With the cloud computing industry continuing to expand, future studies must focus on the long-term effectiveness of these strategies, assessing their impact on energy consumption, carbon emissions, and corporate accountability. The coming years will determine whether cloud providers can lead the way in sustainable innovation or whether further interventions will be necessary to drive meaningful change.

References

Amazon. (2024). *Amazon Sustainability Report 2023 AWS Summary*. https://sustainability.aboutamazon.com/2023-amazon-sustainability-report-aws-summary.pdf

Bijker, W. E., Hughes, T. P., & Pinch, T. J. (2012). *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. The MIT Press.

Cao, Z., Zhou, X., Wu, X., Zhu, Z., Liu, T., Neng, J., & Wen, Y. (2023). Data Center Sustainability: Revisits and Outlooks. *IEEE Transactions on Sustainable Computing*, *9*(3), 236–248. https://doi.org/10.1109/tsusc.2023.3281583

Cloud revenues poised to reach $2 trillion by 2030 amid AI rollout. (2024, September 4). Goldman Sachs. https://www.goldmansachs.com/insights/articles/cloud-revenues-poised-to-reach-2-trillion-by-2030-amid-ai-rollout

Google. (2024). *Google 2024 Environmental Report*. https://www.gstatic.com/gumdrop/sustainability/google-2024-environmental-report.pdf

Hogan, M. (2015). Data flows and water woes: The Utah Data Center. *Big Data & Society*, *2*(2). https://doi.org/10.1177/2053951715592429

Hölzle, U. (2022, November 21). Our commitment to climate-conscious data center cooling. Google. https://blog.google/outreach-initiatives/sustainability/our-commitment-to-climate-conscious-data-center-cooling/

Law, J. (1992). Notes on the theory of the actor-network: Ordering, strategy, and heterogeneity. *Systems Practice*, *5*(4), 379–393. https://doi.org/10.1007/BF01059830

Li, X., & Zhang, S. (2024). Management Mode and Path of Digital Transformation of Power Grid Enterprises Based on Artificial Intelligence Algorithm. *International Journal of Thermofluids*, *21*. https://doi.org/10.1016/j.ijft.2023.100552

Lucivero, F. (2019). Big Data, Big Waste? A Reflection on the Environmental Sustainability of Big Data Initiatives. *Science and Engineering Ethics, 26,* 1009–1030. https://doi.org/10.1007/s11948-019-00171-7

Mandler, C. (2024, September 20). *Three Mile Island nuclear plant will reopen to power Microsoft data centers*. NPR. https://www.npr.org/2024/09/20/nx-s1-5120581/three-mile-island-nuclear-power-plant-microsoft-ai

Marwah, M., Maciel, P., Shah, A., Sharma, R., Christian, T., Almeida, V., Araújo, C., Souza, E., Callou, G., Silva, B., Galdino, S., & Pires, J. (2010). Quantifying the sustainability impact of data center availability. *ACM SIGMETRICS Performance Evaluation Review*, *37*(4), 64–68. https://doi.org/10.1145/1773394.1773405

Mell, P., & Grance, T. (2011). *The NIST Definition of Cloud Computing: Recommendations of the National Institute of Standards and Technology*. NIST. https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

Microsoft. (2024). *How can we advance sustainability? 2024 Environmental Sustainability Report*. https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/documents/presentations/CSR/Microsoft-2024-Environmental-Sustainability-Report.pdf

Monserrate, S. G. (2022). The Cloud Is Material: On the Environmental Impacts of Computation and Data Storage. *MIT Case Studies in Social and Ethical Responsibilities of Computing, Winter 2022*. https://doi.org/10.21428/2c646de5.031d4553

Nakagawa, M., & Smith, B. (2023, May 10). On the road to 2030: Our 2022 Environmental Sustainability Report. *Microsoft*. https://blogs.microsoft.com/on-the-issues/2023/05/10/2022-environmental-sustainability-report/

Oberhaus, D. (2019, December 10). *Amazon, Google, Microsoft: Here's Who Has the Greenest Cloud*. WIRED. https://www.wired.com/story/amazon-google-microsoft-green-clouds-and-hyperscale-data-centers/

Panwar, S. S., Rauthan, M. M. S., & Barthwal, V. (2022). A systematic review on effective energy utilization management strategies in cloud data centers. *Journal of Cloud Computing*, 11(1). https://doi.org/10.1186/s13677-022-00368-5

Rhinesmith, C. (2015). The social shaping of cloud computing: An ethnography of infrastructure in east St. Louis, Illinois. *Proceedings of the American Society for Information Science and Technology*, *51*(1), 1–10. https://doi.org/10.1002/meet.2014.14505101060

Star, S. L. (1999). The Ethnography of Infrastructure. *American Behavioral Scientist*, *43*(3), 377–391. https://journals.sagepub.com/doi/10.1177/00027649921955326

Venkatesan, V., & Karibandi, A. (2024, September 16). Embracing Modernization with a Sustainability Focus. *Amazon Web Services*. https://aws.amazon.com/blogs/migration-and-modernization/embracing-modernization-with-a-sustainability-focus/

Wong, W. (2024, November 6). *Data Center Regulation Trends to Watch in 2025*. Data Center Knowledge. https://www.datacenterknowledge.com/regulations/data-center-regulation-trends-to-watch-in-2025

Yan, D., Chow, M.-Y., & Chen, Y. (2024). Low-Carbon Operation of Data Centers With Joint Workload Sharing and Carbon Allowance Trading. *IEEE Transactions on Cloud Computing*, *12*(2), 750–761. https://doi.org/10.1109/tcc.2024.3396476