Optimizing and Utilizing Key-Value Stores for Cloud Native Infrastructure

The Environmental and Societal Impacts of Data Centers

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
James Draudt

11/08/2024

Technical Project Team Members: N/A

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.

ADVISORS

Caitlin D. Wylie, Department of Engineering and Society

Briana Morrison, Department of Computer Science

Introduction

Data is collected everywhere—on every website, sensor, and video feed—and data collection only continues to grow. Taylor (2023, p.1) predicts that this year, "147 zettabytes... [will be] created, captured, copied, and consumed worldwide," with a projected 23% increase next year. This growth has driven a rapid push to build more data centers to house and analyze this vast amount of information. "Data centers are at the heart of almost every sector at the private, public, and governmental levels" (Hammadi & Mhamdi, 2014, pp. 1-2), making them essential to the economy and national defense. More data centers might seem beneficial, right? Not when we consider the serious environmental and community impacts they have. These massive facilities pollute as much as the aviation industry, comprising 2% of total greenhouse gas emissions worldwide (Bilal et al., 2014, p. 2), and place huge demands on local water and electrical infrastructure (Terri, 2022). Data storage practices are often inefficient and outdated. Within data centers, the Distributed /etc Directory (ETCD) database is widely used but presents significant performance issues and frequent challenges for engineers. This reliance on outdated practices only increases inefficiency. So, not only do data centers harm the environment and communities, but they are also inefficiently run.

To explore potential solutions, I will use the Latent Consequences framework (Helm, 1971), which examines both the intended functions of a technology and the unintended consequences it brings. Many technologies promise positive societal changes but end up having unexpected, often harmful, consequences. I will first analyze the issues with ETCD

and possible solutions, then examine the unintended consequences of the data center boom. Are the benefits of data centers worth the damage they cause locally and globally? How can we mitigate these consequences to achieve only the intended benefits?

Technical Topic

Moving data and processing to the cloud offers substantial advantages, but "[d]espite the many benefits of the cloud, such as high availability and scalability, most on-premise application architectures are not fully prepared to exploit the benefits of this environment" (Balalaie et al., 2016, p. 1). Balalaie (2016) argues that cloud technology still faces challenges before companies can fully realize its benefits. A major issue is that complex, specialized technologies form the backbone of many cloud services, leading to unforeseen consequences that hinder safe and reliable computing. One such technology is ETCD, a database essential to cloud-cluster technology. However, ETCD's complexity and small maintenance team make companies reluctant to adopt it, despite cloud technology's advantages. This hesitation causes inefficiencies, such as excessive maintenance, inefficient power use, and single points of failure. Consequently, companies waste millions of dollars and increase emissions due to their reliance on specialized technologies like ETCD.

This leads to my technical question: How can I optimize or replace ETCD to make cloud-cluster technology more reliable and efficient? To answer this, I will review ETCD's documentation (V3.5 Docs, n.d.) and experiment with configurations. Experimentation with

ETCD is crucial, as "both the Kubernetes control plane and the deployed application depend strongly, and sometimes unexpectedly, on the performance of the ETCD database" (Larsson et al., 2020, p. 1). Larsson et al. (2020) highlight how ETCD impacts various parts of a program, with minor changes often yielding surprising results, making careful analysis important. After assessing ETCD's optimization potential, I will explore alternatives. Some companies are already attempting this by integrating more established database technology, allowing administrators to choose the best database for their cluster's size. If feasible, I could test and benchmark other databases to gauge their suitability for cloud clusters. Success in this area could help mitigate inefficiencies, encouraging more companies to transition to optimized, efficient data centers.

STS topic

The expansion of cloud computing and data collection has led to a rapid increase in data centers. What are the societal and community impacts of this data center boom, and are these setbacks justified by technical advancements and tax revenue? For most, data centers are invisible enablers of advanced technology, but for those whose communities are overtaken by these massive complexes, they are anything but. These facilities have numerous latent consequences, consuming vast amounts of electricity and water while generating significant noise pollution. This strains local infrastructure, as seen in Northern Virginia, where "power supply needs to nearly double in the next 15 years to keep up with demand, driven almost exclusively by data center growth" (Pipkin, 2024, p. 1-2). Increased

consumption also contributes to rising greenhouse gas emissions: contributing 2% of global emissions and is set to surpass the entire aviation industry (Bilal et al., 2014). This issue is pressing because data-driven technologies like AI are projected to drive further data center construction.

Despite these consequences, many argue the benefits outweigh the costs. Companies constructing data centers often emphasize regional economic benefits, job creation, and partnerships with local communities to boost infrastructure. For example, Amazon (Wehner, 2023) claims that data centers provide well-paying jobs and contribute to regional wealth. However, these claims are likely skewed heavily towards big data centers, as they come from companies who directly make or profit off them. Others argue that they can build greener centers, but "there is an urgent need for energy-efficient solutions," which as of now don't exist (Garg, 2011, p1). While I agree that data centers can bring significant benefits, the drawbacks, particularly those unanticipated, can be severe.

Data centers aim to be close to consumers while remaining cost-effective, creating "sweet spots" that attract multiple facilities. Despite community opposition, local governments continue approving data center projects while the state evaluates their impacts. Many state and federal officials consider these decisions to be local matters, though the centers collectively strain the region's power grid and environment (Pipkin, 2024, p. 1). Often, local governments prioritize short-term tax revenue and job creation over the welfare of residents, and some even sell land to shell companies linked to data center developers, concealing plans until it's too late for public opposition. Once

construction is approved, there's minimal effort to shield communities from adverse effects. "Unlike other industries, data centers are largely self-regulating," with no governing agency (Terri, 2022, p. 6-7). This lack of oversight allows data centers to be built near residential neighborhoods. The "growth [of data centers] brings new challenges, including noise pollution and its impact on nearby communities" causing constant "mechanical whine" that residents hear every day and night (Hamilton, 2024, p. 1). Such noise pollution has serious health impacts; one resident described the toll: "She sees the signs of its toll — hypertension, cortisol — but she cannot stop it. No one can, because it does not sleep" (Terri, 2022, p. 5-6).

In places like Bluffdale, Utah, residents face water shortages and power outages because the nearby Utah Data Center "guzzles seven million gallons of water daily to operate" (Terri, 2022, p. 5). This data center, run by the NSA, gets preferential treatment because it is important to "national security," but wouldn't true national security ensure citizens have enough water and electricity to sustain their lives? Companies like Google pledge to be "water-positive" and operate solely on renewable energy. While this is a laudable ambition, there is no way to verify or enforce these goals. Many companies' primary aim is to maximize profit, so if they can cut corners, there's a good chance they will.

Another latent consequence of the data center boom is the vast amount of land needed for these massive structures. In built-up areas like Northern Virginia, open land is scarce. Location is crucial for data centers, but where is this untapped land coming from?

“Recently… data center developers have begun targeting rural land… often adjacent to national parks,” and these rural lands are typically forests or other green spaces. For example, in 2021, Prince William County “voted to add about 60 acres within… Prince William Forest to the data center overlay district,” which is land earmarked to become a national park and is adjacent to a Civil War landmark (Keep Massive Industrial Data Centers Away from Our National Parks, n.d., p. 2). While this article from the National Park Conservation Association draws a very concerning picture it's important to note that it is their job to raise concern so they can raise money and increase shock factor.

Using the framework of latent consequences, I will continue analyzing sources and seek personal accounts from affected communities to assess the full impact of these centers. I will also review the positive benefits that these centers bring to communities to understand both the positive and negative impacts. To conclude, I will examine how similar issues, like the growth of distribution centers for e-commerce, have been addressed in the past. I could then assess if my proposed solutions, discussed in the following paragraph, have been successful before.

Data centers can completely change communities. While they add jobs and tax revenue, they also heavily pollute, strain infrastructure, and take over land essential to the community. I believe latent consequences outweigh the intended benefits, and many reforms are needed to address these issues. The main problem is regulation. Companies should be required to increase infrastructure to the level they need before construction begins. Second, enhanced environmental legislation should set maximum levels for noise

pollution, net emissions, and water usage. Third, communities should have the right to vote on whether they want these facilities in their towns. Lastly, a review board should evaluate any land designated for data centers to ensure it does not disrupt national parks or other sacred sites. With this proposed solution I will look into potential ways of getting this legislation put in place and enforcing regulations going forward. If these changes are implemented, I believe ethical data centers that benefit communities rather than burden them are possible.

Conclusion

Data is the backbone of current technological advancement, but this digital boom has significant, often hidden, societal impacts. Current cloud technologies still face numerous challenges that lead to inefficiencies, lost profits, downtime, and higher greenhouse gas emissions. If these issues, especially those stemming from ETCD, are solved, it will encourage other companies to switch to the cloud, easing the load on massive data centers. But even if this transition is smooth, data centers are growing in both quantity and size, creating community and environmental challenges. These impacts were initially unforeseen but must be addressed if we are to move forward. Companies need to become more transparent and take greater responsibility for their effects on communities. Enhanced legislation and accountability are necessary to ensure that both corporations and communities can grow harmoniously.

References

Balalaie, A., Heydarnoori, A., & Jamshidi, P. (2016). Migrating to cloud-native architectures using microservices: An experience report. In A. Celesti & P. Leitner (Eds.), *Advances in Service-Oriented and Cloud Computing* (Vol. 567, pp. 201–215). Springer International Publishing. https://doi.org/10.1007/978-3-319-33313-7_15

Bilal, K., Malik, S. U. R., Khalid, O., Hameed, A., Alvarez, E., Wijaysekara, V., Irfan, R., Shrestha, S., Dwivedy, D., Ali, M., Shahid Khan, U., Abbas, A., Jalil, N., & Khan, S. U. (2014). A taxonomy and survey on green data center networks. *Future Generation Computer Systems*, *36*, 189–208. https://doi.org/10.1016/j.future.2013.07.006

Garg, S. K., Yeo, C. S., Anandasivam, A., & Buyya, R. (2011). Environment-conscious scheduling of HPC applications on distributed Cloud-oriented data centers. *Journal of Parallel and Distributed Computing*, *71*(6), 732–749. https://doi.org/10.1016/j.jpdc.2010.04.004

Hamilton, S. (2024). Challenges in urban data center design. *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, *269*(2), 716–723. https://doi.org/10.3397/NC_2024_0089

Hammadi, A., & Mhamdi, L. (2014). A survey on architectures and energy efficiency in Data Center Networks. *Computer Communications*, *40*, 1–21. https://doi.org/10.1016/j.comcom.2013.11.005

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

Kass , S. (2019). Building the next generation of sustainable data centers. *Sustainable Development Goals for Data Centers*. SDG-DC, DC. https://www.sdg-dc.com/whitepaper/next-generation-sustainable-data-centers

Keep massive industrial data centers away from our national parks. (n.d.). National Parks Conservation Association; National Parks Conservation Association. Retrieved November 3, 2024, from https://www.npca.org/advocacy/102-keep-massive-industrial-data-centers-away-from-our-national-parks

Larsson, L., Tärneberg, W., Klein, C., Elmroth, E., & Kihl, M. (2020). Impact of etcd deployment on Kubernetes, Istio, and application performance. *Software: Practice and Experience*, *50*(10), 1986–2007. https://doi.org/10.1002/spe.2885

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

Pipkin, W. (2024, June 18). *Energy demands for Northern Virginia data centers almost too big to compute*. Bay Journal. https://www.bayjournal.com/news/energy/energy-demands-for-northern-virginia-data-centers-almost-too-big-to-compute/article_d64ab388-2cb7-11ef-a753-4ffcc056f619.html

Taylor, P. (2023, November 16). *Data growth worldwide 2010-2025*. Statista. https://www.statista.com/statistics/871513/worldwide-data-created/

Terri. (2022, February 14). The staggering ecological impacts of computation and the cloud. *MIT Schwarzman College of Computing*. https://computing.mit.edu/news/the-staggering-ecological-impacts-of-computation-and-the-cloud/

V3. 5 docs. (n.d.). Etcd. Retrieved November 3, 2024, from https://etcd.io/docs/v3.5/

Wehner, R. (2023, October 12). *5 ways AWS data centers benefit local communities*. Amazon. https://www.aboutamazon.com/news/aws/aws-data-center-economic-impact-study