

Economic Viability of Outsourcing Cluster Computing to the Cloud: Insights for Technological Entrepreneurship and Regional Innovation in Honduras

Isaac Zablah¹; Salvador Diaz^{1,2}; Carlos Agudelo-Santos¹; Yolly Molina¹; Melania Madrid¹
Alicia Diaz³; Marcio Madrid^{1,4}

¹ Universidad Nacional Autónoma de Honduras (UNAH), Honduras, jose.zablah@unah.edu.hn, salvador.diaz@unah.edu.hn, carlos.agudelo@unah.edu.hn, yolly.molina@unah.edu.hn, melania.madrid@unah.edu.hn, marcio.madrid@unah.edu.hn

² Secretaría de Estado en los Despachos de Salud (SESAL), Honduras

³ Universidad Católica de Honduras (UNICAH), Honduras, a.sofiadias2002@gmail.com

⁴ Universidad Pedagógica Nacional Francisco Morazán (UPNFM), Honduras

Abstract – *This study evaluates, from a cost and entrepreneurial opportunity perspective, the viability of maintaining a university cluster composed of ten HP ProDesk 600 G1 nodes (40 cores, 60 GB RAM) versus contracting equivalent capacity on Amazon EC2 (C6i.2xlarge instances). A detailed total cost of ownership (TCO) model was developed, incorporating capital investment, maintenance, electricity (USD 0.22/kWh), internet connectivity (200 Mbps), and cloud tariffs under both On-Demand and three-year Savings Plan modalities. At 30% utilization, the local cluster achieves a unit cost of USD 0.025/core-hour, while EC2 ranges from USD 0.048 to 0.105/core-hour. Sensitivity curves indicate that the cloud only becomes cost-effective when utilization exceeds 65% or electricity surpasses USD 0.45/kWh. The analysis further explores implications for deep-tech startups, university spin-offs, and incubator programs: while the cloud provides elasticity and access to advanced hardware, it may erode early-stage margins. The study proposes a reproducible decision-making framework tailored to emerging innovation ecosystems in Latin America, highlighting the strategic value of hybrid infrastructures for balancing cost, flexibility, and regional value retention.*

Keywords– TCO, HPC, IaaS, technology entrepreneurship, Honduras.

I. INTRODUCTION

In today's digital age, compute-intensive (HPC) capacity has established itself as a fundamental pillar for applied research and the development of digital solutions with social impact. However, institutions in the Global South [1], [2] face a critical dilemma: Is it more convenient to invest in the construction and maintenance of on-premises HPC clusters or outsource these needs to cloud providers?, This challenge is compounded by regional constraints such as scarce public-private financing, high energy costs, and uneven connectivity, as documented by Rosa [3]. In addition, the decision transcends the technical, since it directly influences entrepreneurship and regional development by determining the access of startups and local scientists to advanced computational resources. This article takes up the seminal work of De Alfonso et al. [4], which

analysed the costs of physical and virtual clusters and expands it by incorporating an innovation and development perspective.

In addition to the aforementioned studies, recent literature has addressed various facets of this dilemma. Odeh et al. [3] have analysed the specific barriers faced by institutions in the Global South to adopt cloud solutions, highlighting the importance of considering the local context of telecommunications infrastructure and regulatory policies in decision-making.

On the other hand, Juhász et al. [6], provide a detailed comparison of cost and performance between on-premises clusters and cloud-based HPC solutions in various scenarios, providing a practical framework for evaluating the available options and their long-term economic implications. These studies complement the analysis by highlighting both regional challenges and the opportunities that HPC can offer for innovation and economic growth.

Internationally, Smith and Harrell [7] They provide a detailed comparison of cost and performance between on-premises HPC clusters and cloud-based solutions in various scenarios, providing a practical framework for evaluating the available options and their long-term economic implications. These studies complement the analyses by highlighting both regional challenges and the opportunities that HPC can offer for innovation and economic growth.

By assessing the economic, technical, and social trade-offs between local clusters and cloud solutions, this study seeks to guide the strategic decisions of universities, innovation centers, and local governments in the Global South.

II. METHODOLOGY

An explicit comparative design was adopted in two scenarios: an on-premises cluster and a cloud infrastructure. This approach was selected to directly assess the differences in cost, performance, and accessibility between the two options, which is crucial to guide strategic decisions in resource-constrained contexts, following similar published studies [8]. The study combines empirical measurements of electricity

consumption, actual quotes for refurbished hardware, and Amazon EC2 public pricing, following methodologies like those used in compute infrastructure cost analysis [9].

The cluster is at a modest *fab-lab* university where its peaks of use coincide with AI prototyping courses and open innovation challenges at the Faculty of Medical Sciences of the National Autonomous University of Honduras (FCM/UNAH) [10]. This environment is representative of the conditions in which many institutions and startups operate in the Global South, where access to computational resources is critical for the development of innovative solutions [11]. Cluster nodes, configured with HP Prodesk 600 G1 hardware [12] refurbished, consume 70 W at load and 25 W at idle, measured with a Fluke 1730 Energy Analyzer [13]. The Power Usage Efficiency (PUE) factor was estimated at 1.8, given the use of basic cooling schemes, which is common in budget-constrained environments [14].

Region selected *us-east-1* for its lower latency for Central America (< 100 ms), a critical factor for applications that require real-time responses [15]. The C6i.2xlarge instance (8 vCPUs, 16 GiB RAM) was selected for its performance characteristics, including higher throughput per core and higher memory bandwidth, as documented by AWS. Internal observations indicate performance comparable to that of on-premises HP ProDesk G1 nodes [11]. To reflect the local reality with limited capital, Spot Instances and Premium GPUs were excluded, as their use may involve additional costs and management complexities not accessible to small businesses. Equation 1 was used to make the hourly estimates [17]:

$$C_{\text{per core hour}} = \frac{C_{\text{TCO}}}{N_{\text{cores}} + H_{\text{year}} + U} \quad (1)$$

Where:

- C_{TCO} : Total annual cost of ownership
- N_{cores} : Total number of cores
- H_{year} : Hours of the year (8760h)
- U : Average fractional utilization (0–1)

The traditional equation of total cost of ownership (TCO) is extended by adding components specific to the context of startups and innovation environments:

1) *Platform Learning Costs*: 20 hours of initial engineering were estimated for familiarization with the cloud platform, based on studies on the learning curve in adopting cloud technologies [18].

2) *Optional AWS Premium Support*: A cost of 1200 USD/year was included for premium support, an option that can be crucial for startups that require quick technical assistance [19].

3) *Cloud Credits*: Considered possible credits of up to \$5000 that accelerators typically provide to startups, which can significantly reduce the initial costs of cloud adoption [20].

In addition, sensitivity analyses were performed to assess the impact of five utilization levels (10%–90%) and four electricity prices (0.10–0.50 USD/kWh) on TCO, allowing for

robust evaluation under various operating conditions [21]. Equation 2 was used to show how the TCO was calculated [22]:

$$C_{\text{TCO}} = \frac{\text{CAPEX}}{L} + O_{\text{maint}} + P_{\text{internet}} + (PUE \times E_{\text{IT}})P_e + C_{\text{compute}} \quad (2)$$

Where:

- CAPEX : Capital expenditure
- L : Useful life considered for amortization (years)
- O_{maint} : Annual maintenance cost (USD/year)
- P_{internet} : Annual data link fee (USD/year)
- PUE : Laboratory Power Usage Effectiveness
- E_{IT} : Annual energy consumed by the nodes (kWh)
- P_e : Electricity price (USD/kWh)
- C_{compute} : Annual cost of cloud computing

This analysis was carried out using a Monte Carlo simulation approach to model uncertainty in electricity costs and resource utilization, following standard methodologies in risk analysis for infrastructure investment decisions [23].

III. RESULTS

Considering the fixed costs of amortization, maintenance and connectivity, the local cluster reaches an annual TCO of just USD 2592, a figure that represents only 17% of the expense that would be involved in operating the same capacity in EC2 under on-demand mode (USD 14892) and 62% of the TCO associated with a three-year Savings Plan contract (USD 4166). This differential translates into a unit cost of \$0.025/core-hour for the on-premises solution, versus \$0.048 and \$0.105/core-hour for the cloud options evaluated, thus consolidating a relative advantage of 48–76% in favor of the on-premises cluster in the 30% average utilization scenario. This can be seen in Table I.

TABLE I
ANNUAL BUDGET WITH 30% UTILIZATION IN USD/YEAR

Concept	Local Cluster	EC2 On-Demand	EC2 Savings-Plan
Hardware amortization	500.00	-	-
Maintenance	150.00	-	-
Internet (200Mbps)	1200.00	-	-
Energy (U=30%)	742.00	-	-
Compute	-	14892.00	4166.00
Total Annual	2592.00	14892.00	4166.00

Instead, Table II reveals the inverse relationship between utilization level and electricity price required for the on-premises cluster to match the cost of the cloud under *the Savings Plan*. At a moderate load of 50%, the equilibrium threshold drops to only 0.03 USD/kWh, a value practically

unattainable in the Central American energy matrix; however, as utilization increases to 80%, the critical price rises to 0.15 USD/kWh, approaching the regional industrial average. Above 90% utilization, the on-premises cluster would only lose advantage if electricity exceeded \$0.18–\$0.21/kWh, suggesting that in contexts of sustained high demand, the cloud can become competitive even without resorting to promotional credits. These data confirm that the decision to outsource does not depend solely on the absolute energy cost, but on the dynamic combination between electricity tariff and computational usage pattern.

TABLE II
EQUILIBRIUM ELECTRICITY PRICE THAT EQUALS THE LOCAL COST TO THE SAVINGS PLAN

Utilization (U=%)	Balance price/kWh (USD)
50%	0.03
60%	0.08
70%	0.12
80%	0.15
90%	0.18
100%	0.21

The evolution of the cost per core-hour when the demand for the system grows from 10% to 100% utilization. The curve of the local cluster shows a hyperbolic descent: it starts from ≈ 0.08 USD/core-h at 10% and falls to ≈ 0.012 USD/core-h when it reaches 100%. The two horizontal lines correspond to constant cloud pricing: \$0.042/core-h for EC2 On-Demand and \$0.012/core-h under a three-year Savings Plan. It is observed that the local cluster exceeds the cost of the On-Demand option only below $\sim 20\%$ utilization; at that point it is cheaper than On-Demand payment, but it does not match the ultra-reduced rate of the Savings Plan until it reaches full hardware occupancy. This dynamic confirms that, in academic environments with intermittent loads, the *on-premises* solution dominates over the *on-demand* cloud, while the multi-year savings modality maintains a relative advantage except in scenarios of continuous intensive use, as shown in Figure 1.




Fig. 1 Cost per core-hour vs. utilization: Compares on-premises cluster to EC2 On-Demand and SavingsPlan.

Figure 2 plots the total annual cost (TCO) vs. cluster utilization. The green curve, corresponding to local infrastructure, grows in an almost linear fashion from $\sim$ USD 2,100 yr at 10% to $\sim$ USD 4100 yr at 100%, driven mainly by the energy component. In contrast, the purple line (EC2 On-Demand) starts below the cluster at very low loads ($\approx$ USD 1500 per year at 10%), but its slope is six times higher due to the variable price per core-hour; It intercepts the local line at around 20% utilization and drastically exceeds it from there. The blue line (EC2 Savings Plan) exhibits the lowest slope: it remains the cheapest option up to $\sim 65\%$ utilization, at which point it converges with the on-premises solution; above this threshold, the relatively moderate energy increase of the local cluster makes it competitive again. This behaviour shows that the advantage of each modality depends on the combination of annual load and pricing strategy, reinforcing the need for a hybrid approach to balance cost and flexibility.




Fig. 2 Total Annual Cost (TCO) vs. Utilization: Reveals where annual expenses intersect.

The Equilibrium Electric Rate, i.e., the price per kWh at which the cost of the on-premises cluster would equal that of the cloud under the Savings Plan based on utilization. The curve exhibits an almost exponential growth: with a 20% load, a negative rate (economically unrealistic) would be enough to match costs; at 50% utilization the threshold is ≈ 0.03 USD/kWh, still well below the Honduran rate of 0.22 USD/kWh, while at 80% it is close to 0.15 USD/kWh and at 100% it reaches ≈ 0.21 USD/kWh.

These values confirm that, under the electricity prices in force in the country, the local cluster maintains an economic advantage throughout most of the spectrum of use, ceding ground to the cloud only in scenarios of almost continuous occupation and very high energy efficiency, as shown in Figure 3.




Fig. 3 Equilibrium electricity price vs. utilization: Indicates which kWh rate equals both models.

Figure 4 illustrates how the utilization threshold required for the on-premises cluster to match the cost of the Savings Plan decreases as the price of electricity increases. With a low rate of 0.10 USD/kWh, the cluster is still cheaper up to almost 95% occupancy; at the average Honduran price of 0.22 USD /kWh that break-even point is reduced to ~65%. If energy were to rise to 0.35 USD/kWh, a utilization of 45% would be enough for both models to cost the same, and with an extreme rate of 0.50 USD/kWh the tie would already be reached around 35%.

The steep slope of the curve highlights the high sensitivity of the decision to outsource to relatively small variations in energy costs, a critical factor in regions with tariff volatility or government subsidies at stake.




Fig. 4 Balancing Utilization and Power Price: Shows the load level where the local cluster matches the cloud for different energy costs.

The results show that, in the base scenario (30% utilization), the local cluster has a lower cost. However, by incorporating the AWS Savings Plan, which offers discounts for long-term usage commitments, the cloud solution approximates the cost of the on-premises cluster starting at utilization (u) $\geq$ 65%. This shift highlights the competitiveness of the cloud in high-utilization contexts.

A key aspect emerges when considering cloud credits for start-ups, which shift the crossover point to approximately 50% utilization. These credits, often available through innovation support programs, significantly reduce upfront cloud costs,

allowing start-ups to quickly experiment with high-performance computing (HPC) without the need for CAPEX (*Capital Expenditure*) investment.

These findings underscore the viability of the cloud as a flexible and accessible alternative, especially for resource-constrained organizations looking to adopt HPC without significant financial commitments.

IV. DISCUSSION

The findings suggest that universities and governments promoting innovation hubs should maintain a local hardware floor for early-stage validation, while negotiating cloud credits to scale prototypes. Hybrid operating models facilitate the transformation of academic projects into spin-offs, reducing initial financial risk and boosting resilience to demand peaks. From a regional development perspective, local infrastructure feeds regional value chains in services such as maintenance and training, whereas cloud usage transfers spending to external providers. Public policy must balance economic efficiency with value retention. Quantitative results and sensitivity curves enable comparing economic viability of each alternative, aligning with recent literature on high-performance cloud computing [24][25].

Table 1 confirms that, under the baseline scenario of 30% utilization, the local cluster retains an absolute advantage in both annual cost and cost per core-hour compared to the two Amazon EC2 modalities. This finding aligns with the results reported by De Alfonso et al. [4], who demonstrated the competitiveness of modest in-house infrastructures when capital is already amortized, and electricity prices remain below 0.25 USD/kWh.

Figure 1 extends this pattern: the local cluster outperforms the On-Demand cloud option beyond ~20% utilization, while the Savings Plan modality only matches local hardware at full occupancy. These results support the conclusions by Jackson et al. [26], who argue that cloud services are best suited for bursty, short-term workloads rather than for sporadic academic computing.

The energy equilibrium analysis presented in Table 2 and Figure 3 highlights the system's strong sensitivity to electricity pricing. As utilization increases, the breakeven electricity rate converges toward the 0.18–0.21 USD/kWh range, which closely matches average residential and institutional tariffs in Honduras. This observation aligns with broader regional analyses indicating that electricity costs play a decisive role in determining the cost-effectiveness of on-premises clusters versus cloud-based alternatives [27], [28]. In scenarios of increased electricity price volatility, common in Latin America's deregulated markets; cost advantages may shift toward cloud platforms, especially when promotional programs such as AWS Activate are leveraged by eligible organizations [30].

The pattern illustrated in Figure 2 demonstrates that the TCO for EC2 OnDemand scales steeply, becoming prohibitive

above ~25% utilization. In contrast, the Savings Plan trajectory remains nearly flat, intersecting the on-premises cost curve around ~65%. Moreover, Figure 4 reveals that the utilization breakeven threshold falls sharply when the electricity price exceeds ~0.35 USD/kWh, supporting the hybrid strategies recommended for Latin American campuses [30], [31].

The suggested strategy is supported by regional development frameworks. United Nations Development Programme (UNDP) Accelerator Labs emphasize the value of maintaining local infrastructure for early validation, capacity building, and prototyping, while leveraging digital credits and scalable cloud resources for growth. This dual approach balances cost efficiency, innovation flexibility, and value retention within local ecosystems [32].

To obtain, costs and Honduras-specific simulations, with the data on Table I and Figure 4, we can discuss three representative workloads scenarios: (i) product-validation (~50000 core-hours over three months), (ii) a one-week research and development (R&D) sprint (~15000 core-hours), and (iii) a university hackathon (~20000 core-hours). The decisions are the same, with a well-utilized local cluster minimizing average cost; cloud Savings Plans narrow the gap, while On-Demand remains the high-cost baseline. Crucially, startup credits reduce early-stage cash outlays, making cloud attractive for rapid iteration and time-to-market; graphics processing unit (GPU)-backed prototyping further strengthens this case [29], [9].

Environmental sustainability that goes beyond money, in this case, the carbon outcomes also depend on the grid mix and the Power Usage Effectiveness (PUE), in addition to CAPEX and OPEX measures. When local usage is not optimal, selectively transferring workloads to cloud regions with low Power Usage Effectiveness (PUE) and low carbon emissions can help reduce your environmental footprint. On the other hand, a well-used on-premises system is still competitive in terms of both cost and emissions [14], [9].

Readiness of the workforce and ability to innovate, for these are required to have a lasting effect; you need DevOps/MLOps, Financial Operations (FinOps), security, and institutional learning. A step-by-step plan that involves quick certifications and easy-to-follow guides, medium-term platform development and cost management, and long-term partnerships between universities and industries helps small and medium-sized businesses (SMEs) and startups learn and innovate over time [24], [25], [32].

Experience with regional policy levels and startup scenarios shows that sectoral guidance on portability, continuity, and data location, as well as cloud-friendly public procurement and regulatory sandboxes, speeds up diffusion and cross-border scaling [27], [25]. In Honduras, matched cloud vouchers for Technology Readiness Level (TRL) 3–6 and academic connectivity reduce the risk of adoption. Startups often utilize short GPU-assisted experiments, implement elastic bursts around pilots and fundraising, and simplify market access through compliance [32], [29].

IV. CONCLUSIONS

This study provides a comparative techno-economic assessment of local on-premises clusters versus Amazon EC2 cloud alternatives for high-performance computing (HPC) in academic and innovation-oriented contexts in Latin America. Under baseline conditions of moderate utilization (30%), the on-premises cluster demonstrates a clear economic advantage, achieving a total cost of ownership (TCO) 83% lower than EC2 On-Demand and 38% lower than the three-year Savings Plan. Unit cost per core-hour reinforces this advantage, positioning the local infrastructure as 48–76% more cost-efficient in the tested mid-utilization regime.

Sensitivity analyses reveal that this advantage is strongly influenced by the utilization rate and electricity pricing. The breakeven electricity cost that renders cloud options more economical only emerges above 0.21 USD/kWh under full utilization, a threshold closely aligned with Honduras's current industrial tariff. Conversely, when electricity prices rise or workloads approach continuous operation ($\geq 90\%$), cloud solutions, especially with long-term pricing commitments or startup credits; this becomes increasingly competitive.

Figures 1 through 4 highlight critical transition zones in both utilization and power cost dimensions. Specifically, the On-Demand model is consistently outperformed by local infrastructure beyond ~20% usage. The Savings Plan maintains competitiveness until ~65% usage but is overtaken when local infrastructure operates at high efficiency and low energy cost. Promotional credits offered through programs like AWS Activate can shift this breakeven point closer to 50% utilization, making the cloud a viable alternative for early-stage innovation and prototyping.

In short, a well-sized local cluster keeps the average cost per core-hour low when usage is steady, but cloud outsourcing is better for bursts, GPU-assisted prototyping, and quick deployment in a region. Combining clear cost modeling with environmental metrics (like energy use and PUE) and workforce readiness gives universities, small and medium-sized businesses, and startups a way to make decisions that they can act on. At the policy level, matched cloud vouchers for Technology Readiness Level (TRL) 3–6, sector-specific guidance, and regulatory sandboxes, along with targeted human-capital programs, can help people start using the technology early without losing the chance to build up their own skills.

ACKNOWLEDGMENT

We express our deepest gratitude to Professor Antonio Garcia Loureiro, PhD, from University Santiago de Compostela for his invaluable methodological and technological guidance, which significantly shaped the development of this work. We also extend our sincere appreciation to Prof. Dr. Marco Tulio Medina, FAAN, former Dean of the Faculty of Medical Sciences of National Autonomous University of Honduras (2010–2018), for his visionary conceptualization and support in promoting HPC in health research.

REFERENCES

- [1] L. Rodríguez, "Introduction to the thematic dossier: the Global South and the construction of a new International System", *OASIS*, no. 34, pp. 3–10, jul.–dec. 2021. DOI: 10.18601/16577558.n34.02
- [2] S. Haug, J. Braveboy-Wagner and G. Maihold, "The 'Global South' in the Study of World Politics: Examining a Meta Category," *Third World Quarterly*, vol. 42, no. 9, pp. 1923–1944, Sep. 2021. DOI: 10.1080/01436597.2021.1948831.
- [3] F. Rosa, "Internet Interconnection Infrastructure: Lessons from the Global South," *Internet Policy Review*, vol. 10, no. 4, art. 1583, Nov. 2021. DOI: 10.14763/2021.4.1583.
- [4] C. de Alfonso, M. Caballer, F. Alvarruiz, and G. Moltó, "An economic and energy-aware analysis of the viability of outsourcing cluster computing to a cloud," *Future Gener. Comput. Syst.*, vol. 29, no. 3, pp. 704–712, 2013. DOI: 10.1016/j.future.2012.08.014
- [5] M. Odeh, A. Garcia-Perez and K. Warwick, "Cloud Computing Adoption at Higher Education Institutions in Developing Countries: A Qualitative Investigation of Main Enablers and Barriers," *International Journal of Information and Education Technology*, vol. 7, no. 12, pp. 921–927, Dec. 2017. DOI: 10.18178/ijiet.2017.7.12.996.
- [6] Z. Juhász, "Quantitative cost comparison of on-premise and cloud infrastructure based EEG data processing," *Cluster Computing*, vol. 24, pp. 625–641, 2021. DOI: 10.1007/s10586-020-03141-y
- [7] P. Smith, S. Harrell, A. Younts and X. Zhu, "Community Clusters or the Cloud: Continuing cost assessment of on-premises and cloud HPC in higher education," in *Proc. Practice and Experience in Advanced Research Computing (PEARC '19)*, 2019, pp. 1–4. DOI: 10.1145/3332186.3333155
- [8] M. Ferretti and L. Santangelo, "Cloud vs On-Premise HPC: A Model for Comprehensive Cost Assessment," in *Parallel Computing: Technology Trends*, I. Foster et al. (eds.), IOS Press, Mar. 2020, ch. 26, pp. 69–81. DOI: 10.3233/APC200026
- [9] C. Coloma, M. Silva and J. Knobloch, "A cost and energy consumption comparison of on-premises HPC versus AWS EC2 for GPU-accelerated workloads," *EPJ Web of Conferences*, vol. 295, 07022, 2024. DOI: 10.1051/epjconf/202429507022
- [10] Facultad de Ciencias Medicas. UNAH. Available: <https://fcm.unah.edu.hn>. [Accessed: 30-Jul-2025].
- [11] R. B. Vieira, L. P. Bresciani and I. C. dos Santos, "Fab Labs Network in Developing Countries: Knowledge Spillover Effects or Managing Technology Development within the Scarcity Economy?," *Proc. PICMET '17 Portland International Conference on Management of Engineering and Technology*, Portland, OR, USA, Jul. 2017, pp. 1–8. DOI: 10.23919/PICMET.2017.8125265
- [12] HP ProDesk G1. HP. Available: https://www.hp.com/hpinfo/newsroom/press_kits/2013/HPDiscoverBarcelona/ProDesk_600_G1_Desktop_Mini_QuickSpecs.pdf?msocid=34e4a7517efe6d3b34d9b48c7f596c3. [Accessed: 30-Jul-2025].
- [13] Fluke 1730, Fluke Instruments Corp. Available: <https://www.fluke-direct.com/pdfs/cache/www.fluke-direct.com/1730-us/datasheet/1730-us-datasheet.pdf>. [Accessed: 30-Jul-2025].
- [14] R. Miller, "Uptime Institute: The Average PUE is 1.8," *Datacenterknowledge.com*, Available: <https://www.datacenterknowledge.com/uptime/uptime-institute-the-average-pue-is-1-8> [Accessed: 31-Jul-2025].
- [15] "CloudPing Latency – Latin America (Mexico City) to us-east-1 (Northern Virginia)," *CloudPing*, 2025. Available: <https://www.cloudping.co>. [Accessed: Jul. 30, 2025].
- [16] "Amazon EC2 C6i Instances", Amazon Web Services. Available: <https://aws.amazon.com/ec2/instance-types/c6i/>. [Accessed: 31-Jul-2025]
- [17] Broadcom, "How is Hourly Cost calculated in the AWS EC2 Instance Asset report?," *CloudHealth Knowledge Center*. Available: Knowledge Center Article #282671. [Accessed: Jul. 30, 2025].
- [18] M. Rahman and D. Kirby, "Learning-Curve Applications for Emerging Products and Services," *Proceedings of the 8th North American International Conference on Industrial Engineering and Operations Management*, Houston, TX, USA, Jun. 2023. DOI: 10.46254/NA8.20230024
- [19] Amazon Web Services, "Compare AWS Support Plans", AWS Premium Support FAQs and Pricing, Available: <https://aws.amazon.com/premiumsupport/pricing/> [Accessed: Jul. 30, 2025]
- [20] Spacelift Blog, "AWS Activate Program – Credits for Startups: \$5,000 to \$100,000 depending on partner and stage," Spacelift, Available: <https://spacelift.io/blog/aws-startup-credits>. [Accessed: Jul. 30, 2025].
- [21] CloudOptimo, "Understanding Total Cost of Ownership (TCO) in Cloud Computing," CloudOptimo Blog, Available: <https://www.cloudoptimo.com/blog/understanding-total-cost-of-ownership-tco-in-cloud-computing>. [Accessed: Jul. 30, 2025].
- [22] J. Hand, "Total Cost of Ownership (TCO) Model for Storage," SNIA. Available: https://www.snia.org/sites/default/files/SSSI/SNIA_TCO_Whitepaper_03-2021.pdf. [Accessed: Jul. 30, 2025]
- [23] A. Cunha Jr., R. Nasser, R. Sampaio, H. Lopes y K. Breitman, "Uncertainty quantification through Monte Carlo method in a cloud computing setting," *arXiv preprint*, Available: <https://arxiv.org/abs/2105.09512> (Accessed: Jul. 30, 2025)
- [24] Ontario's Universities, "Supporting Ontario's Regions through Innovation," Ontario Universities, Available: <https://ontariosuniversities.ca/news/supporting-ontarios-regions-through-innovation/>. [Accessed: Jul. 30, 2025].
- [25] Connected Places Catapult, "Hubs of Innovation Report," United Kingdom, Available: <https://cp-catapult.s3.amazonaws.com/uploads/2021/03/Hubs-of-Innovation-Report-final-PDF.pdf> [Accessed: Jul. 30, 2025].
- [26] K. R. Jackson et al., "Performance analysis of high-performance computing applications on the Amazon Web Services cloud", *Proceedings of the 2nd International Conference on Cloud Computing Technology and Science*, Indianapolis, IN, USA, 2010, pp. 159–168. DOI: 10.1109/CloudCom.2010.75
- [27] M. Whyman, "Cloud Computing and Data Localization in Latin America," *Center for Strategic and International Studies (CSIS)*. Available: https://csis-website-prod.s3.amazonaws.com/s3fs-public/2023-06/230607_Whyman_Cloud_Computing_1.pdf. [Accessed: Jul. 30, 2025].
- [28] Organización Latinoamericana de Energía (OLADE), "The Rise in Energy Prices in Latin America and the Caribbean Remains Below the Increase in Overall Inflation," Available: <https://www.olade.org/en/noticias/the-rise-in-energy-prices-in-latin-america-and-the-caribbean-remains-below-the-increase-in-overall-inflation/>. [Accessed: Jul. 30, 2025].
- [29] Amazon Web Services, "AWS Activate – Credits and Support for Startups," Available: <https://aws.amazon.com/activate/>. [Accessed: Jul. 30, 2025].
- [30] B. Martens, M. Walterbusch, and F. Teuteberg, "Costing of Cloud Computing Services: A Total Cost of Ownership Approach," in *Proc. 45th Hawaii Int'l Conf. System Sciences (HICSS-45)*, 2012. doi:10.1109/HICSS.2012.186
- [31] Dell Technologies, "The Public Cloud Breakeven Point: Choosing Between On-Premises vs the Public Cloud for Workload Deployment," Dell Technologies InfoHub, Available: <https://infohub.delltechnologies.com/l/choosing-between-on-premises-vs-the-public-cloud-for-workload-deployment/the-public-cloud-breakeven-point-4/>. [Accessed: Jul. 30, 2025]
- [32] United Nations Development Programme (UNDP), "Learning from Innovation in Latin America and the Caribbean," Accelerator Labs regional knowledge brief, Available: https://www.undp.org/sites/g/files/zskgk326/files/2023-11/acclabs_4-en.pdf. [Accessed: Jul. 30, 2025]