Green Cloud Computing: An Analysis of Energy-Saving Policies, Procedures, and Technologies

Rana Veer Samara Sihman Bharattej Rupavath

rupavathrana@gmail.com

MS in Computer Science

Abstract— The term "green computing" refers to the use of computers and the resources connected with them in an eco-friendly and ecologically responsible manner. This use is related to the information technology industry and services. The practice of using the least number of resources possible is another component of green cloud computing. For the purpose of controlling the waste of resources, a distributed service system that includes a job dispatcher and a physical server's virtual machine modules for generating queueing models, enforcing ISN policies, and modelling SN policies are all put to good use. It is well known that virtual machines are installed on physical machines and the number of hosts is reduced through the use of methods like virtualization, consolidation, and migration as well as metrics and methods for green cloud and algorithm strategies. There has been much research on this subject..

Keywords—machine learning, green computing, classification, regression, clustering, digital marketing, big data,

I. Introduction

In cloud computing, performance enhancement is demanded, yet energy consumption and carbon dioxide emissions are rising, raising environmental and economic costs. A cooling system uses power and generates heat to cool data center processor chips. Data centers cooling equipment like chillers and air conditioners are expensive and energy-intensive. It may not schedule bigger tasks.[1] Virtualization, which divides a real server into several virtual servers operating simultaneously, has also improved data center efficiency. Consolidating virtual machines, virtualizing servers, adjusting the voltage and frequency on the fly, live migrating across servers, and switching between resources are all part of green computing. In the Green Cloud Computing Survey's combined approach to resource allocation and server consolidation, a small number of servers or physical machines PMs are placed into energy-saving sleep mode until a virtual machine VM is assigned to them. The goal is to run v VMs in a dynamic setting with fewer PMs. Live VM migration is used for server consolidation. Timely and sequential allocation of resources is essential. [2].

II. Objective

The research aimed to fulfill the following objectives:

- The energy consumption of data centers from the viewpoint of the system
- Individual software-level data center power modeling
- At the application level

III. Methodology

Power-hungry yet crucial, data centers provide Internet services on a massive scale. If we want to create and improve energy-efficient operations in data centers and cut down on wasteful energy use, we need power utilization models. Energy efficiency in data centers has become more important and difficult to achieve in recent years.

High data availability can only be ensured if the data center's many components are working as intended and using as little power as necessary. The core of any IT system is the technical infrastructure, which includes things like power supply, technical cooling, and technical security. Even a temporary disruption in the physical infrastructure may have a significant effect on the availability of IT services. A green data center must be both energy efficient and environmentally friendly. A green or sustainable data center is one in which all of the systems, notably the mechanical and electrical frameworks, are designed to reduce energy consumption. Carbon emissions are reduced, money is saved, and productivity is boosted. These green data centers aid cutting-edge businesses in cutting costs and carbon footprints.

IV. THE ENERGY CONSUMPTION OF DATA CENTERS FROM THE VIEWPOINT OF THE SYSTEM

This section evaluates the whole data center to determine the necessary levels depending on the amount of power used. The model of a data center that was used for this research. Both the hardware and the software of a computer system make up the whole of any computer system. In addition, a data center is comprised of two basic elements: its software and its hardware.[3]

These layers have the potential to be improved or optimized in order to reduce the amount of electricity that is used by data centers. Three distinct layers make up the software layer: the OS, the virtualization layer, and the application layer. A software taxonomy is presented as a potential tool for application during the development of a sustainable data centre. The goals of green cloud computing may be accomplished via the implementation of a variety of solutions at the level of individual programmers. In addition to software-based measures, non-data-center-specific elements, such as regulations and rules set by the government, organizations, and renewable energy sources, are also taken into consideration in order to accomplish the objectives of green cloud computing. [4]

V. INDIVIDUAL SOFTWARE-LEVEL DATA CENTER POWER MODELING

How much power can be saved by implementing changes to the operating system, the virtualization layer, and the applications themselves?

At the OS Level

The OS sits between the application layer and the hardware layer. Applications are primarily responsible for increasing the demand for resources, and it is the OS's responsibility to allocate and control those resources. While the hardware itself is the primary energy drain, [5]it is crucial to monitor OS-level power events if software-level optimization of data center energy use is to be achieved. the power consumption of the OS's many functions. It was shown that data route and pipeline topologies, which permit numerous challenges such as out-of-order execution, squander 50% of the total power of the OS processes analysed. The clock alone consumes 34% of the total energy, with the different cache levels using the remainder.

Operating systems use OSPM, or operating system power management, to manage and switch between the platform's several power modes. Switching to the most efficient power mode for the platform or system, which may be utilised by all of the devices and components inside the platform or system, is made possible by Operating System Power Management (OSPM). This idea also goes by the moniker "Operating System-Directed Configuration and Power Management" (OSPM).

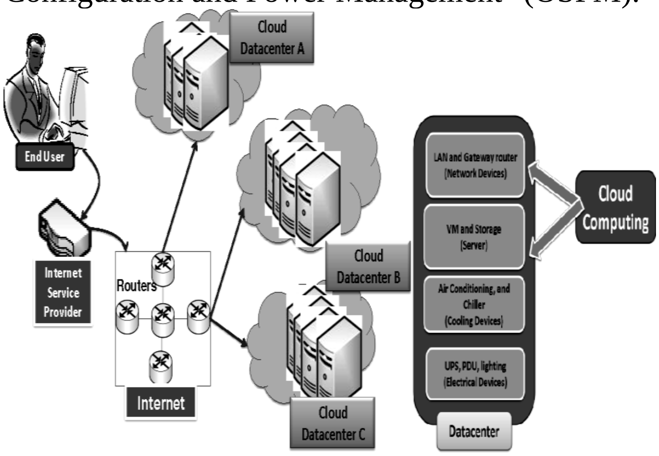

FIG 1. GREEN CLOUD COMPUTING

Since running voltage and energy management were mostly relocated from the hardware and software level to the operating system, the trade-off between quality and power efficiency has been the subject of much study and analysis. Polar, available for Linux, is a programme that can detect energy-saving setups automatically. No changes need to be made to the applications themselves since Polar will automatically choose the best settings for them based on data collected from the whole system. They don't only provide effective system control but also account for bank shots (alterations to configuration [6] parameters unrelated to power management). Developers of operating systems saw energy as a resource on par with time. Once energy was seen as just another resource available to the OS, changes to the internals of the OS (such as locking mechanisms) were made to create energy-aware OSes. Without limiting the frequency-scaling mechanism's usage to periods when no real-time tasks are executing, we may create energy-aware real-time scheduling by modifying the deadline scheduler using the curfew subsystem. They described the structural challenges they had while attempting to bring the GRUB-PA method to a real OS like Linux. Testing using a multi-core ARM architecture validated the efficacy of their proposed solution.

VI. AT THE APPLICATION LEVEL

Since cloud computing has a perpetual problem with energy efficiency, several suggestions have been made to enhance it on the application level. However, a deeper analysis at the application level is required to allow adaptive controls and to limit global energy usage in light of the increasing complexity of contemporary workloads. In an effort to reduce the carbon footprint of data centers, many batch programmes are now being run on virtual machines. They examined the app's functionality, estimated the amount of power required for each function, and arrived with an overall figure.[7] The evaluation seeks to find the most effective use of resources by measuring software efficiency in terms of performance and power consumption per job, especially while working with shared resources and in a variety of circumstances that need different energy profiles. Data-intensive applications were separated from communication-intensive ones. Energy efficiency has not been a priority in the execution of data-intensive systems that produce, process, and transport vast amounts of data. Problems with data management, transfer, and storage might result in significant energy use. One or more dependent services compose a communication-intensive application, with their respective communication flows often being separate from one another. Power is used at high rates due to communication traffic, yet there are several methods of dynamic power management that may improve efficiency.

SaaS, or Software as a Service, describes cloud services as a software delivery method that provides rights on demand and sits atop the cloud computing architectural stack. In general, SaaS providers provide supplementary cloud computing layers, enabling them to host client data and adapt software in response to user feedback. This drastically reduces the out-of-pocket cost of launching a brand-new IT system. Customers aren't expected to pay for or keep up any kind of site infrastructure. A quick internet connection is all they need to get to their applications quickly. The infrastructure and applications used by SaaS companies are shared among their clientele. This strategy may eliminate the need for new equipment and is definitely more energy efficient than deploying several copies of software on different infrastructures. The more reliable the forecast, the more cash may be saved by reducing energy use. Since most software as a service (SaaS) businesses sell applications that are stored in data centers or run on the infrastructure supplied by IaaS providers, It is essential for these companies to model and evaluate the energy efficiency of the conception, implementation, and deployment of their programme. The cloud infrastructure that the SaaS provider uses is intended to be both kind to the environment and simple to use. It is especially important in the realm of social networking and gaming platforms to educate individuals about the effects their activities have on the environment. SaaS firms may also provide environmentally friendly software services that are housed in data centers that produce less trash and have fewer copies.[8]

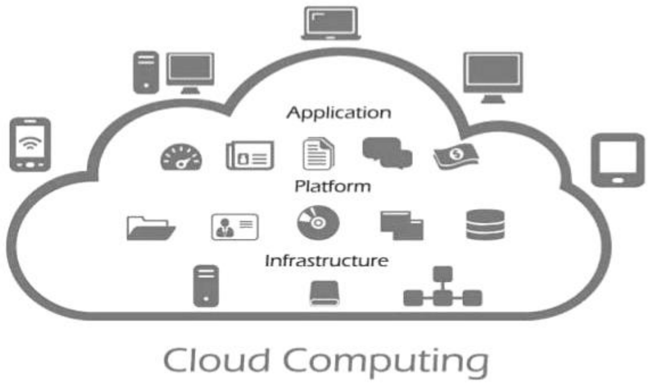

FIG 2. GREEN CLOUD COMPUTING

CONCLUSION

The use of the green algorithm led to a considerable reduction in pollutants and lower consumption of energy. The use of the VM consolidation approach results in a reduction in the amount of energy consumed as well as the number of migrations in a cloud. Consequently, it has been shown that the use of energy by virtual data servers is thirty percent lower than that of real servers. In comparison to the system in which DVFS was deployed but virtual machines were not consolidated, the power consumption of the decentralized method for resource allocation was reduced by 66%. Energy consumption and cloud migrations may both be decreased by using a VM consolidation strategy. And focus on the Weight Round instead. Robin Scheduling Algorithm allows cloud computing to provide a decision-making optimization problem. With the help of pay-as-you-go pricing and self-service alternatives, companies may encourage more efficient work habits and boost life-cycle management. When many companies or departments share a same resource, it's called multitenancy. Efficiency-boosting software that automates tasks, maximizes consolidation, and optimizes use. As these proportions improve, the amount of physical infrastructure that can be eliminated does, too, allowing server virtualization to save even more money and resources.

REFERENCES

[1] C. Kachris and I. Tomkos, “Energy-efficient optical interconnects in cloud computing infrastructures,” Sustainable Practices, pp. 463–478. doi:10.4018/978-1-4666-4852-4.ch026

[2] R. Jansen and P. R. Brenner, “Energy Efficient Virtual Machine Allocation in the cloud,” 2011 International Green Computing Conference and Workshops, 2011. doi:10.1109/igcc.2011.6008550

[3] H. Mouftah and B. Kantarci, “Energy-efficient cloud computing: A green migration of traditional it,” Handbook of Green Information and Communication Systems, pp. 295–330, 2013. doi:10.1016/b978-0-12-415844-3.00011-5

[4] A. Katal, “Energy Efficient Virtualization and consolidation in Mobile Cloud Computing,” Green Mobile Cloud Computing, pp. 49–69, 2022. doi:10.1007/978-3-031-08038-8_3

[5] M. Kaur and P. Singh, “Energy efficient green cloud: Underlying structure,” 2013 International Conference on Energy Efficient Technologies for Sustainability, 2013. doi:10.1109/iceets.2013.6533383

[6] P. Priya and G. Kaur, “Green Cloud Computing Based Future Communication Systems,” Green Communication Technologies for Future Networks, pp. 193–213, 2022. doi:10.1201/9781003264477-11

[7] A. Jain, M. Mishra, S. K. Peddoju, and N. Jain, “Energy efficient computing- green cloud computing,” 2013 International Conference on Energy Efficient Technologies for Sustainability, 2013. doi:10.1109/iceets.2013.6533519

[8] Geetanjali and S. J. Quraishi, “Energy savings using Green Cloud Computing,” 2022 Third International Conference on Intelligent Computing Instrumentation and Control Technologies (ICICICT), 2022. doi:10.1109/icicict54557.2022.9917654