

Fixed-cost Virtual Private Server using KVM Virtualization

Piyush Dhall¹, Parkhi Malhotra¹, Ashish Sharma², Garima Gupta², Sandeep Tayal², Yogesh Sharma²

¹Student, Department of Computer Science Engineering, Maharaja Agrasen Institute of Technology

²Assistant Professor, Department of Computer Science Engineering, Maharaja Agrasen Institute of Technology

Abstract - In this paper, we propose a solution for making bandwidth usage more feasible in cloud computing for low throughput requirements. Our solution involves using a fixed-cost virtual private server (VPS) that utilizes KVM virtualization. We conducted user interviews and alpha testing with our proposed solution, which received positive feedback. We propose that the initial bandwidth transfer rate of the server be set at 1Gbps, and after a subsequent amount of data transfer, such as 2TB for the basic plan, the transfer rate is rate-limited to a fixed number, such as 50Mbps. This allows for cost savings compared to current cloud providers, such as AWS EC2, AWS LightSail, DigitalOcean, Linode, and Vultr, which charge a high price on a per GB basis for additional data transfer. We demonstrate the feasibility and cost-effectiveness of this solution through analysis and comparison with other available solutions.

Key Words: Fixed Cost VPS, Virtualization, Limited Bandwidth but Unmetered Transfer, 16+ TB at 50Mbps, Cloud Computing, Bandwidth pricing, Cloud computing cost optimization, Data transfer fees, KVM virtualization, Low throughput requirements, Virtual private server (VPS)

1. Introduction

The rise of the internet and the increasing demand for online presence has led to a need for reliable and cost-effective solutions for hosting websites and applications. With the COVID-19 pandemic driving even more businesses and individuals online, this need has become more pressing than ever before^[1]. While public cloud providers like AWS and GCP offer elastic solutions that cater to the needs of small and medium-sized enterprises (SMEs) and large enterprises, these options may not be suitable for smaller projects with limited budgets and technical expertise.

In order to address this gap, we propose a cloud-based solution that is targeted towards small businesses, students, and professionals with small and non-funded projects. Our solution aims to provide a fast and cost-effective option for hosting servers that is similar to the "Jio-fied" system, which offers an alternative to overage fees for users with low data streams but high total bandwidth requirements. This will be particularly beneficial for Internet of Things (IoT) applications, where data usage may be low but total transfer needs across the month are high.

To reduce the technical overhead and make it easier for developers to host their websites and applications on the internet, we will offer a range of managed and unmanaged services. These will include tools and features to help manage and monitor servers and applications, as well as flexible pricing options that allow users to scale resources up or down as needed. By offering these services, we aim to provide a convenient and cost-effective solution that meets the needs of small businesses, students, and professionals without sacrificing business goals.

Our solution will be a welcome alternative to traditional options like procuring and managing dedicated physical servers, which can be expensive and time-consuming. In addition to the upfront costs of purchasing the servers, there are also ongoing expenses related to maintaining and monitoring them, as well as the risk of downtime due to technical issues. With our proposed cloud-based solution, these costs and risks can be significantly reduced, making it a more attractive option for small businesses, students, and professionals.

Overall, we believe that our proposed solution will fill a critical need in the market and provide a valuable resource for those looking to host their websites and applications on the internet. By offering a fast and cost-effective option that is tailored to the needs of small businesses, students, and professionals, we hope to make it easier for developers to bring their projects online and reach a wider audience.

1.1 Current Offerings in the IaaS Service Model

In the current IaaS service model, cloud providers assign a Virtual Private Server to the customer, providing them with fixed disk space, CPU as well as Memory (RAM). However, the bandwidth remains a flexible option.

To see and compare the VPS options available in the industry, let's keep the base configuration as 1vCPU, 1 GB RAM, 20GB Disk and 1TB per month transfer. We will then calculate for the scenarios where there is a need for a higher amount of transfer. The pricing models found for different service providers are mentioned below:

1.1.1 AWS EC2

Amazon Web Services Elastic Compute Cloud, referred to as AWS EC2, is their base IaaS offering, where they provide virtual cloud servers. AWS, as reported in our User Feedback Survey (section 2.2), is complex for a beginner to get familiar with, for the base configuration, if we calculate the price on AWS Price Calculator^[2], it turned out to be more than \$100 a month.

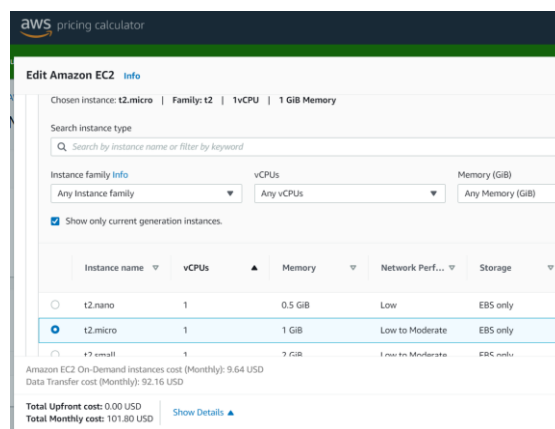


Fig 1.1 AWS EC2 pricing

1.1.2 AWS LightSail

While AWS EC2 might be complicated for beginners and people with low requirements, AWS also offers a simpler VPS offering, named AWS LightSail^[3]. It is designed to be easy to use for developers, making it a good choice for those who are new to using VPS services. Lightsail can be more cost-effective than other VPS solutions, especially for those who only need a small number of resources.

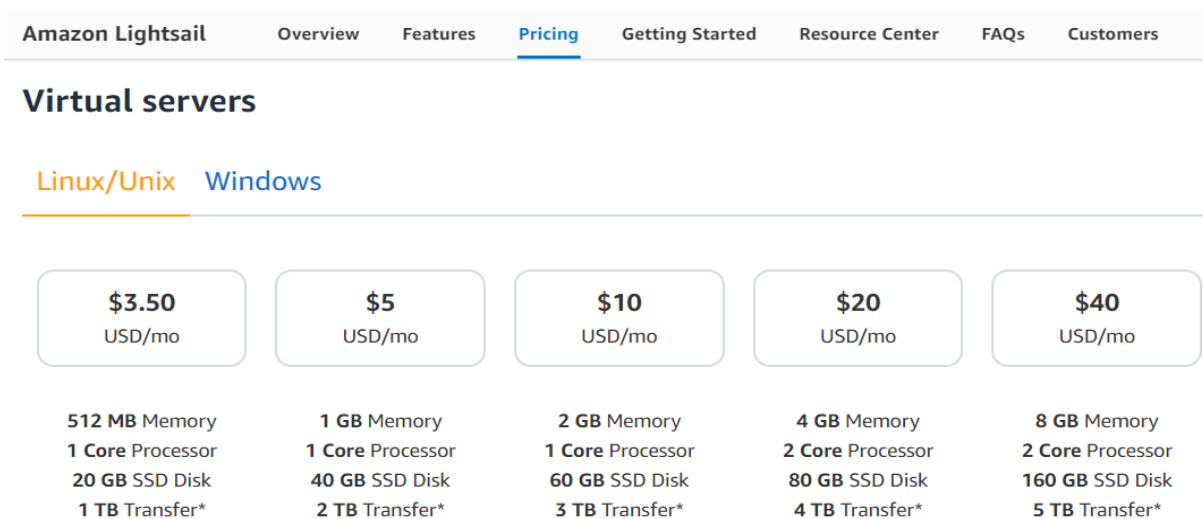


Fig 1.2 AWS LightSail pricing

For our base configuration, Lightsail offers a plan worth \$5 dollars a month which offers 1 Core processor, 1GB RAM and 2 TB transfer. This seems promising, however, when the user exceeds that base amount of transfer, the overage pricing can be a huge factor in calculating the total expenses.

Overage - Unique to each region

	US East (N. Virginia)	US East (Ohio)	US West (Oregon)	Canada (Central)	Europe (Frankfurt)	Europe (Ireland)	Europe (London)	Europe (Paris)	Europe (Stockholm)	Asia Pacific (Mumbai)	Asia Pacific (Singapore)	Asia Pacific (Sydney)	Asia Pacific (Tokyo)	Asia Pacific (Seoul)
Storage	\$0.023 USD/GB	\$0.023 USD/GB	\$0.023 USD/GB	\$0.025 USD/GB	\$0.025 USD/GB	\$0.023 USD/GB	\$0.024 USD/GB	\$0.024 USD/GB	\$0.023 USD/GB	\$0.025 USD/GB	\$0.025 USD/GB	\$0.025 USD/GB	\$0.025 USD/GB	\$0.025 USD/GB
Transfer	\$0.09 USD/GB	\$0.09 USD/GB	\$0.09 USD/GB	\$0.09 USD/GB	\$0.09 USD/GB	\$0.09 USD/GB	\$0.09 USD/GB	\$0.09 USD/GB	\$0.09 USD/GB	\$0.109 USD/GB	\$0.12 USD/GB	\$0.114 USD/GB	\$0.114 USD/GB	\$0.126 USD/GB

Fig 1.3 AWS LightSail transfer overage pricing

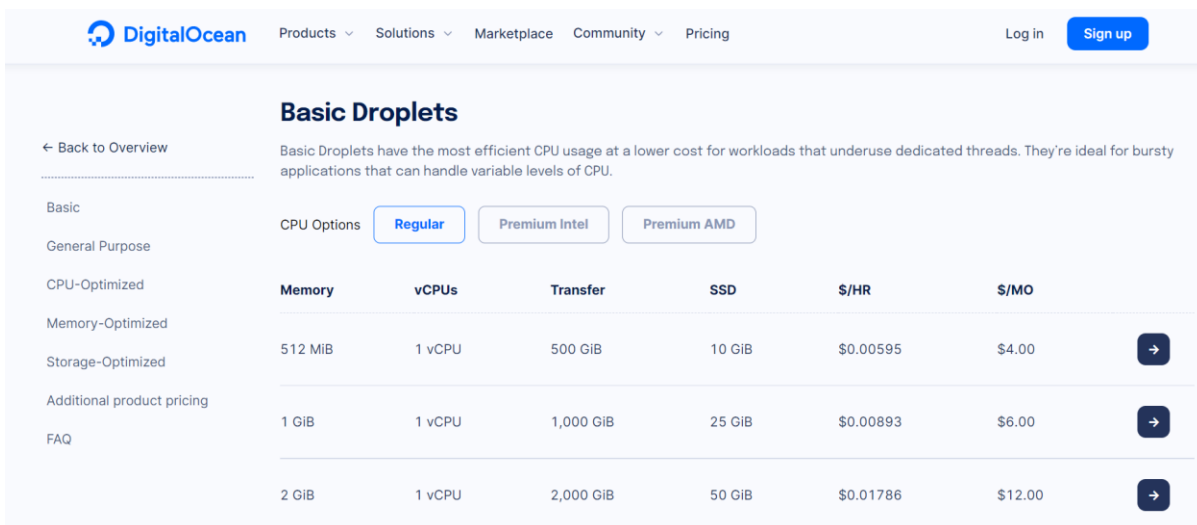
Considering the overage pricing of LightSail, on average it will charge a user 9 cents per GB, which is calculated for a 3 TB overage, the overage counts to be 270 dollars:

since 3TB = 3000 GB

3000 GB x 9 cents = 27000 cents

27000 / 100 = \$270 Dollars

1.1.3 DigitalOcean



The screenshot shows the DigitalOcean website interface for 'Basic Droplets'. It includes navigation links like 'Products', 'Solutions', 'Marketplace', 'Community', and 'Pricing'. The main content area is titled 'Basic Droplets' and describes them as having efficient CPU usage. Below this, there are tabs for 'CPU Options' (Regular, Premium Intel, Premium AMD) and a table of pricing options. The table lists configurations for Memory, vCPUs, Transfer, SSD, \$/HR, and \$/MO, with buttons to view details for each configuration.

Fig 1.4 DigitalOcean's Droplet Pricing

DigitalOcean is a cloud computing platform that offers virtual private servers (VPS) known as "Droplets". Their pricing model is based on a pay-as-you-go model^[4], which can be cost-effective for those who only need a small number of resources. Although their tagline says, "Simple, Predictable pricing", It's not "fixed". For our base configuration, the plan amount comes to be \$6 dollars per month, however, it is not fixed and for bandwidth overage, it will be increased further.

Industry-leading bandwidth pricing

Keep your bandwidth costs low with generous transfer quotas and outbound overage at just \$0.01/GiB.

[Go to bandwidth calculator →](#)

Fig 1.5 DigitalOcean's Transfer Overage pricing

Considering the overage pricing of DigitalOcean, it will cost a user 1 cent per GB, which is calculated for a 3 TB overage, the overage counts to be \$30 dollars :

since 3TB = 3000 GB

3000 GB x 1 cent = 3000 cents

3000 / 100 = \$30 Dollars

Hence, It is important to monitor your data transfer usage to avoid exceeding your limit and incurring additional charges. You can view your data transfer usage in the DigitalOcean control panel.

If you exceed the data transfer limit defined in your plan, you will be charged additional fees depending on the region and the amount of data transferred, the rate for additional data transfer is \$0.01 per GB.

Which, if calculated for a 5 TB overage, turns out to be \$50 USD additional.

1.1.4 Other Misc. Providers

Other than the providers mentioned above, there are also some VPS providers which also provide servers based on a Pay-As-You-Go basis, however, when the transfer exceeds the base plan limits, the user might get a bill shock at the end of the month when their invoice comes.

The providers having similar models with the same pain points are:

- Linode^[5] \$5 per month for the base plan than \$0.01 for each GB outbound transfer. (\$10 dollars per TB excess usage)
- Vultr^[6] \$5 per month for the base plan than \$0.01 per GB for bandwidth used in excess of your quota. (\$10 dollars per TB excess usage)

In addition to the above Provider's study, we also conducted User Research, interviewing them on their usage, as someone who is beginning into the Cloud field while coming from a Developer background, what issues they face in the current cloud infrastructure.

2. User Research

Primarily, We wanted to understand the core of the problem we were trying to solve and also figure out what the needs of the people would be.

To do so, we conducted user surveys and user interviews with a certain identified set of participants. The goal was to understand the goals and frustrations that people experience when they try to host their applications.

2.1 Survey and Interviews conducted with a certain identified audience.

We prepared an interview transcript to conduct the interviews in a smooth and efficient manner. Here is the interview transcript:

Date:

Author: Parkhi Malhotra, Piyush Dhall

Participant Name:

Email:

Time/ Date:

Location:

Hello, Thank you for participating!

Is it okay with you if we do an audio or a video recording of this interview? This is so that we can go back to the video or audio to make sure we got everything right. The tapes will only be used internally and will never be shared anywhere with anyone.

We will not be evaluating you, but we want to learn from you about what you think about **Bandwidth offerings in Cloud Computing**, what works for you and what doesn't, and with this hope, we can make better products for people like you.

There are no right or wrong answers to the questions I ask in this study. We are simply interested in knowing about your experience and your thoughts.

This would just take 2-3 minutes to complete. Thank you for your help!

1. Are you a student or a working professional?
2. If working professional, how many years of IT experience do you have?
3. How would you describe your field/ role? (Example: Software Developer, System Engineer, Product Manager, CTO, CIO, CEO, Cloud Architect)
4. Number of employees at your company, if any?
5. Which industry does your organization cater to?
6. What tech stacks do you/ your organization generally use?
7. Does your organization use On-premises Deployment?
8. Does your organization use Cloud Deployment?
9. What type of Cloud Service Model does your organization use?
10. Which infrastructure, if any, do you/ your organization generally use as a service provider?
11. Which platform, if any, do you/ your organization generally use as a service provider?
12. Where do you prefer hosting your applications?
13. How would you describe the design of your application?
14. What is the traffic trend of your website or application?
15. How much traffic do you think your servers utilize on a daily basis?
16. How much traffic do you think your servers utilize on a monthly basis?
17. Where are the majority of your users located?
18. What time of the day do you receive most of the traffic? (Note: IST is Indian Standard Time, GMT +5:30)

19. Do you use cloud servers to host projects for your organization?
20. Do you use cloud servers to host your personal projects?
21. Depending on your organization's industry and service type, do you need to comply with any security regulations such as HIPAA, GDPR, PCI DSS or SOX?
22. Do you face any difficulties in moving your application from development (localhost) to the cloud? If yes, please specify.
23. Are you satisfied with the support from your current provider? If not, which of your needs are not being met, please specify.
24. How much do you think you have to pay for bandwidth usage other than the computing and storage costs?
25. Would you use a service in which there will be fixed lower Bandwidth but Unlimited Transfer? (Unlimited @ 50Mbps instead of 1TB @ 1Gbps)
26. Which of your use cases/applications do you think will fit into a similar fixed Bandwidth but Unlimited Traffic offering?

2.2 Some quotes from participants

"I actually do not like AWS because it is way too complex to achieve simple operations."

"For AWS scaling is a little tricky, let's suppose you have a node, you want to increase in size, so for that, you have to increase its RAM usage so for that there are no custom options, so you can scale only a certain way that will work for you but it doesn't work always for you."

"For AWS, it is complex and has too many features and most of which we do not care about"

"Kubernetes sometimes becomes too complex, sometimes your organization just don't need Kubernetes."

"I have not personally used Digital Ocean but I think the pricing is not fair and some offerings are not meant for everyone. It does not make sense for startups."

"AWS is great for big organizations with a lot of resources, but for small businesses and startups, it might be overkill. The cost can add up quickly if you're not careful."

"Kubernetes is powerful, but it can be difficult to set up and manage. It's not always the best choice for small teams or projects."



Fig. 2.1 Interview conducted with Mr Arun Yadav (Cisco TAC Engineer)^[7]

We collected information on the daily network trends and the applications that this model is best suited for. Some applications might not work with this strategy since they are particularly latency-sensitive, and limiting any form of traffic

might harm them (for example NetApp storage servers). For small businesses, freelancers and student developers, whose needs are not too high, this technique, however, can work. Additionally, consumers won't have any problems in the event of an unanticipated overage.

2.3 User Personas

2.3.1 Rashmi

Rashmi is a 22-year-old computer science student at a university. She is passionate about technology and enjoys building websites and applications as a hobby. Rashmi recently developed a website or application as part of a school project and is now looking to host it on the cloud so that it is accessible to the public. She is on a tight budget and is hesitant to host her project on the cloud because she is worried about the fluctuating costs of cloud hosting services. Rashmi wants to find a cloud hosting solution that is affordable and fits within her budget but is frustrated by the lack of transparency and predictability in cloud hosting costs.

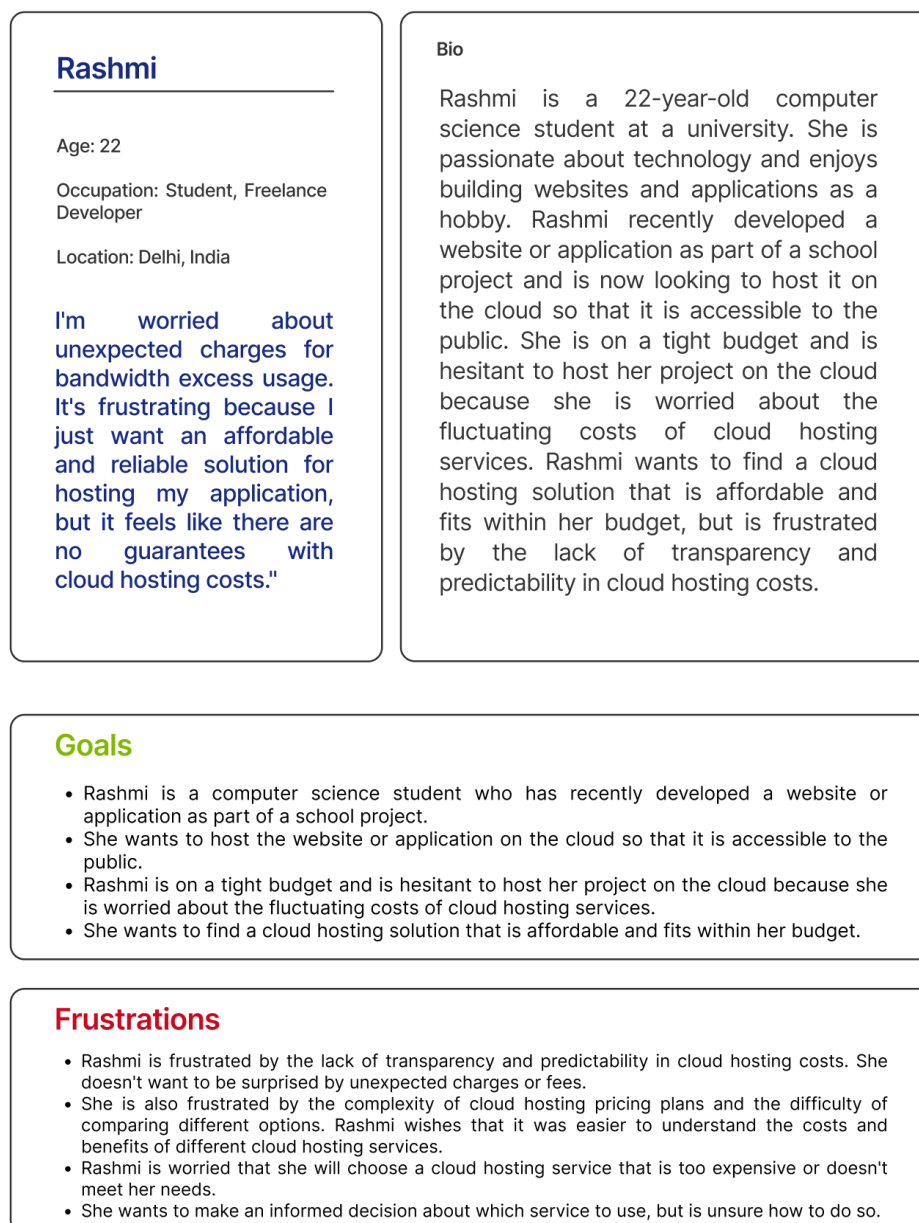


Fig 2.2 Rashmi User Persona

2.3.2 Rishabh

Small business owner Rishabh just created a website or application to advertise his goods or services. He is hesitant to use the cloud to host the website or application because he is concerned about the variable pricing of cloud hosting providers. Because the charges are opaque and unpredictable, Rishabh is having problems finding a cloud hosting solution that falls within his limited budget. Despite his best efforts, Rishabh is having trouble deciding which cloud hosting provider to utilize because of these problems.

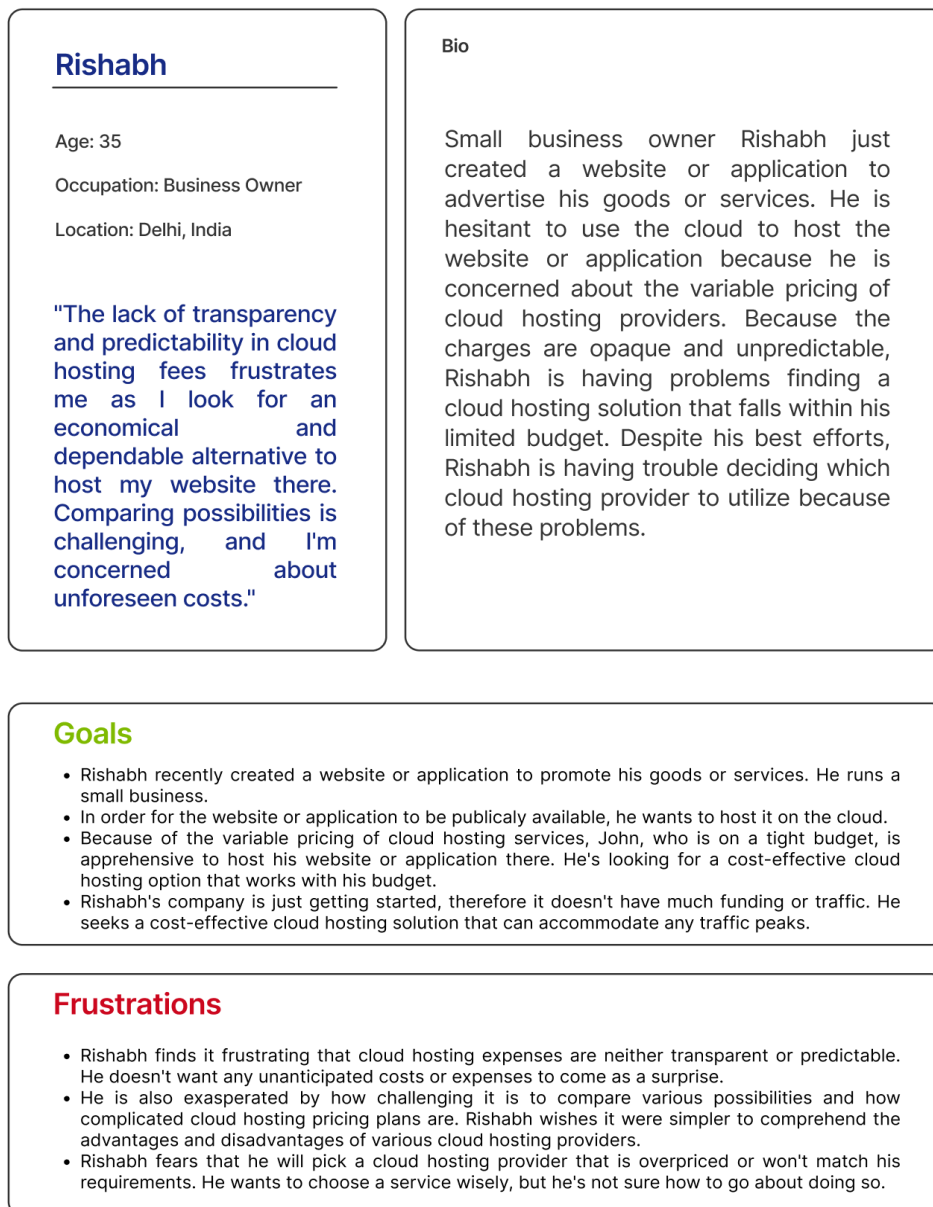


Fig 2.3 Rishabh User Persona

3. Implementation

Dedicated Servers can be rented with providers like Hetzner^[8] or can be set up at a colocation provider like Serverius^[9], where the bandwidth limit is either Unmetered capped at 1Gbps, or high enough for our needs.

On top of them, a Virtualization stack can be installed in order to facilitate the VPS provisioning. Since the dedicated servers come with bulk bandwidth. Running multiple VPSs under the network line is feasible.

Before we can host VPSs on a public network, we need IP addresses in the public internet space for the Virtual Machines to be able to connect to the internet and act as a server. For that, we need public IPv4 and IPv6 addresses. IPv4 is relatively harder to procure due to the limited availability of IPv4 in the world^[10], still, it can be procured at around \$2-3 dollars a month per IP from most providers^[11].

IPv6 on the other hand is available freely and even in bulk as most providers allocate a /64 block to each dedicated server, which can be enough even for the whole Organization to run their network.

For the Virtualization stack, KVM can be leveraged as it is resource efficient in terms of computing power as well as efficient in management. A similar approach can be followed in other Virtualization stacks though.

To assign IP addresses to the VMs after they are created, Cloud-init as well as DHCP is used for the IP address assignments. Cloud-init is a set of tools and configuration scripts that are used to automate the initial setup and configuration of cloud-based servers, such as those used in Infrastructure as a Service (IaaS) environments. Cloud-init is commonly used to perform tasks such as setting the hostname, installing packages, writing files, and configuring networking on cloud servers when they are first launched.

Although Cloud-init itself can be utilized for IP assignments, however, while implementing the said solution, we faced difficulties in some network configurations of the data centres where the IP Gateway is outside the IP Subnet, such as OVHCloud. In those cases, static routes need to be defined along with the IP, Subnet and Gateway. Since we couldn't implement the static routes directly in the cloud-init, we leveraged an internal DHCP server to run on the Bridge which is responsible for making the VMs connect with the Network line.

For installing the Operating Systems, a set of templates can be made for installing some common Operating Systems. Cloud-init can also be used for that. As it is typically included as part of the cloud images that are provided by cloud providers and is run automatically when a new server is launched. The configuration of cloud-init is typically specified using a configuration file in a standard format, such as YAML or JSON, which can be passed to the cloud provider when the server is launched.

Once the VPS is live and serving, we need a way to monitor its bandwidth utilization over the month. For that RRD is used along with the KVM to fetch the metrics. RRD (Round Robin Database) is a system for storing and analyzing time-series data, such as network bandwidth usage.^[12] RRD is often used to monitor the bandwidth of a Virtual Private Server (VPS) running on a Kernel-based Virtual Machine (KVM) because it provides a simple and efficient way to store and visualize the data.

The RRD store can then be periodically monitored to see how much transfer a VM has done over the month. Once the transfer crosses a certain limit according to the plans (like 1TB, 2TB etc.), the VM's bandwidth is throttled to a minimum number which is set to 50Mbps for now. But can be configured while implementing this procedure by the network administrator, according to the availability and sustainability inside the data centre due to regional factors.

3.1 Rate Limiting the bandwidth and their results

The rate limit is applied by setting the KVM machine's parameters by the following command. Note the "rate=12" in the CLI, which is for setting a rate limit of around 100 Mbps.

```
root@n11:~# qm set 100 -net0 virtio=8E:29:6A:47:E3:3B,bridge=vibr0,rate=12
update VM 100: -net0 virtio=8E:29:6A:47:E3:3B,bridge=vibr0,rate=12
```

Fig 3.1 Applying a Throttle of 100Mbps

3.2 Disengaging Throttle back to 1Gbps

The rate limit is disabled by setting the KVM machine's parameters by the same command. Note the "rate=120" in the CLI, which is resetting the rate limit back to 1Gbps.

```
root@n11:~# qm set 100 -net0 virtio=8E:29:6A:47:E3:3B,bridge=vibr0,rate=120
update VM 100: -net0 virtio=8E:29:6A:47:E3:3B,bridge=vibr0,rate=120
```

Fig 3.2 Resetting the throttle back to 1Gbps

3.3 Speedtests

For conducting speed tests, Ookla’s Speedtest CLI [13] is used. It returns the results in the console as well as a URL image. As have been represented below:

```
OpenSSH SSH client
root@jen:~# root@jen:~# speedtest

Speedtest by Ookla

Server: Claranet Benelux B.V. - Amsterdam (id: 30847)
ISP: Serverius Holding B.V.
Idle Latency: 1.87 ms (jitter: 0.12ms, low: 1.78ms, high: 1.96ms)
Download: 959.27 Mbps (data used: 433.6 MB)
13.46 ms (jitter: 0.43ms, low: 1.78ms, high: 16.66ms)
Upload: 1003.53 Mbps (data used: 453.4 MB)
1.73 ms (jitter: 0.11ms, low: 1.59ms, high: 2.23ms)
Packet Loss: 0.0%
Result URL: https://www.speedtest.net/result/c/3c2ef242-4904-44f6-a3da-a6a227b7a767
```

Fig 3.3 Conducting a Speedtest by Ookla’s CLI

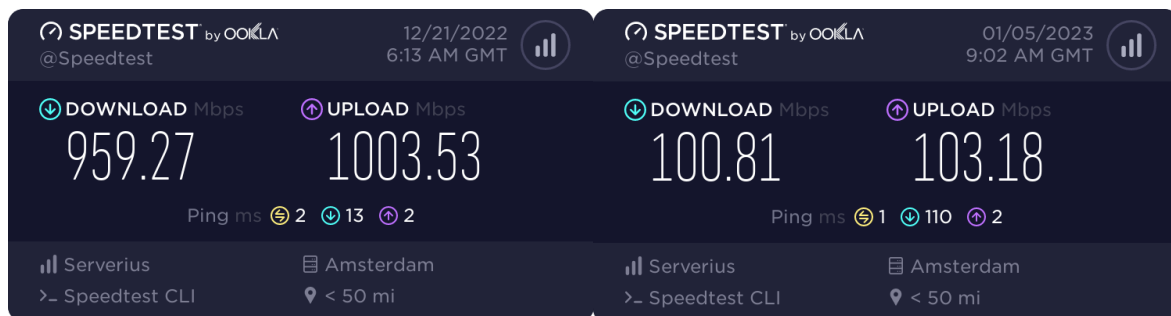


Fig 3.4 Normal Speed (set to 1Gbps) vs Throttled speed (100 Mbps)

3.4 Calculating what 50Mbps can give over a calendar month

1 Megabit per Second: When we talk about 1 Megabit per Second (1 Mbit/s), we’re referring to a rate of data transmission that equals 1,000,000 bits per second. A megabit is a unit of digital information that consists of 1,000² (or one million) bits. Meanwhile, a second is the standard unit of time used in the International System of Units (SI). Therefore, 1 Mbit/s is equivalent to 1,000,000 bit/s.

If we assume a monthly data cap of 1 Terabyte (1 TB), which translates to 8,000,000,000,000 bits in total (using the base unit), the average data transfer rate can be estimated at around 3086419.7531 bits per second. This is based on the assumption that a month has 30 days or 2592000 seconds (using the SI base unit). Therefore, we can say that 1 TB per month is roughly equivalent to 3086419.75308642 bit/s.

$$\begin{aligned}
 \text{50 Mbps monthly transfer} &= 50 \text{ Mbps} \times 30 \text{ Days} \\
 &= 50 \text{ Mbps} \times 720 \text{ Hours} \\
 &= 50 \text{ Mbps} \times 43,200 \text{ Minutes} \\
 &= 50 \text{ Mbps} \times 2,592,000 \text{ Seconds}
 \end{aligned}$$

- = 129,600,000 Megabits per month
- = 16,200,000 MegaBytes per month (*since 1 Byte = 8 bits*)
- = **15.6 TeraBytes per month** (*Divide by 1024 twice*)

This calculation is to give an idea of what 50Mbps sounds like!

However, in a general situation where the bandwidth is getting utilized in full, the total bandwidth will exceed this result (contrary to general norms). This is because, initially the server will have access to the full 1000 Mbit/s of bandwidth, and the rate limit will only be applied once the base limit exceeds.

3.5 The architecture of the Final Solution

A portal can be developed in a web development framework like Django or MERN stack can be created so that the user can easily create and manage their VPSes and similarly the provider admins can set the bandwidth limits explicitly if they want.

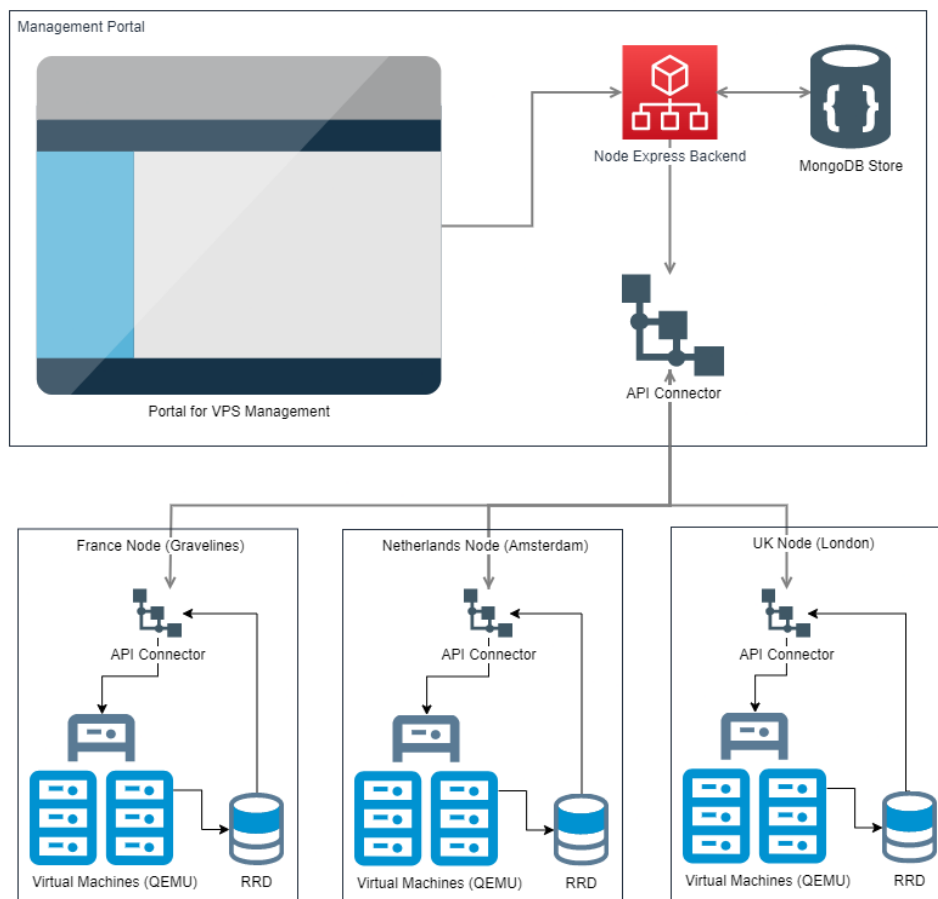


Fig 3.5 Final Architecture of the Cloud Nodes and a Management Portal

The final architecture of this cloud computing solution consists of three dedicated nodes running Proxmox, a virtualization platform based on KVM. These nodes are connected to a centralized portal, which is made using the MERN stack (MongoDB, Express, React, and Node.js). The portal serves as the interface for users to access and manage their virtual private servers (VPSs) and connects to the nodes using REST API Connectors.




Fig 3.6 Statistics of a cloud node hosted in Netherlands

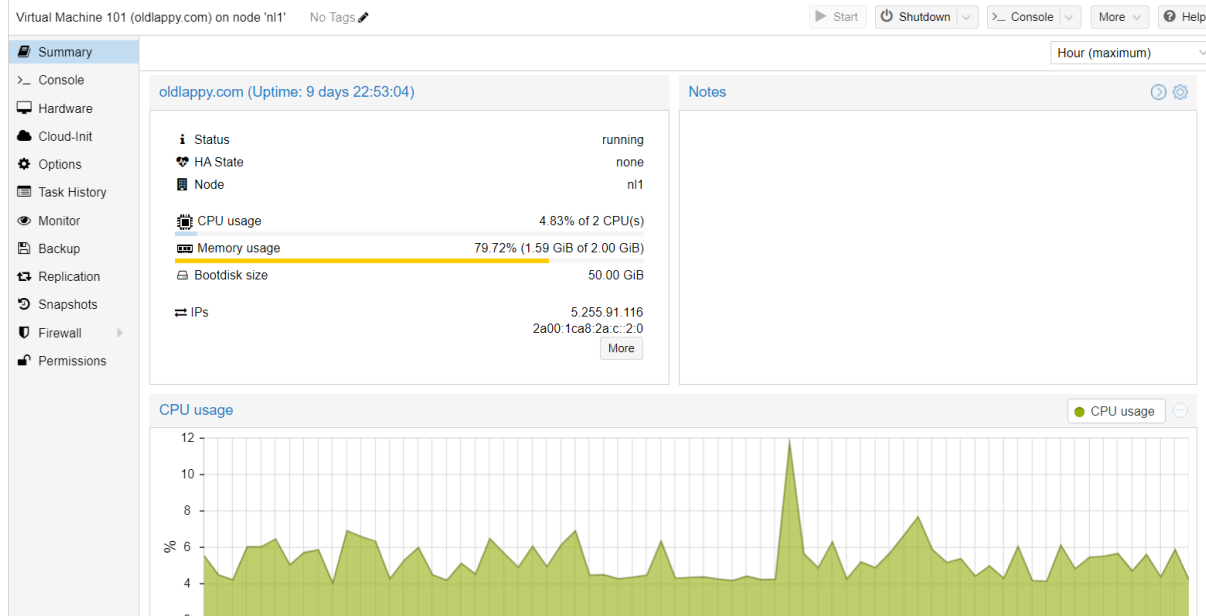


Fig 3.7 A virtual machine running on the NL1 node

The VPSs are created using the resources of the three Proxmox nodes, by cloning the specific cloud-init templates as per the user's requirements (Debian, Ubuntu etc.) and are configured to have a fixed cost and a rate-limited transfer rate after a certain amount of data transfer, as described in section 3.1, through a background check which checks the RRD for the monthly transfer. When the new calendar month starts, the limits are reset, as described in section 3.2. This allows for optimized costs as compared to other cloud providers that charge a high price on a per GB basis for additional data transfer.




Fig 3.8 Network statistics of a Virtual Machine running on a cloud node

Users can access the portal and create, configure, and delete their VPSs as needed. The portal also provides real-time monitoring and usage statistics for each VPS, as well as alerts for any issues or thresholds that are exceeded.

4. User Validation/ Alpha Testing

Based on the user research conducted, it has been determined that this server is beneficial for IoT applications where data streams are low, but overall bandwidth requirements are high. The server currently caters only to the IaaS model, but future iterations may include PaaS models such as Docker container runtimes and managed Kubernetes pods on a similar fixed-cost model. In order to host VPSs on a public network, public IPv4 and IPv6 addresses are necessary. IPv4 is harder to obtain due to limited availability but can be purchased for around \$2-3 per month per IP. IPv6, on the other hand, is readily available and often provided in bulk by providers. The KVM virtualization stack was chosen for its resource efficiency and ease of management. Cloud-init and DHCP were used to assign IP addresses to VMs, with an internal DHCP server used to connect VMs to the network in certain configurations where the IP gateway was outside the IP subnet. Templates were used to install common operating systems, and cloud-init was used to configure them. RRD was used to monitor bandwidth utilization over the course of a month, and VMs with transfers exceeding certain limits were throttled to a minimum of 50Mbps. User feedback has indicated that this server is reliable and easy to use, with helpful customer support.

4.1 Alpha Testing Feedback

- "I use the platform to host personal projects and also to host VPN."
- "It works pretty well and has the speed that I want"
- "I am not looking for high speeds, it's not in my use case, even in the projects in a start-up they don't need much bandwidth initially."
- "Bandwidth capacity is not the problem."
- "Since it is unlimited I don't have to care about how much capacity I have used."
- "Due to the nature of my current project, I benefit more from having a lower speed and unlimited traffic than the other way around."

5. Conclusion

Based on the user research conducted, the server will be beneficial for IoT applications, where the data is in low streams, however, the bandwidth requirement in total is more. Currently, this server caters only to the IaaS. We can switch to PaaS models in the next iterations, allowing Docker container runtimes and managed Kubernetes Pods, on a similar fixed-cost model.

This cloud model is designed to address the needs of small businesses, students, and professionals who are working on small or non-funded projects and require a server solution for them or their IoT applications. IoT applications often have low data streams but require a high total bandwidth. The current model offers infrastructure as a service (IaaS) and in the future, the same can be expanded to offer platform as a service (PaaS) options such as the ability to run Docker containers and managed Kubernetes pods. This will allow users to take advantage of the fixed-cost model while still being able to easily deploy and manage their applications.

In addition to its usefulness for IoT applications, this model also aims to provide an alternative to traditional cloud solutions for users who may be hesitant to use them due to their non-fixed costs. It offers a simple and cost-effective solution for professionals and students who need a reliable server to test, run applications, and deploy on the internet.

To manage the server and assign IP addresses to the VPSs, a virtualization stack like KVM is used in conjunction with tools like Cloud-init and DHCP. Cloud-init automates the initial setup and configuration of cloud-based servers, while DHCP allows for the automatic assignment of IP addresses to client devices on a network. In some cases, static routes may also need to be defined in order to properly connect the VPSs to the network. To monitor the bandwidth usage of the VPSs, a system like RRD (Round Robin Database) can be used to store and analyze time-series data such as network bandwidth usage. If a VPS exceeds a certain data transfer limit, its bandwidth can be throttled to a minimum amount in order to prevent overuse and ensure fair usage for all VPSs on the network.

6. Future Scope

There are several potential avenues for future development. Some possible areas for exploration include

Integration with additional tools and technologies: The solution could be enhanced by integrating with other tools and technologies that are commonly used in the development and management of websites and applications. This could include integration with version control systems, containerization technologies like Docker, and frameworks for automating deployment and scaling of applications (like 1 click connections with VSCode, 1 click deployments like WordPress, LAMP stack etc.)

Improved security and compliance: Ensuring the security and compliance of hosted applications and data is always a top priority. The solution could be further developed to include additional security measures, such as non-password authentications, Key Based authentications, Single Sign On (SSOs) as well as compliance with industry regulations and standards.

Enhanced monitoring and performance optimization: To ensure the optimal performance of hosted applications, the solution could be enhanced with additional monitoring and performance optimization features. This could include real-time monitoring of server resources and application performance, as well as tools for identifying and addressing bottlenecks and other issues.

Expansion to additional regions: In order to make the solution more accessible to a wider global audience, it could be expanded to cover additional regions around the world. This would involve establishing new data centers and partnerships with local providers to ensure low latency and high performance for users in these regions.

Diversification of services offered: The solution could be further developed to offer a wider range of hosting and related services beyond just server provisioning. This could include hosting of specific types of applications or technologies (such as e-commerce platforms or machine learning models), as well as additional tools and resources for developers, such as documentation, training materials and support forum.

References

1. Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis. 2020. The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic. In Internet Measurement Conference (IMC '20), October 27–29, 2020, Virtual Event, USA. ACM, New York, NY, USA. <https://doi.org/10.1145/3419394.3423658>
2. *AWS Pricing Calculator*, <https://calculator.aws/#/estimate>. Accessed 26 March 2023.
3. "VPS, web hosting pricing—Amazon Lightsail—Amazon Web Services." *Amazon AWS*, <https://aws.amazon.com/lightsail/pricing/>. Accessed 26 March 2023.
4. "Droplet Pricing." *DigitalOcean*, <https://www.digitalocean.com/pricing/droplets>. Accessed 26 March 2023.
5. Mitchel, Justin. "Cloud Computing Services." *Linode*, <https://www.linode.com/pricing/>. Accessed 26 March 2023.
6. "High Performance, High Frequency, Bare Metal, Affordable Cloud Computing." *Vultr.com*, <https://www.vultr.com/pricing/>. Accessed 26 March 2023.
7. Interview Recording with Mr Arun Yadav (Cisco TAC Engineer): <https://www.youtube.com/watch?v=OwEULmAEwJE>, Conducted on 21st December 2022.
8. "Dedicated Root Server Hosting." *Dedicated Root Server Hosting*, <https://www.hetzner.com/dedicated-rootserver>. Accessed 26 March 2023
9. "Colocation data center, server colocation, twin colocation." *Serverius*, <https://serverius.net/colocation/>. Accessed 26 March 2023.

10. Babatunde, Olabenjo & Al-Debagy, Omar. (2014). A Comparative Review Of Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6). International Journal of Computer Trends and Technology. 13. 10.14445/22312803/IJCTT-V13P103.
11. Providers repeatedly increasing the IPv4 pricing due to limited space left - <https://blog.ovhcloud.com/additional-ipv4-new-pricing/>
12. Singh, M., Mewada, H., Tahilyani, M., Malviya, J., Sharma, R. & Shrivastava, S. S. (2022). RRDTTool: A Round Robin Database for Network Monitoring. Journal of Computer Science, 18(8), 770-776. <https://doi.org/10.3844/jcssp.2022.770.776>