VULTR

The ultimate guide to multicloud strategies with Vultr

Cloud computing has dramatically improved how users interact with software applications over the past decade. It's changed how applications are hosted and used, how businesses store and access data, and even how companies operate. Organizations no longer store all their data and software on physical servers, buying more hardware whenever their businesses scale up. Instead, cloud computing allows organizations to scale flexibly by accessing various services online and on demand.

The great thing about cloud computing is that it benefits companies of all sizes. Companies can set up servers on their premises and create a private cloud environment that they manage, update, and maintain themselves or use a public cloud environment. This is done through a cloud service provider or vendor which provides IT services over the internet.

Cloud vendors offer services such as application development platforms, servers, storage, and fully managed virtual machines (VMs). These services are available as infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS). Companies that only want to depend partially on the public cloud or cannot move all their workloads to the public cloud use a hybrid cloud environment. This is comprised of both private and public clouds.

The many available cloud providers cater to a myriad of different business needs. Each has its strengths and weaknesses. There's no "one size fits all" vendor for everything a business needs. Consequently, you must consider many factors when choosing a cloud provider that best suits your needs.

With the many options and advancements in infrastructure as code (IaC), organizations can choose from different cloud providers and run their workloads in various cloud environments as needed. This is called a multicloud approach. A multicloud environment uses cloud services from more than one cloud vendor. This prevents vendor lock-in and increases the organization's options in the cloud.

Vultr is a leading independent cloud simplifying cloud infrastructure deployment with its advanced cloud platform. This whitepaper will provide everything you must know to run a multicloud approach using Vultr.

Vultr makes multicloud easy and affordable

Cloud computing gives businesses access to nearly infinite storage and computing capabilities that allow them to scale without worrying about managing hardware. However, all this flexibility comes at a cost. Cloud expenditure is proving a challenge as companies in the IT industry are struggling to keep their cloud costs under control. The International Data Corporation (IDC) estimated that nearly 30% of overall IT budgets go to cloud expenditure, which is trending upward quickly. The total public cloud spending has risen from $307.7 billion in 2020 to $383.6 billion in 2021. It's estimated to exceed $1.3 trillion by 2025.

The cost of hyperscalers

Hyperscale computing refers to computing resources that can scale up and down to reasonable levels based on the demand placed on a system. Hyperscaling is about seamlessly providing computing resources – typically computing, memory, networking, and storage – at a massive scale. It relies on massively scalable server architectures and virtual networking.

Hyperscalers apply hyperscaling to data centers and the cloud to meet the fluctuating cloud demands and provide large-scale cloud infrastructure and services. This is the driving force behind technologies like video streaming, social media, cloud computing, and big data storage. Hyperscale computing is necessary to build a robust and scalable cloud or a distributed storage system. It's often associated with the infrastructure required to run large distributed sites such as Google, Facebook, Twitter, Microsoft, and Amazon. Some of these companies not only build the infrastructure for themselves but also allow their customers to use parts of it. Examples include Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP), the three most popular cloud computing hyperscalers.

Organizations hoping to upscale their workloads and reduce costs tend to look to cloud hyperscalers. However, different businesses don't enjoy the same benefits. When relying on hyperscalers, companies don't always use all the resources vendors provide. Companies must configure their environments based on the vendor's specifications. This often leads to companies paying for more resources than they need.

Furthermore, billing models tend to be complicated, with hidden charges based on cloud usage patterns. Predicting cloud costs tends to be complex for organizations that rely on these providers. This has led to the development of many tools to reduce and make sense of cloud costs.

Hyperscaler use cases

Hyperscalers are cloud providers with data centers with massive computing power, large amounts of storage, and extensive network infrastructure.

Corporations can use hyperscalers to migrate all their workloads to the cloud and deliver end-to-end solutions through the cloud. Hyperscalers support mass migrations to the cloud. They provide migration strategies and options for reconfiguring a platform for the cloud.

Hyperscalers can store and process massive amounts of data with network speeds of up to 40 gigabytes per second (Gbps), with 800 Gbps and even 1600 Gbps coming soon. These are ideal for companies that continuously produce large amounts of data and perform big data analytics, and store that data in the cloud.

Hyperscalers can support the development and testing of large-scale, resource-heavy workloads. This is due to hyperscalers' abilities to scale tremendously to meet very high demand. Using a scale-up or scale-out mechanism, automation can increase or decrease infrastructure resources to meet changing needs. Applications that experience varying amounts of traffic at certain intervals benefit significantly from infrastructure scaling.

These large-scale cloud providers are optimized for these cases because they have large data centers with significantly more servers. This provides the advantage of economies of scale. Large-scale cloud providers also emphasize certain engineering practices, such as automation, modularity, and resource disaggregation. Therefore, certain computing resources can be mixed and matched.

Other use cases

Moving workloads to these giant platforms without determining the right infrastructure is not a sound business decision. Not all workloads work well in a hyperscale environment. Hyperscalers optimize their environments for highly scalable workloads and profit by renting out their standard platforms at scale. Moving workloads carelessly could lead to underutilization of the vendor's resources and unaccounted cloud costs.

Businesses must avoid deploying low-intensity workloads onto hyperscale clouds. These workloads don't have specific computational requirements. They typically run on the default configuration of the cloud provider. These are usually typical web applications and servers, non-IT-related SaaS applications, and low-impact internal portals. Such workloads are typically static, with predictable attributes such as resource requirements, traffic demands, and downtimes.

Components essential for your application's security and well-being are usually not factored into the cost of cloud services. These are sold separately from your hosting plan and include network security and application health monitoring services for day-to-day operations. These costs can add up swiftly and unpredictably, making it very expensive to run small workloads.

Hyperscaler latency

Public cloud providers simplify deploying and scaling web applications for businesses. They offer multiple availability regions worldwide for better availability and performance of applications on a global scale. However, expectations for application performance continue to rise. Many new services like online gaming, augmented reality (AR), virtual reality (VR), and the Internet of Things (IoT) demand the lowest latency possible.

This is a challenge for traditional cloud architecture. People now connect to the cloud in more remote places than ever, further compounding this challenge. Traditionally, hyperscalers have been inherently centralized. Until recently, they have not prioritized geographic distribution. As a result, they lack proper worldwide coverage. Because their data centers are located mainly in large urban areas, remote areas are at a great disadvantage. The cost is latency.

At this point, you might be thinking about content delivery networks (CDNs). However, they are ineffective in solving this challenge as well. A CDN distributes data to a network of nodes closer to end users, where it is cached and served. This results in low latency. But the CDN can only serve cached data and has to call the centralized data center to serve new data. Therefore, the latency advantage is obsolete.

Consequently, the choice of a cloud provider for organizations with international teams can make or break how effectively they work. This depends on how close they are to the provider's data centers. Teams or individuals in remote areas are often left underserved by the company's choice of provider.

A quest for better coverage and low latency within regions has led organizations to move to smaller, regional – also known as local – cloud providers. These providers have

locations within the region they serve. As the data travels a shorter distance, there is low latency. Therefore, these providers can offer better network connectivity than hyperscalers.

Vultr: cloud hosting at the edge

Vultr is a cloud infrastructure provider with a global footprint and provides superior offerings with its advanced cloud platform. Vultr provides various products to suit the customer's different business needs.

Cloud Compute offers virtual private servers (VPSs) with shared resources. These VMs share some resources with other VMs on the same virtual network, such as CPUs. You can deploy a server and run your website on these VMs.

Optimized Cloud Compute offers VMs that provide fully dedicated resources. These aren't shared with other VMs on the network. Your options include:

- General Purpose cloud servers that provide a balance of CPU, memory, and Non-Volatile Memory Express (NVMe) storage.
- CPU Optimized servers that provide proportionally more CPU than random access memory (RAM) with a 1:2 CPU: RAM ratio. This means you get one virtual CPU (vCPU) for every 2GB of RAM.
- Memory Optimized servers that provide proportionally more RAM than CPU with a 1:8 CPU: RAM ratio.
- Storage Optimized servers that provide proportionately more NVMe storage with a 1:8 CPU: RAM ratio. The smallest storage space available is 4TB.

Bare metal provides access to dedicated physical servers with no virtualization layer. They're perfect for resource-intensive applications and workloads that require single-tenant or non-virtualized environments.

The Vultr Kubernetes Engine is a fully managed container orchestration tool for deploying and scaling containerized applications.

Block Storage is expandable storage volumes. These can be mounted directly onto your compute instances to add more storage space to your cloud server instead of upgrading your compute instance. You can add up to 40TB of additional storage with a minimum of 40GB for hard disk drive (HDD) block storage. You can also add up to 10TB of additional storage with a minimum of 10GB for NVMe storage.

Object Storage provides scalable storage on demand for your application's data. It allows storing and accessing static files, such as media files. It also supports the Amazon Simple Storage Service (Amazon S3) API for storing, retrieving, listing, and deleting objects in Amazon S3.

Vultr has been expanding worldwide. It now has cloud computing services available in 29 locations worldwide that you can spin up in under 60 seconds. When you deploy, the Vultr cloud orchestration spins up your instance in your desired data center. With the most extensive worldwide network, Vultr lets you quickly scale a low-latency infrastructure solution wherever your customers are. With Vultr, you don't get locked in by long-term contracts. The pricing is straightforward and affordable. You only pay for what you use.

Vultr use cases

One of the most significant challenges cloud computing users face is unnecessary complexity, cost, and billing unpredictability. However, Vultr believes simplicity matters. Vultr is committed to providing businesses worldwide with the best price-to-performance of any cloud computing platform.

Vultr has more data center locations worldwide than any other independent cloud provider. This provides a platform for businesses to scale their applications globally. While providing all this, Vultr maintains a single price per instance in almost every location. The reduced complexity of Vultr's cloud hosting options makes building simple, reducing surprises and risks.

Vultr is the preferred option for businesses looking for avenues around hyperscalers. When working with hyperscalers, you have access to hundreds of services and pay a premium for accessing all those services. As hyperscalers grow, they lack personalized attention, and their program rules become increasingly complex.

Vultr is optimized to fill those gaps and serve customers' needs without compromising performance. Vultr offers fundamental capabilities, such as VMs, storage options, and container orchestration. The Vultr platform is highly customizable. Affordable plans provide flexibility, simplicity, and value. This ensures developers can deploy new cloud instances of any type and configuration in any location in as little as 60 seconds. This way, developers can start building, testing, and deploying workloads as quickly as possible.

Multicloud with Vultr

When it comes to cloud vendors, there's no universal solution because every business has different needs. As a result, a multicloud approach has become a common practice among companies that rely on the cloud. Access to a wide range of cloud vendors allows businesses to choose the key features of the individual cloud platforms that best suit their business needs.

This approach also means that organizations can embrace a wider range of cloud vendors, including alternative and specialist clouds, without being limited to the small set of hyperscale cloud providers. Also, with this approach, users can now use infrastructure as code (IaC) technologies to be independent of any cloud provider. Consequently, they can take advantage of the different benefits of all cloud vendor types.

Use Vultr for low-cost deployment and management of high-performance servers. You can run all your workloads that don't require hyperscaler resources, like infinite scaling. You can still keep all critical enterprise production systems running on hyperscaler infrastructures to benefit from the services and options they provide.

Optimizing your cloud infrastructure with Vultr

Hyperscalers are necessary to run large distributed workloads on the cloud. They can achieve significant scale and offer massive scaling capabilities for distributed workloads of up to thousands of servers. Hyperscalers can accomplish this because of the large data centers that can achieve horizontal scaling. These house thousands of servers that are capable of high-performance levels.

Hyperscalers are best suited for big data and workloads that require massive scaling. However, what they provide is usually overkill for smaller workloads and can lead to unnecessary costs. Fortunately, for these kinds of workloads, there are alternatives.

For workloads running on the cloud that don't require the massive scaling capabilities of hyperscalers, Vultr offers a more cost-effective and efficient platform. Vultr offers all the core features of a cloud computing platform without all the complexity of hyperscalers. Doing so allows developers to build and run applications on the cloud faster and more efficiently. This is especially if you have workloads with varying requirements because the different cloud providers have optimized for particular workloads.

Using various providers through a multicloud approach ensures you get the best of both worlds. You can run specific workloads on the cloud provider that best suits your business needs.

When deciding to run workloads on the cloud, choose Vultr, not the hyperscalers.