



Market Insight Report Reprint

Vultr rides cloud GPU wave, widens ambitions

July 17 2023

by **Agatha Poon**

Positioned as an alternative to major cloud hyperscalers, Vultr is aiming to solidify its presence in the cloud GPU market with fractionalized offerings and large-scale GPU clusters to meet the diverse needs of customers in both developed and developing markets, and is selling directly to enterprises.

S&P Global
Market Intelligence

This report, licensed to Vultr, developed and as provided by S&P Global Market Intelligence (S&P), was published as part of S&P's syndicated market insight subscription service. It shall be owned in its entirety by S&P. This report is solely intended for use by the recipient and may not be reproduced or re-posted, in whole or in part, by the recipient without express permission from S&P.

Introduction

Independent cloud infrastructure provider Vultr has been quietly expanding its cloud footprint globally. With the opening of new cloud datacenters in Tel Aviv and Manchester, UK, the company now has a cloud footprint spanning 32 locations across the Americas; Europe, the Middle East and Africa; and Asia-Pacific. Positioned as an alternative to major cloud hyperscalers, Vultr is aiming to solidify its presence in the cloud GPU (graphics processing unit) market with fractionalized offerings and large-scale GPU clusters to meet the diverse needs of customers in both developed and developing markets. Additionally, the company has ramped up its sales and marketing resources to grow beyond the developer community.

THE TAKE

Vultr hopes to carve out a niche in the cloud services market by focusing on ease of use and price performance as a technology foundation. Its experience and expertise in providing support to the developer community certainly helps. And the timing seems right because hybrid and multicloud deployments are taking hold across the broader business market, particularly in the enterprise segment. Vultr is now eyeing opportunities in the cloud GPU market, which is expected to become increasingly democratized as competition heats up. Growing beyond the developer community, the company will compete squarely with major cloud hyperscalers and their partners, many of which focus on creating value using artificial intelligence and are pushing into key verticals.

Context

Founded in 2014 by David Aninowsky, the founder and executive chairman of The Constant Company, Vultr was Constant's flagship cloud platform designed to support the developer community from day one. Presented as completely bootstrapped and remaining private, Constant sought to establish intellectual capital with Vultr. The company has roughly 175 employees, and is expanding globally through a network of partners.

As a cloud computing platform and company brand, Vultr helps generate recurring annual revenue of more than \$125 million, according to Constant. In 2022 alone, Vultr grew at a rate of 40%. Developed prior to the use of cloud infrastructure for production workloads and machine-learning/AI initiatives, Vultr has evolved from its foundation that enables immediate deployment with on-demand elasticity and transparent pricing to a cloud platform that is optimized for HPC workloads, advanced analytics and AI workloads.

Delivering Vultr out of 32 worldwide locations, including nine locations in Asia-Pacific (Tokyo, Osaka, Seoul, Delhi NCR, Mumbai, Bengaluru, Singapore, Sydney and Melbourne), the company continues to work with global datacenter partners like Equinix Inc. and Digital Realty Trust Inc., as well as various local datacenter providers, to accelerate growth. Vultr says it has more than 250,000 active customers in 185 countries.

Technology and product

As cloud deployment becomes an integral part of organizations' IT strategies, and as hybrid and multicloud deployments gain momentum, Vultr sees itself as comparable to leading cloud platforms in terms of cloud footprint, and is uniquely positioned for customers looking to diversify financial and operational risks. On average, the company claims to be 40%-50% more cost effective than major cloud platforms, and the platform is MACH-certified.

MACH — Microservices-based, API-first, Cloud-native SaaS, and Headless — is a vendor-neutral consortium of companies that promotes open standards, interoperability and vendor choice. The alliance aims to provide education and support to business organizations looking to transition from a legacy infrastructure to a truly open and composable cloud-native infrastructure.

To that end, MACH has implemented a certification program to provide a set of standards, with four principles as technology foundations. These are composable connectivity (apps connect easily via APIs); infinite scalability (seamless upgrades without needing human intervention); ultimate extensibility (making changes on the fly without disrupting front-end user interface); and flexibility and transparency (no vendor lock-in and highly adaptable). Vendors and integrators that provide MACH-certified services follow MACH standards.

As far as product strategy is concerned, Vultr has refreshed its cloud platform with a range of high-performance VMs, including Optimized Cloud Compute, Cloud Compute, Cloud GPU and Bare Metal. Optimized Cloud Compute is best suited for customers that want to run their workloads on dedicated AMD vCPUs instead of the shared vCPUs delivered by typical cloud compute instances. Customers can choose from multiple server types — general purpose, CPU optimized, memory optimized and storage optimized — to ensure the right balance of resources for their workloads.

Targeting customers deploying ML/AI and graphics-intensive applications, Vultr has made Cloud GPU available in collaboration with NVIDIA Corp. The company believes that what sets it apart from the competition is its flexibility — ranging from providing fractions of a single GPU to a large-scale cluster of dedicated NVIDIA HGX H100. For resource-intensive applications, the company says its bare-metal servers fit the bill. Delivered as a single-tenant dedicated server, Bare Metal gives its customers direct access to all server resources without any virtualization layers. For high-speed networking, customers can take advantage of burstable network connections (up to 25 Gbps).

Using the company's purpose-built control panel, customers can deploy and configure cloud instances from any device and manage server health and cloud spending with real-time price monitoring. For ease of use, Vultr provides a suite of one-click installers for common web applications and managerial tools such as cPanel, WordPress and Magenta, as well as data-science- and MLOps-optimized capabilities like Domino Data.

With the snapshot feature, users can duplicate their server setups on new instances across the Vultr network. Other key features available from the control panel include service utilization and availability monitoring, billing and subscription management, and access control of prebuilt marketplace images. For deploying and scaling containerized applications, there is Vultr Kubernetes Engine. Delivered as a fully managed offering, VKE will manage the control plane and worker nodes, and provide integration with other managed services such as load balancers, block storage and DNS. Vultr offers both hourly and monthly pricing.

Customers and partners

Much of the company's growth has been built on its experience and expertise in meeting the demand from developer-centric business. Typical use cases include web and application hosting, video gaming, SaaS applications, and video streaming. Among its reference customers are Overstock, Harvard Medical School and Raytheon Technologies Corp. Most of its customers sign up for at least three services — cloud compute, networking and storage. New customers, particularly enterprises, are driving demand for Cloud GPU, Kubernetes and managed database offerings.

Under the Vultr brand, the company primarily sells through channel partners to expand its global reach and credibility. It has developed a global network of more than 500 partners, mostly value-added resellers and managed service providers, while working with local telecom operators to reach a wider market. Through its strategic partnership with Airtel Ltd., Vultr is able to reach SMBs and large businesses in India. The cloud offerings are hosted in Airtel's datacenters in Bengaluru, Mumbai and Delhi NCR.

As hybrid and multicloud deployments come to fruition, the company believes it has a bigger role to play in the enterprise segment. To that end, since the third quarter of 2022, Vultr has invested in building a direct sales force, and claims a solid sales pipeline since then.

Strategic partnerships have been an important part of the company's product strategy. Vultr emphasizes its global partnership with NVIDIA as a key foundation for its market position and differentiated offerings in the GPU market. Other technology partners include Cloudflare, HashiCorp Inc. and Wasabi Technologies.

Strategy

Vultr has carefully defined its niche in the cloud infrastructure market with a focus on ease of use and price performance. Building on its partnerships with a number of cloud partners, including Cloud 66, Backblaze, Domino Data Lab and Console Connect, the vendor has formed the Cloud Alliance to provide pre-integrated cloud services (IaaS, PaaS and SaaS).

At present, the partnership with Console Connect is available via the Console Connect portal or by contacting the Vultr sales team. With Domino Data Lab, the Control Plane and Data Plane offerings are available on Vultr Marketplace. Served as a “storefront” for customers to gain access to a suite of pre-integrated services from Cloud Alliance partners, the marketplace is integrated with Vultr’s Control Panel.

Competition

Positioned as an alternative to the major cloud hyperscalers, Vultr is in direct competition with AWS, Microsoft Corp. (Azure) and Google for public cloud offerings. For hybrid and multicloud propositions, it will bump up against a broader set of players, including IBM Corp., Oracle Corp., British Telecom Group, T-Systems AG, Orange Business Services SA, Telefonica SA and NTT Ltd.

Alternative cloud offerings provided by the likes of 1&1 Internet, OVHcloud, DreamHost, Digital Ocean Inc., Ingram Micro Cloud and UpCloud are gaining ground in the SMB segment. In Asia-Pacific, the firm already has a cloud footprint in Japan, South Korea, Singapore and Australia. It is also eyeing new growth opportunities in India, in addition to its three existing cloud datacenter locations in the country. This puts it in competition with a number of regional cloud providers, including Telstra Ltd., SingTel ADR, Tata Communications Ltd. and IJ Inc., as well as Chinese cloud hyperscalers such as Alibaba Cloud and Tencent.

SWOT Analysis

STRENGTHS Constant, the parent company of Vultr, has a wealth of experience in the developer community, and benefits from the open-source advantages of cost savings and community. The company is among the first in the cloud GPU market. As late adopters engage with the cloud, the case for simpler and more cost-effective public cloud services becomes stronger.	WEAKNESSES Vultr is a small player compared with much larger competitors, many of which are focused on the broader use of AI to create added value and are pushing into key verticals.
OPPORTUNITIES As hybrid and multicloud deployments continue to take shape with a growing demand for cloud GPUs, the time has come for providers like Vultr to leverage partners to drive global expansion and target customers that may not need/want the complexity of hyperscalers’ offerings.	THREATS Growing beyond the developer community, the company will compete squarely with major cloud hyperscalers and their partners for enterprise opportunities.

CONTACTS

The Americas

+1 877 863 1306

market.intelligence@spglobal.com

Europe, Middle East & Africa

+44 20 7176 1234

market.intelligence@spglobal.com

Asia-Pacific

+852 2533 3565

market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

Copyright © 2023 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON “AS IS” BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT’S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence’s opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not endorse companies, technologies, products, services, or solutions.

S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global’s public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.