



Market Insight Report Reprint

As cloud native takes hold, VMs lose their grip

June 2 2021

by **Jean Atelsek, Melanie Posey**

Nearly half of container adopters in our recent survey say they are deploying new or refactored applications on purpose-built container instances rather than on VMs. Solving for the complexity of Kubernetes could mean winning the ultimate prize: hosting the management plane that enables visibility and control of distributed applications in heterogeneous environments.

451 Research

S&P Global

Market Intelligence

This report, licensed to Chronosphere, 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

In what may portend a shift from VMs to clusters as the default unit of IT operations, a recent survey found that almost half of container adopters are deploying new or refactored applications on purpose-built container platforms rather than on VMs. Whether truly ‘invisible’ infrastructure will take broader hold – i.e., with workloads managed exclusively at the application level rather than via hardware (virtualized or not) – remains to be seen, since VMs are still the primary mode of deployment even for net-new and modernized applications.

Cloud providers, systems vendors and software suppliers are all solving for the complexity of Kubernetes in hopes of winning the ultimate prize – hosting the management plane that enables visibility and control of distributed applications across heterogeneous environments.

THE TAKE

Applications are taking on new forms. To leverage cloud capabilities, enterprises are onboarding new modes of development (like DevOps and GitOps), deploying software using cloud-native architectures (containers and serverless) and tapping new sources of data (e.g., from edge devices and social media), all in the service of faster product development and better customer outcomes. Recent 451 Research surveys have shown that, despite the ability to have containerized applications drive the infrastructure via Kubernetes, rather than operate from the infrastructure up, VMs retain a solid presence as the default unit of IT management. This is partly a reflection of Kubernetes complexity, which is being addressed by stakeholders across the spectrum with more easily adoptable platforms built on opinionated stacks.

The success of this approach is revealed in survey data showing that nearly half of organizations are deploying net-new and modernized workloads without VMs. With enterprises prioritizing investment in moving off of legacy infrastructure and refactoring applications, it may be only a matter of time before IT operations reach a tipping point, whereby VMs become a vestigial (albeit still useful) component of IT environments. The question remains: How much time, and what will ultimately shift the balance?

Businesses prioritize IT modernization

As organizations pivoted to support remote working and schooling through the COVID-19 pandemic, many quickly learned to appreciate the value of ‘virtuality’ – the ability to use professionally managed services for running IT on a pay-per-use or per-user basis. According to 451 Research’s Voice of the Enterprise: Cloud, Hosting & Managed Services, Budgets and Outlook 2020 survey, fielded in Q4 2020, 40% of organizations are using more public cloud services due to the pandemic, and a large proportion of them expect changes in public cloud usage made during the stay-at-home orders to become permanent.

This shift, combined with an explosion in data being collected, stored and used to inform business decision-making and customer engagement, has resulted in an uptick in app modernization investment. According to 451 Research’s Voice of the Enterprise: Digital Pulse, Digital Strategy 2021 survey, 56% of companies consider digital technology a high strategic priority, and 46% consider modernizing legacy systems, infrastructure and applications as a major investment area for implementing that strategy in the next 12 months. In fact, modernization was second only to security (51%) as an investment priority, and 36% of respondents expect modernization to have a transformational impact on their businesses over the next two to three years.

A key component of this effort is containerization. Container adoption is growing but far from universal: 35% of organizations using public clouds in 451 Research’s forthcoming Voice of the Enterprise: Cloud, Hosting and Managed Services, Workloads & Key Projects 2021 survey said they do not plan to use containers. That compares with 37% that are using VMs to run containers, and 23% that plan to replace VMs with containers.

New workloads are shifting to container platforms

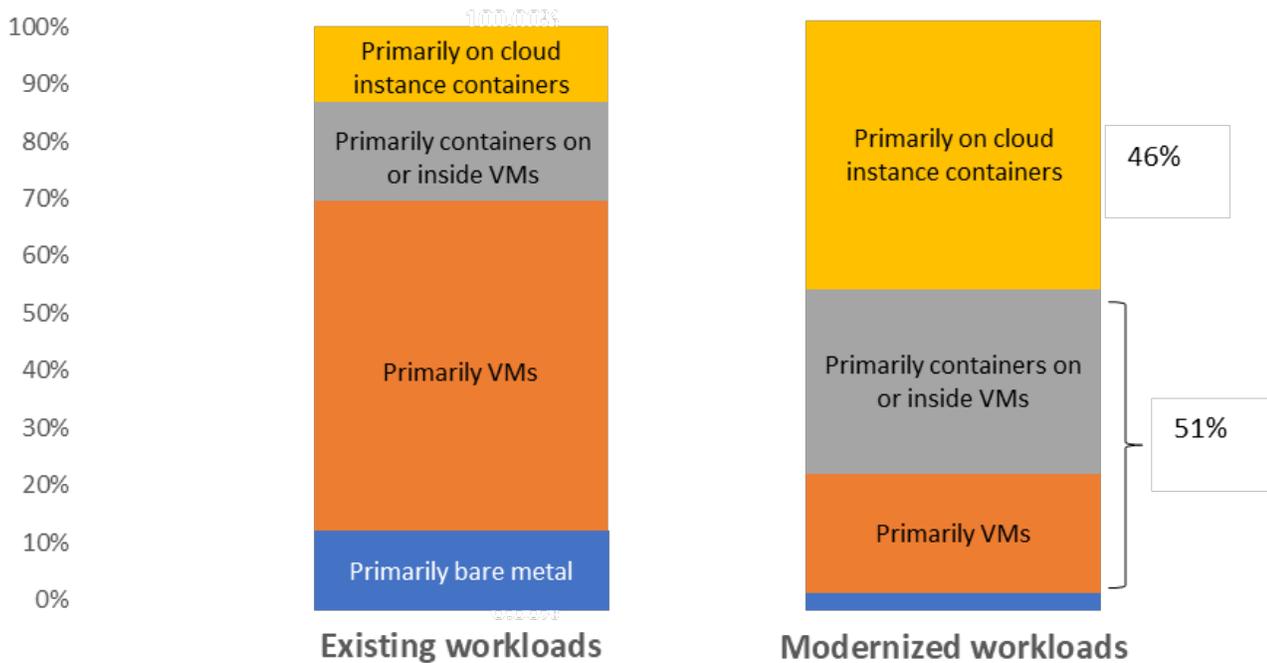
The advantages of containers are compelling. Their relatively small footprint versus VMs, and more efficient resource utilization, enable portable and economical app deployment. Widespread adoption of the open source Kubernetes orchestration platform makes it possible to deliver services across multiple on-premises and off-premises environments, including VMs and bare metal.

That doesn't mean it's easy. Enthusiasm about the promise of Kubernetes has met with the difficulty in implementation. The very extensibility that makes it so appealing also creates a multitude of configuration options that must be carefully set and managed. For born-in-the-cloud organizations this is less complex. The difficulty is compounded when legacy infrastructure and applications need to be brought on board. Enterprise modernization efforts are one way of solving for this complexity, but both approaches (managing legacy workloads via Kubernetes and refactoring applications for cloud-native management) require significant skills and investment, and provide uncertain returns.

Cloud providers and ISVs have created serverless container platforms and software to offer a third option for net-new workloads. These systems are built on opinionated stacks that leave it to the service provider to provision underlying infrastructure based on application demands. In this scenario, container orchestration programs such as Kubernetes effectively move addressable infrastructure up a layer, from the server to the application.

These purpose-built platforms are gaining ground, judging by the deployment profiles of existing workloads versus net-new/modernized workloads in our recent survey (see figure below).

Packaging of Existing Versus Net-New or Modernized Workloads



Q. Which of the following best describes your organization's approach to packaging environments for existing and modernized/net-new workloads?

Base: Current or future container users

Source: 451 Research's Voice of the Enterprise: Cloud, Hosting and Managed Services, Workloads & Key Projects 2021

Among the services available for this type of deployment are:

- **AWS Fargate.** AWS launched Fargate in 2017 to enable serverless management of containerized applications for Amazon's proprietary Elastic Container Service. At re:invent 2019, it released Amazon Elastic Kubernetes Service (EKS) on AWS Fargate into general availability. Fargate uses lightweight Firecracker micro-VMs based on the KVM hypervisor.
- **Azure Container Instances.** Also introduced in 2017, Azure Container Instances are designed for workloads that can operate in isolation, such as basic applications and task automation. Container Instances cannot run existing Kubernetes pod definitions, but they can interact with Azure Kubernetes Service clusters via virtual nodes.
- **Google Cloud Run.** Released in 2019, Cloud Run is based on Google-developed Knative serverless technology, and can be invoked via web requests or Pub/Sub events. It can be deployed independently or as part of a GKE cluster, and aims to help developers get automatically scalable clusters up and running with a minimum of fuss.
- **GKE Autopilot.** Launched in February 2021, GKE Autopilot is a mode of operating GKE clusters, where Google manages the infrastructure and control plane. It locks down certain configuration details based on best practices developed by Google in running the GKE fleet for production workloads.
- **IBM Code Engine.** Introduced in 2020, IBM Code Engine is designed to run serverless, container or batch jobs via a unified platform. It relies on open source foundations including Knative, Tekton (a framework for creating CI/CD pipelines) and Paketo (for language runtime support).
- **Spot Ocean by NetApp.** Ocean, launched in 2018, is a serverless container SaaS available on AWS's EKS, Azure's AKS and Google's GKE Kubernetes services. Ocean-managed clusters can also be set up in AWS ECS, Red Hat OpenShift and Kops.

For a price comparison of some of these services versus self-managed Kubernetes environments under various circumstances, see *The Economics of Cloud Containers and Kubernetes*.

Even VMware – the company with arguably the most to lose from a transition away from VM-based deployment – is evolving to adapt to cloud-native. Its Tanzu portfolio is key to the company's transition from a single-platform virtualization vendor to a multicloud/hybrid platform provider. In late April, it announced that vSphere 7.0 would support a new Virtual Machine Service, offering Kubernetes-native provisioning of VMs and guest operating systems using the declarative model that's central to Kubernetes management.

Clearly, VMs are far from passé. The isolation, security, familiarity and manageability of VMs gives them staying power. In the absence of a standard for serverless container operation, providers and vendors are implementing them in unique ways. This creates a danger of lock-in – the very circumstance many businesses are counting on Kubernetes to avoid.

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 © 2021 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.