Enterprise Strategy Group | Getting to the bigger truth.™

ESG SHOWCASE

Google Cloud VMware Engine with NetApp Cloud Volumes Service

Date: March 2022 **Author:** Rob Strechay, Sr. Analyst

ABSTRACT: Organizations are starting to take a "cloud-when-right" approach to transforming their on-premises applications to cloud-native applications. Using Google Cloud VMware Engine with NetApp Cloud Volumes Service can be a cost-effective and performant middle ground on the journey to becoming cloud-native. NetApp brings a plethora of capabilities to this Google Cloud native service that makes migration easier.

Overview: IT Transformation Is the First Step in Digital Transformation

As organizations continue to make strides with their digital transformation initiatives, IT modernization remains a work in progress. Part of many organizations' IT modernization process is moving more workloads to the cloud. According to ESG research, 77% of applications are strong candidates or potential candidates to move to public cloud services over the next 5 years.[1] Organizations are not solely focused on a cloud-first strategy, but rather a cloud-when-right strategy, moving their workloads to the cloud when it makes sense for the application to live there. Seventy-nine percent of organizations embrace a modernization strategy for their data center, which consists of some combination of shifting adoption to public cloud resources, increasing interoperability to public cloud resources, and investing to achieve a cloud-like experience on-premises.[2]

Cloud infrastructure is essential to a modern IT strategy. As application migration candidates become mission-critical applications, the risk grows, and it becomes essential for the IT decision makers to leverage the right technology. Organizations will have to make choices on how they migrate to the cloud and what infrastructure the application will live on. Many "born-in-the-cloud" applications are constructed with a "Lego building block" approach (i.e., every new service has a choice of underlying services, such as picking between an object, block, or file service to build on top of).

Many times, these choices come with limitations on performance, resiliency, or durability. Differentiating among the cloud service choices and understanding how well your applications may perform is not straightforward. Migrations to the cloud for applications that are built for the enterprise can be complex and costly. This is especially true if an organization's applications are moving from on-premises virtualized servers into cloud-native instances. One option to consider is moving applications from on-premises virtualized servers to a cloud-native VMware service, such as Google Cloud VMware Engine.

---

[1] Source: ESG Complete Survey Results, *2022 Technology Spending Intentions Survey*, November 2021.
[2] Source: ESG Research Report, *Application Infrastructure Modernization Trends Across Distributed Cloud Environments*, to be published.
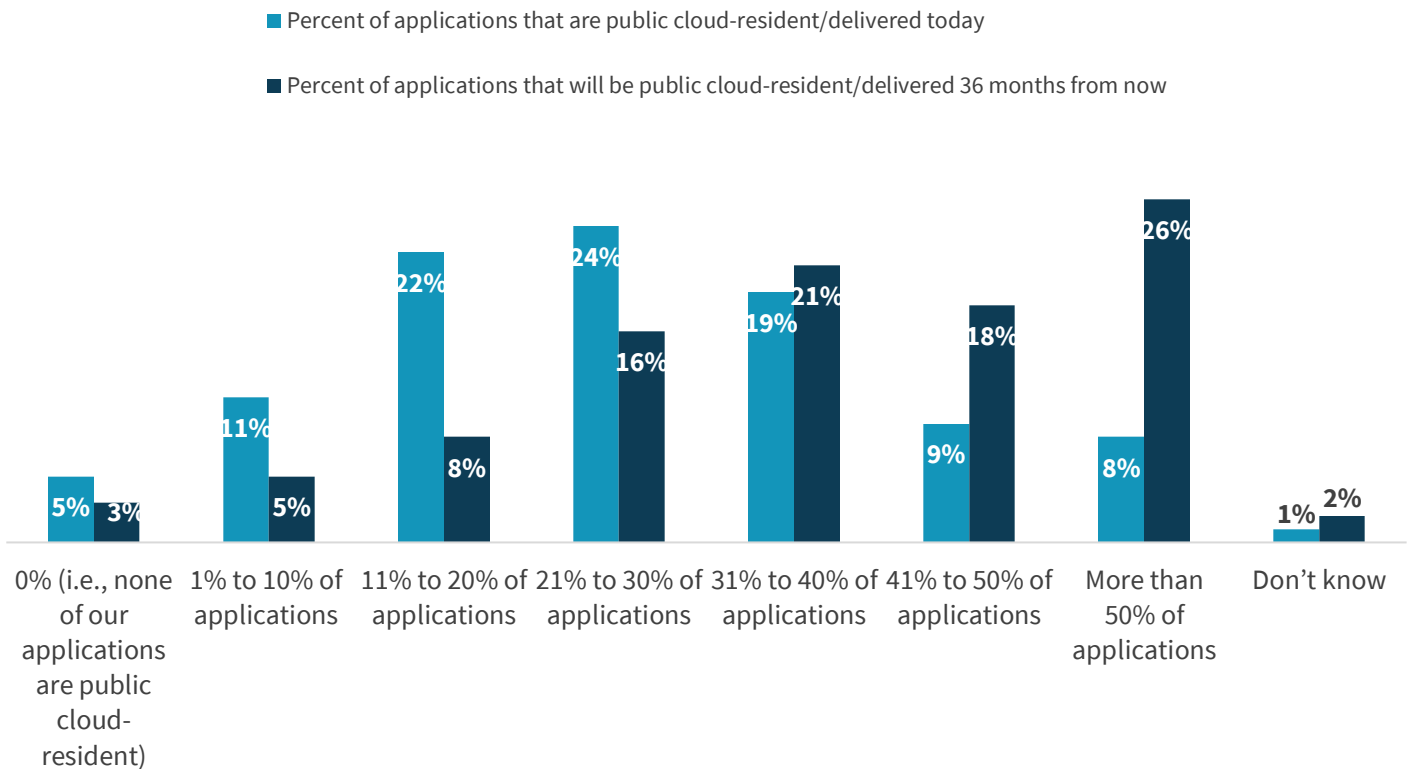
Organizations Want to Be Cloud-First but Need Performance, Resilience, and Cost-Effectiveness

Almost all organizations, 95% of them, are currently using public cloud services, and the vast majority (82%) of organizations are leveraging more than one public cloud infrastructure service provider. In addition, only 8% of organizations reported that 50% or more of their applications are currently public cloud–resident/delivered, with 26% of respondents reporting that 50% or more of their applications will be public cloud–resident/delivered 36 months from now (see Figure 1).[3]

Figure 1. Percentage of Applications that Are Public Cloud–Resident/Delivered

Of all the business applications used by your organization, approximately what percentage is currently public cloud–resident (i.e., SaaS apps and/or custom apps that run on public cloud infrastructure)? How do you expect this to change – if at all – over the next 36 months? (Percent of respondents, N=706)

■ Percent of applications that are public cloud-resident/delivered today

■ Percent of applications that will be public cloud-resident/delivered 36 months from now

Category	Today	36 months
0% (i.e., none of our applications are public cloud-resident)	5%	3%
1% to 10% of applications	11%	5%
11% to 20% of applications	22%	8%
21% to 30% of applications	24%	16%
31% to 40% of applications	19%	21%
41% to 50% of applications	9%	18%
More than 50% of applications	8%	26%
Don't know	1%	2%

Source: ESG, a division of TechTarget, Inc.

According to an ESG research survey, more than two-thirds of organizations say that it takes at least a week to migrate an application to the cloud, with 38% saying it takes multiple weeks or longer.[4] Even though organizations with a cloud-first approach to workload placement outnumber those with an on-premises first approach by 4:1 (44% vs. 9%), 47% of organizations are taking a cloud-when-right approach, considering both on-premises technology resources and public cloud services equally when considering how to deploy new applications (see Figure 2).[5]
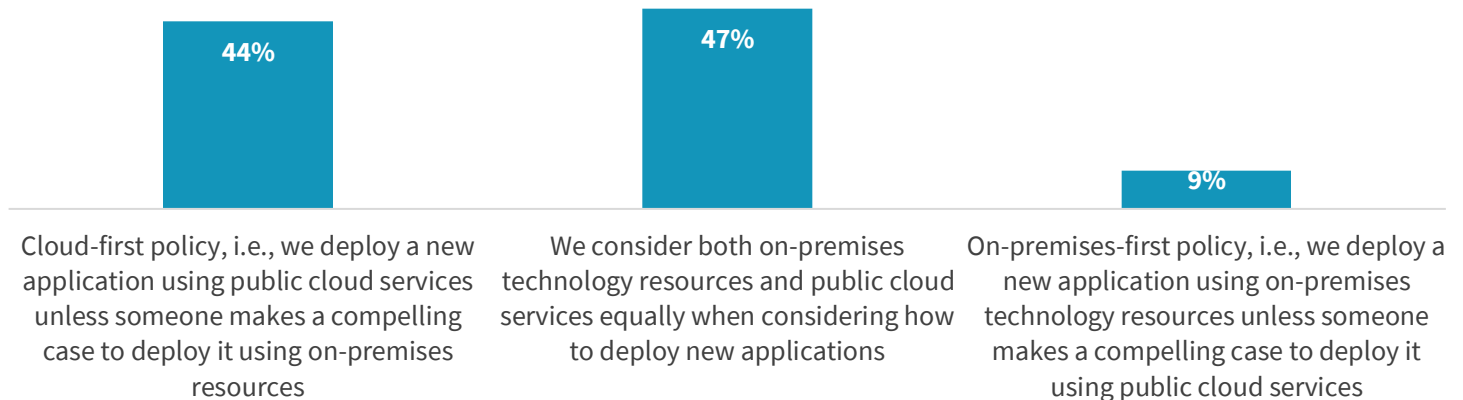
---

[3] Source: ESG Complete Survey Results, *2022 Technology Spending Intentions Survey*, November 2021.
[4] Source: ESG Research Report, *Application Infrastructure Modernization Trends Across Distributed Cloud Environments*, to be published.
[5] Source: ESG Complete Survey Results, *2022 Technology Spending Intentions Survey*, November 2021.

Figure 2. Percentage of Applications that Are Public Cloud–Resident/Delivered

Which of the following best describes the approach your organization takes when it comes to new application deployments? (Percent of respondents, N=670)

44%	47%	9%
Cloud-first policy, i.e., we deploy a new application using public cloud services unless someone makes a compelling case to deploy it using on-premises resources	We consider both on-premises technology resources and public cloud services equally when considering how to deploy new applications	On-premises-first policy, i.e., we deploy a new application using on-premises technology resources unless someone makes a compelling case to deploy it using public cloud services

Source: ESG, a division of TechTarget, Inc.

What and Why to Consider Moving to Google Cloud VMware Engine?

With the cloud-when-right approach becoming more prevalent, organizations are revisiting which applications should become cloud-native and which should remain as-is on premises. One option can be to move "as is" in a virtual machine (VM) form to the cloud, then refactor once in Google Cloud VMware Engine. Google has deployed the VMware software-defined data center (SDDC) software stack, including a hypervisor, tools, management capabilities, support, and monitoring, on the Google Cloud hardware stack. The service is built and sold by Google Cloud, is available in the Google Console, and is a Google-native service that integrates with third-party storage, disaster recovery, and backup ISV solutions that an organization might already use on premises.

What to Think About When Evaluating Google Cloud VMware Engine

With Google Cloud VMware Engine, customers can move many more "as-is" applications in a VMware configuration. This approach can provide considerable advantages based on organizations knowing how the applications are configured for performance with the amount of storage, CPU, and memory. Internal resources in the organization will also be able to understand more quickly what availability and data protection requirements are needed. Manageability is another advantage of using Google Cloud VMware Engine, with organizations using the same tools they use today to manage their VMware on-premises environment.

With a service like Google Cloud VMware Engine, organizations have choices to make in how they configure the VMs they migrate. This means rightsizing the VMs before they are moved to the cloud service.

Don't Go It Alone: Engage with a Trusted Subject Matter Expert

Thirty-two percent of organizations say moving apps/data between data centers and multiple public cloud services is one of the greatest challenges they face as a result of using multiple CSPs.[6] To avoid these migration challenges, organizations should engage in a conversation with a value-added reseller (VAR) to assist in the move to Google Cloud VMware Engine. Additionally, they should find a VAR that can help map their current IT architecture and operations model to the right Google Cloud VMware Engine services and that has experience with the networking and migration aspects of moving to Google Cloud VMware Engine. Importantly, the VAR should understand both the on-premises networking architecture and how to use a capability like VMware NSX-T and HCX to bridge a layer-2 network to the Google Cloud VMware Engine service.

This not only includes rightsizing the CPU and memory but ensuring an organization is using the right storage. By default, Google Cloud VMware Engine uses VMware vSAN-based storage in a three-node minimum configuration. As expected with a cloud service, organizations pay for a minimum amount of storage but have choices in how they can add storage to the VMs. By default, Google Cloud VMware Engine scales by adding another node to get more CPU, memory, or storage. This can be inconvenient if an organization only needs more storage for more VMs to be deployed but has plenty of CPU and memory.

NetApp Cloud Volumes Service as a Native Google Cloud Service

NetApp Cloud Volumes Service for Google Cloud delivers file systems with high throughput, IOPS, data protection services, and performance levels of NetApp's latest all-flash storage arrays in Google Cloud. NetApp Cloud Volumes Service is natively available in the Google Cloud Console. Also, because it uses NetApp ONTAP software, native NetApp storage services are available, such as NetApp Snapshot copies, volume clones, and volume replication. As a native service in Google Cloud, NetApp Cloud Volumes Service is available to regular Google Cloud instances and now to Google Cloud VMware Engine for both in-guest storage and to support VMware datastore storage. Currently, support for VMWare datastore support by Cloud Volume Service in Google Cloud VMWare Engine is available in preview. NetApp Cloud Volumes Service comes with three service levels: Standard (16MB/s per terabyte), Premium (64MB/s per terabyte), and Extreme (128MB/s per terabyte).

Typical Deployments of Google Cloud VMware Engine with NetApp Cloud Volumes Service

The service level for the NetApp Cloud Volumes Service an organization uses will depend on the use case. For use with Google Cloud VMware Engine as the storage repository for VMware datastores, an organization may want to start with the Premium service to provide cost-effective, high-throughput storage for the mixed workloads of multiple VMs. A key advantage of Cloud Volumes Service is that an organization can mix and match storage types and change from one storage type to another to meet workload storage requirements with some of the VMs that live on a datastore with Standard level of service, some VMs on internal Google Cloud VMware Engine VMware vSAN storage, and some VMs that have their home directories on NetApp Cloud Volumes Service using the Standard SW service level. This way the organization can deliver the correct service level based on the cost and performance.

The advantage of using Google Cloud VMware Engine with NetApp Cloud Volumes Service is the ability to protect resident VMs with NetApp Snapshot copies and data replication. Using the NetApp Cloud Volumes Service filesystem for data protection will offload the compute nodes from having to perform these activities, achieving a higher density of VMs per node. It is important to note that the number of VMs per node will vary based on the applications in the VMs. From a manageability capability, using NetApp Cloud Volumes Service for data protection leverages the integration with VMware's

---

vSphere APIs for Array Integration to make the data protection activities transparent and managed out of the Google Cloud VMware Engine VMware vCenter console.

Figure 3. Native NetApp Cloud Volume Service Servicing Google Cloud VMware Engine
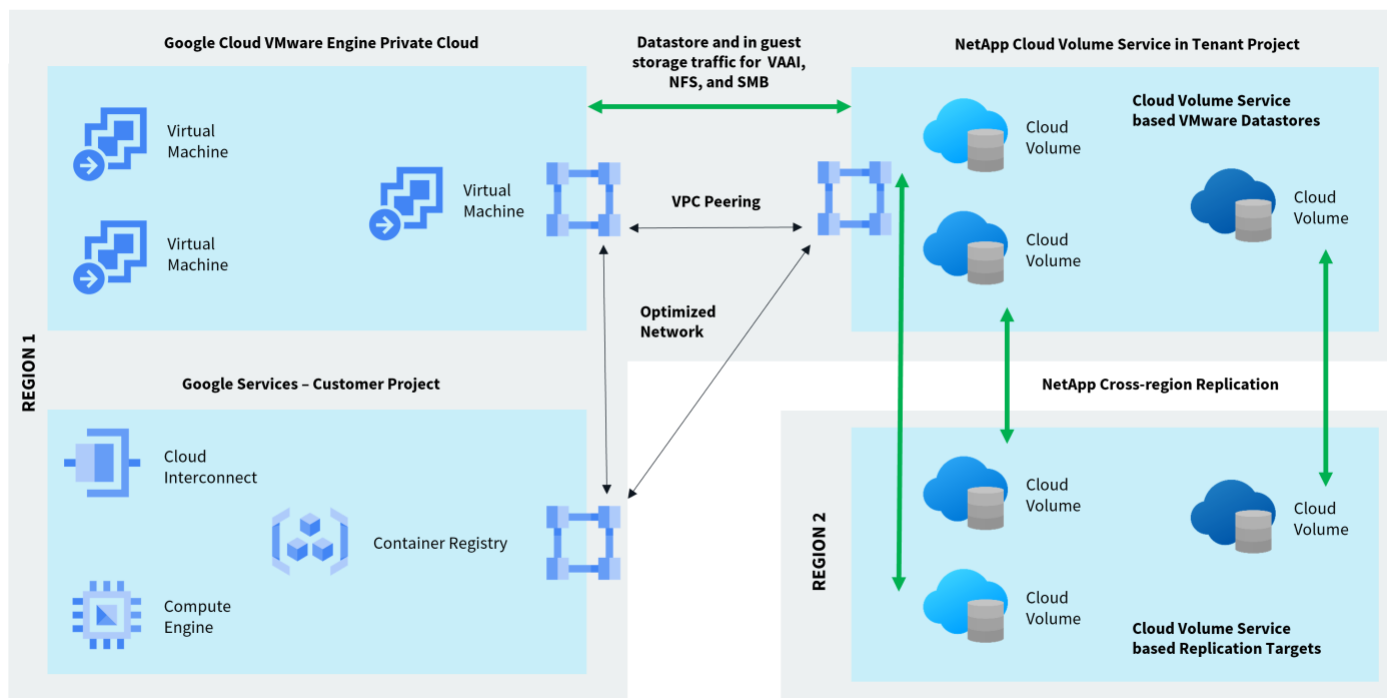
Source: ESG, a division of TechTarget, Inc.

Another advantage to using the native NetApp Snapshot technology and data replication is that an organization will have a disaster recovery site with a minimal configuration. For example, an organization might decide to move some of its VMs to Google Cloud VMware Engine before it ultimately rebuilds some of the applications into cloud-native ones running on native Google Cloud. Once the applications are in Google Cloud VMware Engine, the organization can use the integrated NetApp cross-region replication (CRR) to protect those VM workloads in another NetApp Cloud Volumes Service filesystem in another region. The organization does not have to instantiate a Google Cloud VMware Engine cluster in that region until it wants to test or perform a disaster recovery operation.

As part of the migration use case from on-premises to Google Cloud VMware Engine with NetApp Cloud Volumes Service, customers can use NetApp Cloud Sync across the VMware HCX network connection. This is a free data movement option when customers go from on-premises NetApp ONTAP to NetApp Cloud Volumes Service.

The Bigger Truth

Today, organizations still lean toward being cloud-first but understand that some applications are not built to be cloud-native "as-is" and may never be rebuilt into cloud-native versions. This is even with organizations telling ESG that 77% of their applications are strong candidates or potential candidates to move to public cloud services over the next 5 years.[7] With native VMware in Google Cloud, organizations can now take the first step in moving applications "as is" to the cloud in a manner that preserves many of the ITOps capabilities for their current on-premises VMware estate.

---

[7] Source: ESG Complete Survey Results, *2022 Technology Spending Intentions Survey*, November 2021.

Now with the combination of the Google-native Google Cloud VMware Engine and NetApp Cloud Volumes Service, organizations can make the journey to the cloud even simpler. The addition of NetApp Cloud Volumes Service for Google Cloud VMware Engine VMware datastore and in-guest storage allows organizations to move to a storage platform that might be more common to their current environment. Using NetApp Cloud Volumes Service with Google Cloud VMware Engine enables organizations to use VMware's vSphere APIs for Array Integration to manage storage, monitor performance, and enable data protection that is offloaded from the Google Cloud VMware Engine nodes' VMware vSAN.

Organizations need to consider how they want to achieve their goals, stage their migrations, and plan for what moves when. This is where bringing in a VAR partner who has been there and done that will make things run more smoothly—especially one that has subject matter expertise in both Google and NetApp. Leveraging the technical capabilities of the VMware networking stack and NetApp data protection and movement capabilities can make migrations of "as is" VMs a snap. Once an organization gets to normal day-to-day operations in Google Cloud VMware Engine on NetApp Cloud Volumes Service storage, it can leverage NetApp data protection and movement capabilities to do cross-region replication cost-effectively and performantly to provide disaster recovery. Organizations that are comfortable with their current VMware infrastructure and looking for an on-ramp to the cloud should take a serious look at the combination of Google Cloud VMware Engine with NetApp Cloud Volumes Service.

Enterprise Strategy Group is an integrated technology analysis, research, and strategy firm that provides market intelligence, actionable insight, and go-to-market content services to the global IT community.

www.esg-global.com   contact@esg-global.com   508.482.0188