Google Cloud

# Google Cloud VMware Engine Performance Migration & Benchmarks

# Table of Contents

Google Cloud


Chapter 1

## Executive Summary

Introduction 4

Chapter 2

## On-premises and Google Private Cloud Connectivity

Overview 5

End-to-end connectivity using Megaport 8

VPN connectivity 10

Chapter 3

## Virtual Machine Migration between On-premises and GCP

Virtual machine migration yool 10

Virtual machine migration test procedure 12

Virtual machine types 12

Virtual machine waves 13

Virtual machine migration test scenarios 13

Virtual machine migration test process 14

Chapter 4

## Benefits of Google Cloud & Google Cloud VMware Engine Networking

Dynamic routing mode 16

Google VPC network peering 18

Multi-VPC connectivity 19

Cloud DNS 20

Chapter 5

## Conclusion

Chapter 1

# Executive Summary

Google Cloud VMware Engine (GCVE) allows a user to deploy a managed VMware environment within an Enterprise Cloud Solution. Utilizing Google Cloud allows an individual to tie into existing services and cloud capabilities; one of those services and solutions mentioned within this document is our Hybrid Cloud Extension, also known as HCX. HCX provides the end user a seamless transition from on-prem to the cloud, allowing a systems administrator to quickly deploy a private cloud and scale their needed Virtual Machines accordingly. The key objective within this white paper is to showcase the attributes of a Google Cloud VMware Engine migration journey and the technical attributes and capabilities it is capable of. The proposed referenced solution is well suited for organizations looking to begin their cloud migration journey and understand the technical requirements within the process.

## Introduction

There are three key objectives of this white paper. First, analyzing the performance of migration of virtual machines from an on-premises data center to Google Cloud with Google Cloud VMware Engine private cloud. Second, providing detailed technical guidance on establishing end-to-end connectivity between on-premises data centers and Google Cloud. Third, explaining the benefits thatGoogle Cloud offers to the virtual machines that are difficult to achieve in an on-premises data center environment. Organizations planning on migrating workloads to Google Cloud can use the performance benchmark results and architecture guidance offered in this white paper to optimize their workload migration while minimizing disruptions.

The performance tests and analysis developed in this white paper provide insight into several aspects of workload migration from on-premises to Google Cloud VMware Engine. First, it helps organizations learn a robust migration architecture based on Google, VMware, and Megaport best practices. The architecture presented in this white paper offers optimized workload migration from on-premises to Google Cloud. Second, using the data presented in the whitepaper, organizations can determine the length of time it would take to migrate their workloads from on-premises to Google Cloud. Third, organizations can learn the benefits that Google Cloud offers to improve the service delivery of business applications.

Organizations either in the evaluation phase or in the execution phase of workload migration from on-premises to Google Cloud can benefit from the content presented in this white paper. Organizations face a number of unknowns while migrating workloads to Cloud environments, including cost of migration, effort, and risk of business application outage. Migration costs include establishing end-to-end connectivity between on-premises to Google Cloud and engaging professional services. The effort includes several organizational tasks including planning and preparations, change management, testing and validation, and training and skills development. The risks include the type and duration of application outages prior, during, and after the workload migration. Organizations may also experience technical and business challenges to be able to extend networking and security between on-premises and Cloud environments for non-disruptive workload migration. Performance aspects of workload migration from on-premises to Google Cloud is of quintessential importance as they determine the cost, effort, and risks of migration. One migration approach is to shut down the virtual machines prior to migrating to Google Cloud to reduce migration-related overheads. However, such a “cold migration” approach is disruptive to the business. Shutting down the virtual machines and starting them back up causes the applications to go offline, affecting business operations. “Hot” or “live migration” of virtual machines does not require powering off virtual machines, hence it doesn’t affect business operations.

VMware provided an on-premises lab environment to test workload migration to Google Cloud. The lab environment was connected to Google Cloud using Megaport private connectivity. Several test scenarios were developed — each consisting of virtual machines of various configurations — to measure the performance benchmark data. A team of experts — highly skilled and certified on both VMware and Google technologies — were engaged to perform the migration tests, analyze the results, and build documentation.

Once workloads are migrated to Google Cloud, customers can utilize the power of Google Cloud, including dynamic routing mode, Google VPC peering, multi-VPC peering, and Cloud DNS. The benefits of each of these technologies are analyzed and presented in this white paper. Performing testing and validation of the benefits of each of the above technologies is beyond the scope of the white paper. However, inferences were made on this white paper based on previously validated, proven, and documented features of the Google Cloud.

Chapter 2

# On-premises and Google Cloud Connectivity

## Overview

End-to-end connectivity was established between the on-premises data center and Google Cloud as a prerequisite to migrating virtual machines between them. An on-premises lab was secured by VMware with all required VMware products installed. The lab also had end-to-end private connectivity to Google Cloud using Megaport. The lab had enough capacity (storage, RAM, and CPU) to store the virtual machines identified in the test scenarios. Furthermore, a 10GB port was configured with Megaport to handle a maximum bandwidth of 10GB between on-premises and Google Cloud VMware Engine. Note that, even though 10GB connectivity was established by Megaport, only 2GB maximum bandwidth was utilized for the test scenarios due to a limitation in HCX.

The following product versions were utilized in the on-premises lab environment

| Product | Version | Build |
|---|---|---|
| vCenter | 7.0.2 | 17958471 |
| ESXi | 7.0.2 | 17867351 |
| VSAN | 7.0.2 | 18426014 |
| NSX | 3.1.2 | 17883596 |
| HCX | 4.2 | 18422312 |

**Table 1: VMware product versions utilized in the on-premises lab**
On-premises CPU: Intel(R) Xeon(R) CPU E5-2630 v4
On-premises ESXi Host: PowerEdge R730xd

The following product versions were utilized in the Google Cloud VMware Engine environment for some of the test runs marked as * in the benchmark table later this chapter.

| Product | Version | Build |
|---|---|---|
| vCenter | 7.0.1 | 18392253 |
| ESXi | 7.0u1 | 17168206 |
| VSAN | 7.0u1 | 16850804 |
| NSX | 3.0.2 | 16887200 |
| HCX | 4.0.2 | 17881554 |

**Table 2: VMware product versions utilized in the Google Cloud VMware Engine**
Google Cloud VMware Engine CPU: Intel(R) Xeon(R) Gold 6240 CPU
Google Cloud VMware Engine Host: PowerEdge R640

The following product versions were utilized in the Google Cloud VMware Engine environment for some of the test runs.

| Product | Version | Build |
|---|---|---|
| vCenter | 7.0.2 | 18895595 |
| ESXi | 7.0.2 | 18836573 |
| VSAN | 7.0u1 | 16850804 |
| NSX | 3.1.2 | 17883600 |
| HCX | 4.2.2 | 18868175 |

**Table 3: VMware product versions utilized in the Google Cloud VMware Engine**
Google Cloud VMware Engine CPU: Intel(R) Xeon(R) Gold 6240 CPU @ 2.60GHz
Google Cloud VMware Engine Host: PowerEdge R640

## End-to-end connectivity using Megaport

End-to-end connectivity between on-premises and Google Cloud VMware Engine was established using Megaport and Google Cloud Interconnect. This setup allowed Layer 3 network connectivity between the VMware network's on-premises data center and Google Cloud.

Megaport provides a portal to configure the connectivity between on-premises and Google Cloud. Using the Megaport portal, one must first create a physical port to establish connectivity between a physical on-premises data center and Megaport.

It only takes a few steps. Log into your Megaport portal and order a port in the required data center location. Provide the Letter of Authorization to your DC provider to facilitate the physical cross connect.

The next step is the creation of a cloud router using Google Cloud Console's Hybrid Connectivity link. Under Hybrid Connectivity, choose Interconnect. Once the Google Partner Interconnect is created, a pairing key is generated, which must be entered into Megaport's Google Connectivity page.

After that, you must create a Virtual Cross Connect (VXC) for Google using the pairing key generated from Partner Interconnect. A VXC is an L2 circuit that provides connectivity between Megaport and Google Partner Interconnect. Once the data center provider has completed the physical cross connect, you can build out connectivity to Google Cloud.

01 Either click on the Google Cloud tile or "Add Connection:"

02 Select "Google Cloud"

03 Put in the pairing key from Google Cloud Partner Interconnect

04 Select the country and on-ramp location where you want to connect to Google Cloud, and click "Next"

05 Add a name for the VXC

06 Select your rate limit

07 Select a VLAN corresponding to A-End on-premises Port VLAN

08 Next, add VXC and click "Order." Now we see this is in a deployable state

09 In your Google Cloud console, accept the connection

Assuming the on-premises device is configured for BGP, you should then see BGP turn into an established state.

Upon successful validation, both Partner Interconnect and the Megaport service will show that the connection between Google Cloud and Megaport is valid.

The following figure depicts the Megaport components needed to configure end-to-end connectivity between the on-premises data center and Google Cloud.



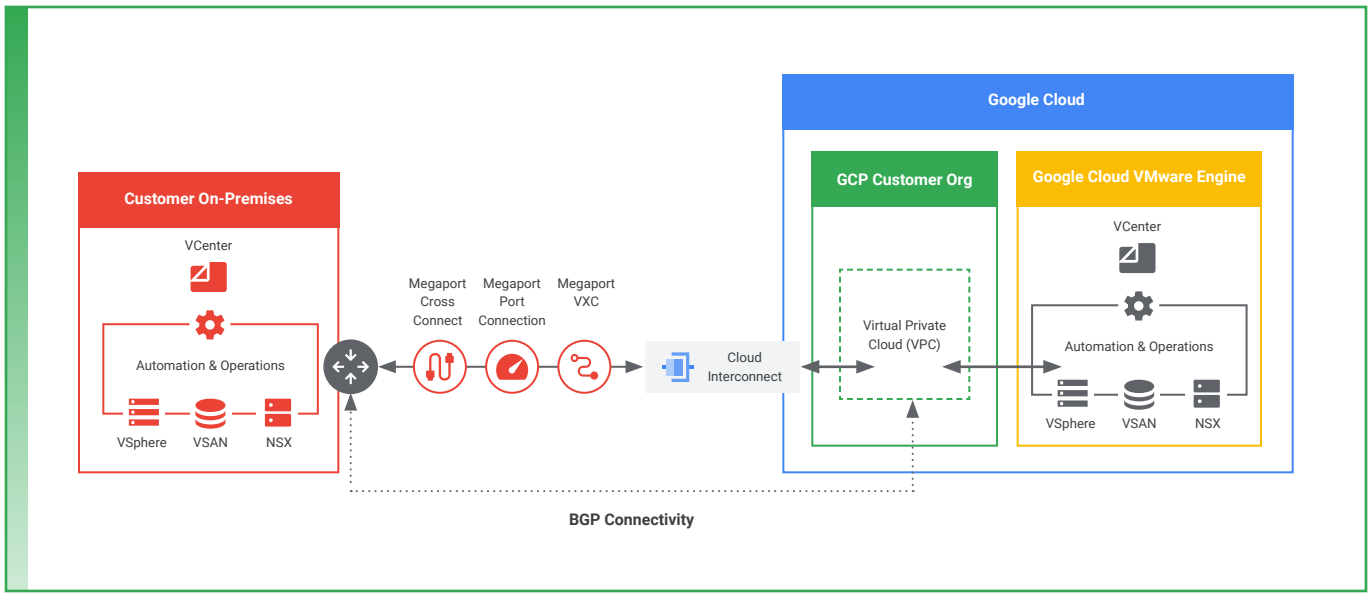


**Figure 1: Megaport Components for L3 Connectivity between on-premises and Google Cloud**

### VPN connectivity

Google offers Cloud VPN to connect on-premises data centers and Google Cloud using internet based IPsec tunnels. Cloud VPN is useful for low volume data transmission between on-premises and Google Cloud, hence it's useful for migration of a small number of VMs from on-premises to Google Cloud VMware Engine. Cloud VPN establishes L3 communication between on-premises and Google Cloud VMware networks via the internet. L3 connectivity is sufficient for the migration of VMs from on-premises to Google Cloud. Detailed analysis of the use and performance of Cloud VPN for the migration of virtual machines from on-premises to Google Cloud VMware Engine is beyond the scope of this white paper.



Chapter 3

# Virtual Machine migration between on-premises and Google Cloud VMware Engine

### Virtual machine migration tool

VMware HCX was chosen as the migration tool to migrate virtual machines between on-premises and Google Cloud VMware Engine. This tool is bundled as part of the Google Cloud VMware Engine private cloud. It provides several key features on top of virtual machines migration, including WAN Optimization, Replication Assisted vMotion, L2 Extension, and VM Mobility Platform. During the creation of a Virtual Private Cloud (VPC) in Google Cloud, VMware HCX Manager and HCX IX appliances are deployed and configured in Google Cloud VMware Engine.

To migrate virtual machines between on-premises and Google Cloud VMware Engine, HCX Manager and HCX IX appliances must be deployed and configured in the on-premises environment per VMware recommended best practices. HCX WAN Optimization appliance is optional but is recommended. A license key must be downloaded from the Google Cloud VMware Engine private cloud console and applied during the on-premises installation of HCX.

HCX set-up on-premises requires configuration of several network profiles to enable communication between on-premises vCenter and Google Cloud VMware Engine vCenter. Note that HCX utilizes the backend physical connectivity between on-premises and Google Cloud VMware Engine private cloud, which was created using Megaport. The HCX WAN Optimization appliance was deployed and configured in the VMware lab environment. The WAN Optimization service improves performance characteristics of the private lines or internet paths by applying WAN optimization techniques like data de-duplication and line conditioning.

Once virtual machine migration is initiated, HCX deploys VMware Mobility Platform on both on-premises and Google Cloud VMware Engine environments. The VMware Mobility Platform contains a virtual ESXi host with a VMware Virtual Processor. Once connectivity is established between the two virtual ESXi hosts — one in on-premises and another in Google Cloud — virtual machines are migrated from on-premises ESXi hosts to the virtual ESXi host on-premises, and then to the virtual ESXi host in Google Cloud VMware Engine followed by migration to the Google Cloud VMware Engine private cloud ESXi hosts. VMware Mobility Platform eliminates vMotion CPU compatibility issues between on-premises and Google Cloud VMware Engine ESXi hosts.

The following diagram depicts the VMware products and components involved in the migration of virtual machines from on-premises to Google Cloud VMware Engine.

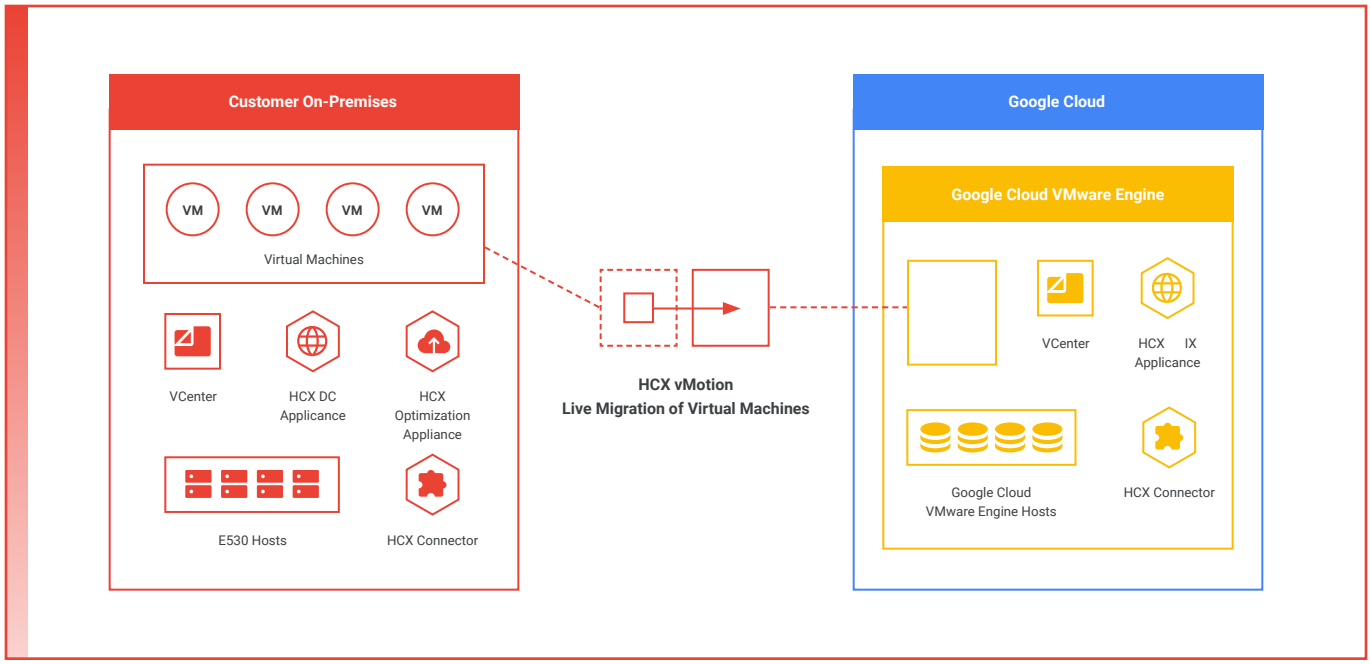


**Figure 2: VMware components (HCX and vSphere) required for virtual machine migration**

## Virtual machine migration test procedure

To perform virtual machine migration benchmarking, a test plan and procedure were developed. Both test plan and procedure were developed to mimic the real-world experience, the details of which are explained in the following sections:

## Virtual machine types

Three types of virtual machines were chosen for the performance benchmarking:

**Windows**

**Ubuntu**

**TinyCore Linux**

Though virtual machines containing RedHat or CentOS Operating System are more common than Ubuntu, Ubuntu was chosen as a pre-packaged template was already available.

**Windows virtual machines**
The Windows virtual machine type was selected for this white paper as Windows is a widely used operating system by most customers. Also, the Windows operating system is utilized for critical IT servers including Active Directory, DNS, File Servers, Web Servers, and Database Servers (SQL). The virtual machines had the Windows Server 2019 operating system installed utilizing 4vCPUs, 16G RAM, and 90GB hard disk. The virtual machines were cloned from a VMware template using the vCenter "Clone virtual machine from Template" wizard.

**Ubuntu virtual machines**
The Ubuntu virtual machine type was selected to represent the Linux Operating System in the testing and benchmarking process. These virtual machines contained Ubuntu 20.04 Operating System utilizing 4 vCPUs, 8G RAM, and 32GB hard disk. The virtual machines were created from a VMware template using the vCenter "Clone virtual machine from Template" wizard.

**TinyCore virtual machines**
The TinyCore virtual machine type was not chosen to represent real-world workloads, but rather to test the end-to-end connectivity and the volume of virtual machine migration from on-premises data centers to Google Cloud and vice versa. These virtual machines contained the TinyCore 11.1 operating system utilizing 0.1 vCPUs, 512MB RAM, and 256MB hard disk. The TinyCore virtual machines were created using clTopus. The clTopus tool has a pre-packaged and hardened VM containing the TinyCore operating system. clTopus also offers a drag-n-drop wizard to deploy these virtual machines in minutes.

## Virtual machine waves

Batches (or waves) of virtual machines where each batch contains 50 or 100 virtual machines were prepared to perform the virtual machine migration benchmark test. Wherever a batch of 100 virtual machines was included in a test case, 50 were from Windows and the other 50 from TinyCore. A higher number of virtual machine migrations couldn't be tested due to the limitation of resources in the on-premises lab.

## Virtual machine migration test scenarios

Several test scenarios were chosen to represent real-world scenarios of virtual machine migration from on-premises data centers to Cloud environments.

**HCX Bulk:** This approach allows cold migration of virtual machines between on-premises data centers and Google Cloud. HCX Bulk migration performs virtual machine migration much faster than the hot (or live) migration using HCX vMotion. This approach may be useful for customers that are able to shut down the virtual machines for the purpose of migration to Google Cloud VMware Engine. HCX Bulk migration also allows users to schedule a group of virtual machines to be migrated over a scheduled date and time.

**HCX vMotion:** This approach allows hot (or live) migration of virtual machines between on-premises data centers and Google Cloud. Live migration using HCX vMotion eliminates the issues caused by shutting down virtual machines. However, it increases the time needed to perform the migration.

**vMotion:** This approach allows hot (or live) migration of virtual machines between on-premises data centers and Google Cloud without using HCX. vMotion is a feature of vCenter and allows cross-vCenter migration of virtual machines. vMotion is suitable for use cases where ESXi host on both on-premises and Google Cloud VMware Engine are vMotion-compatible. Furthermore, vMotion is suitable for use cases where organizations don't want the overhead of configuring HCX on-premises.

**On-premises to Google Cloud:** This is the most common scenario hence applicable to all customers that are migrating workloads into Google Cloud.

**Google Cloud to on-premises:** This is not a common scenario. However, this test scenario represents use cases where virtual machines need to be migrated back from Google Cloud to on-premises data centers for troubleshooting.

## Virtual machine migration test process

Prior to performing the benchmark testing and monitoring the results, several end-to-end connectivity tests were performed to identify any bottlenecks in the system. There are several connection points where bottlenecks can occur, including between physical servers to edge devices in an on-premises data center, on-premises edge devices to Megaport physical port (1GB or 10GB), and Megaport port to Google Cloud via a Virtual Cross Connect (varies based on selected VXC data rate). Performance benchmarking experts with domain skills were engaged to remove all bottlenecks and system configuration issues prior to running the test scenarios.

The on-premises environment was dedicated only for benchmark testing for a few weeks to ensure that the results are not affected by operations outside of the virtual machine migration. Furthermore, the health of vCenter on the on-premises data center and Google Cloud VMware Engine were continuously monitored during the entirety of the virtual machine migration to ensure that the results were not affected by the sub-optimal health of vCenters. In addition, bandwidth utilization in vCenter and Google Interconnect was continuously monitored during each test run.

Several other precautionary measures were taken to ensure that the test results are not affected by external factors. For example, all of the virtual machines were pre-configured and stored inside the same storage (VSAN) for all test scenarios. Furthermore, all dependencies to external connectivity like VPNs were removed.

An attempt was made to perform at least two tests for each test scenario to evaluate the consistency of the results. During testing and benchmarking, the results of some tests were discarded as they were not consistent with the other test results. For such scenarios, it was assumed that external factors affected the test results.

The following table shows the test results between the on-premises data center and Google Cloud VMware Engine using various migration methods:

| Test # | Migration Method | Direction | VM Type | # VMs | Run 1 Time (hh:mm) | Run 2 Time (hh:mm) | Average Throughput vCenter (Mbps) |
|---|---|---|---|---|---|---|---|
| 1 | HCX Bulk | On-prem -> GCP | Windows VM<br>TinyCore VM | 50<br>50 | 2:06* | 2:12* | 2183 |
| 2 | HCX Bulk | GCP -> On-prem | Windows VM<br>TinyCore VM | 50<br>50 | 2:13* | 2:08* | 2196 |
| 3 | HCX Bulk | On-prem -> GCP | Ubuntu VM | 50 | 0:57* | 0:59* | 1859 |
| 4 | HCX Bulk | GCP -> On-prem | Ubuntu VM | 50 | 0:53* | 0:52* | 1585 |
| 5 | HCX vMotion | On-prem -> GCP | Windows VM | 50 | 6:18 | 7:09 | 253 |
| 6 | HCX vMotion | GCP -> On-prem | Windows VM | 50 | 10:29 | 9:44 | 260 |
| 7 | HCX vMotion | On-prem -> GCP | Windows VM<br>TinyCore VM | 50<br>50 | 10:49 | 11:30 | 248 |
| 8 | HCX vMotion | GCP -> On-prem | Windows VM<br>TinyCore VM | 50<br>50 | 14:27 | - | N/A |
| 9 | HCX vMotion | On-prem -> GCP | Ubuntu VM | 50 | 7:21 | 6:16 | 147 |
| 10 | HCX vMotion | GCP -> On-prem | Ubuntu VM | 50 | 5:26 | 5:32 | 128 |
| 11 | HCX vMotion | On-prem -> GCP | TinyCore VM | 50 | 0:14* | - | 112 |
| 12 | HCX vMotion | On-prem -> GCP | Windows VM<br>Ubuntu VM | 50<br>50 | 11:53 | 11:09 | N/A |
| 13 | HCX vMotion | GCP -> On-prem | Windows VM<br>Ubuntu VM | 50<br>50 | 15:02 | 13:09 | 108 |
| 14 | vMotion | On-prem -> GCP | Windows VM | 50 | 1:49 | - | N/A |
| 15 | vMotion | GCP -> On-prem | Windows VM | 50 | 1:48 | - | N/A |
| 16 | vMotion | On-prem -> GCP | Windows VM<br>TinyCore VM | 50 | 2:02 | - | N/A |
| 17 | vMotion | GCP -> On-prem | Windows VM<br>TinyCore VM | 50<br>50 | 1:50 | - | N/A |

**Table 4: Test scenarios for virtual machine migration**

The test results that are marked with * were performed with an older configuration of Google Cloud VMware Engine.

N/A represents “Not Available.” For some of the benchmark data, accurate bandwidth data was not obtainable. For example, for the vMotion tests (#14 till #17), the test runs were 2 minutes or less. Hence, it was not possible to get accurate bandwidth utilization data for those tests. For two other HCX vMotion tests (#8 and #12), bandwidth utilization data couldn’t be calculated reliably.

**Note:** Though 10GB connectivity was established between on-premises and Google Cloud, only 2GB could be utilized due to a limitation from HCX.



Chapter 4

# Benefits of Google Cloud & Google Cloud VMware Engine Networking

## Dynamic routing mode

When configuring VPC networks, you can set the dynamic routing mode to match how you would like the Cloud Router of that VPC network to dynamically advertise and propagate routes, it can be set to either regional or global mode. This feature tremendously helps organizations utilize Google Cloud as it eliminates additional networking tasks like configuration of static routes. Network connectivity between on-premises and Google Cloud can be established either via Cloud VPN or Interconnect. If Cloud VPN is utilized, the dynamic routing mode feature of Google Cloud Router advertises the subnet ranges to the on-premises VPN gateway. If Interconnect is utilized, the dynamic routing mode feature of Google Cloud Router advertises the subnet ranges to the on-premises router that is configured with the Interconnect.

Note that a VPC network must be used and configured with Private Services Access to connect to Google Cloud VMware Engine.

**https://cloud.google.com/vmware-engine/docs/networking/howto-setup-private-service-access**

Google Cloud VMware Engine routes can then be imported with a VPC Peering connection, but it should be noted that they will not be automatically advertised to on-prem without creating a custom route advertisement.

The following diagram depicts the end-to-end connectivity and advertisement of VMware networks from Google Cloud to the on-premises router using Cloud Router and Cloud Interconnect.

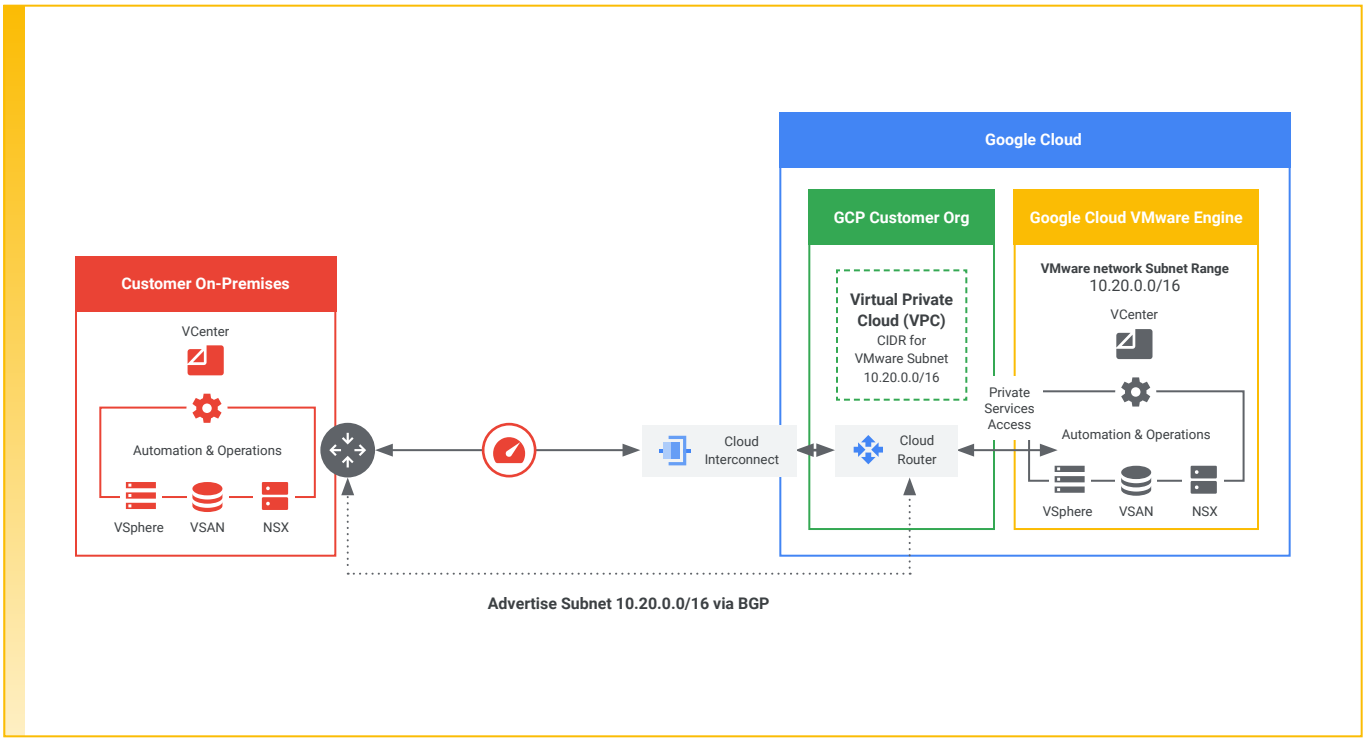


**Figure 3: Dynamic Routing Mode advertising subnet range to on-premises router**

## Google VPC network peering

Google VPC Network Peering is a capability offered by Google, similar to other CSP offerings but with the added advantage of Global VPCs, which allows workloads within a VPC to communicate with workloads from other VPCs using internal private IP addresses, thereby eliminating the need for using external or public IPs, even when crossing regions. The core advantages of Google VPC Network Peering are reduction of network latency and cost, and increased network security. When using VPC Network Peering, workloads — such as virtual machines and VMkernal adapters — residing on networks in one Google Cloud private cloud can communicate with workloads of another Google Cloud private cloud instance without needing VPN or external IP addresses, thereby simplifying network connectivity while improving end-to-end security, even if they reside in different regions.

Note that VPC peering is not transitive, and while peered networks can communicate with each other, they do not receive routes to any networks the other peer may be peered to. A VM appliance can be deployed to bridge that gap. More information on VMs with NICs in different VPCs and the available routing options can be found in the following Google documentation:

**https://cloud.google.com/vpc/docs/vpcpeering?hl=en#multiple_network_interfaces_per_instance**

The following diagram shows VPC peering between each consumer VPC and it's Google Cloud VMware Engine Private Cloud (each in a different region) where network connectivity between those VPCs is established between two separate Google Cloud instances connected to both VPC networks.

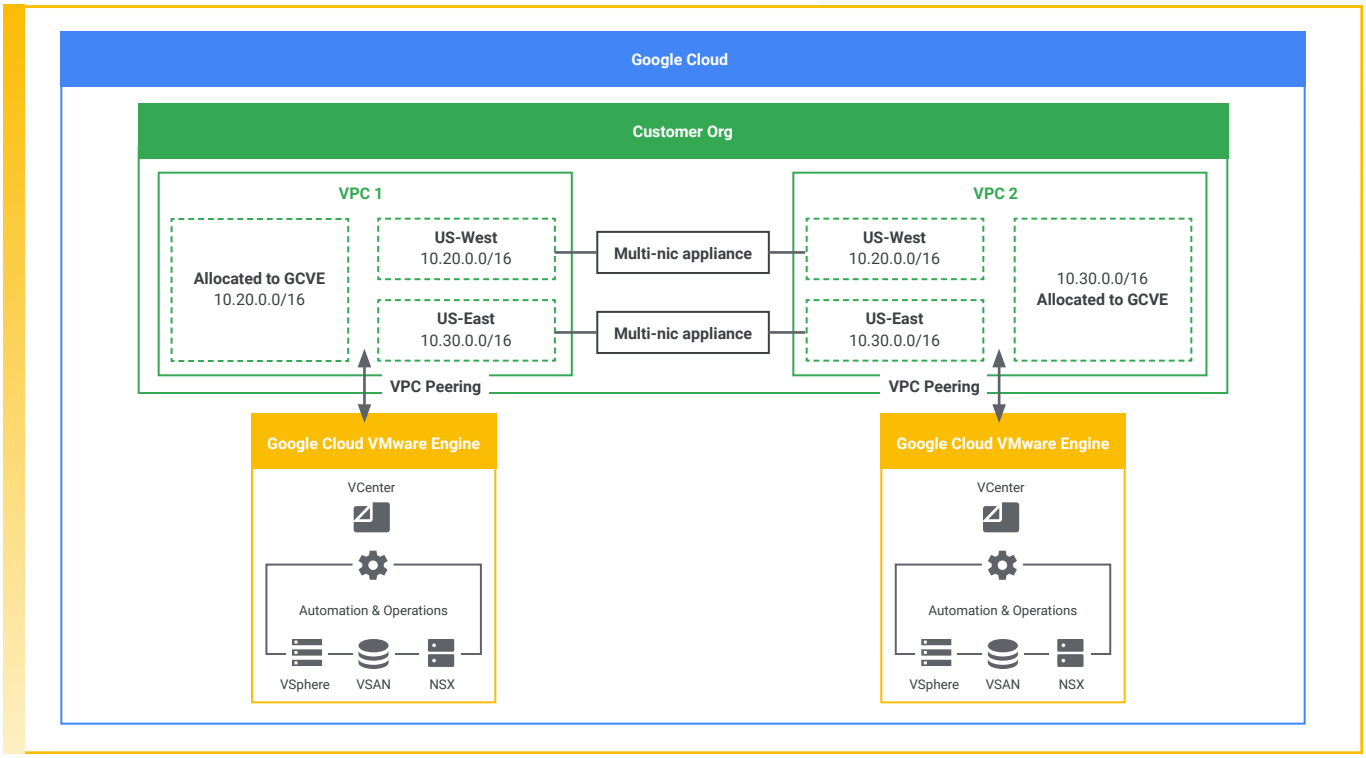


**Figure 4: VPC Network Peering between Google Cloud VMware Engine instances**

## Multi-VPC connectivity

By default, Google Cloud VMware Engine allows access to the same private cloud from different VPC networks without additional configuration to those VPCs, but for scenarios where you would have separate VPC networks to segment traffic for testing and development as an example, Multi-VPC connectivity can be useful. Multi-VPC connectivity is an extension of the concept of the VPC Network Peering exclusive to Google Cloud VMware Engine, which allows a consumer VPC to peer with multiple VPCs (3 by default per region) as shown on the diagram below. The use of Multi-VPC connectivity essentially eliminates network connectivity-related overhead and complexities from all Google Cloud VMware Engine instances within an organization. Organizations gain unprecedented value from their VMware workloads by allowing them to communicate with each other without requiring any tedious network connectivity. Organizations that have multiple data centers across various geographical regions require layers of network connectivity including VPN, edge networking, and firewalls to allow for communication among VMware workloads. Multi-VPC connectivity also lets you access the same private cloud from different GCP VPC networks without the need to change your VPC network architecture.

The value of Multi-VPC Peering is also realized in hybrid cloud scenarios where some of the VMware virtual machines are on-premises and the rest are on Google Cloud. On-premises VMware infrastructure and workloads can communicate with Google Cloud VMware infrastructure and workloads using private IPs without requiring any new network infrastructure.

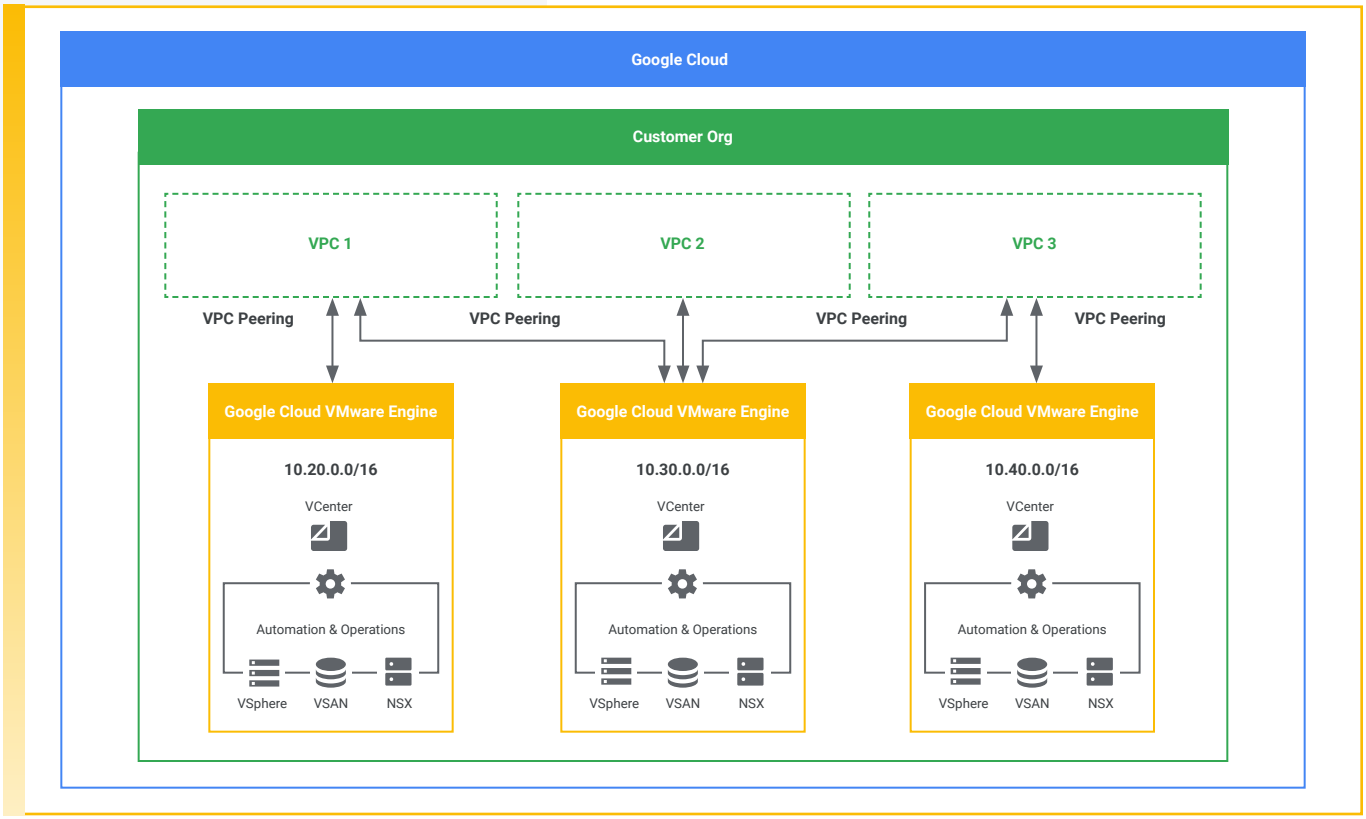


**Figure 5: Multi-VPC connectivity**

## Cloud DNS

Google Cloud DNS provides 100% availability and scales up to millions of records. Organizations with multiple geographical data centers running VMware virtual machines grapple with DNS resolution from one data center to another. This challenge goes away with the migration to Google Cloud. Google Cloud DNS performs DNS resolution across multiple GVCE instances. When a new Google Cloud VMware Engine instance is brought online, DNS resolution can be performed for the new resources with minimal configurations. FQDN is an essential component of VMware infrastructure services as the use of IP addresses is not recommended for system-to-system integration within VMware. Cloud DNS allows FQDN access to VMware infrastructure in Google Cloud from anywhere, including on-premises.

Chapter 5

# Conclusion

Google Cloud VMware Engine provides an optimal experience and solution when delivering a migration cloud solution that can adapt to a variety of customer environments and needs. SPJ Solution's use case scenario demonstrated the ease of adoption and timeframe when presented with the challenge of migrating VMs to and from your on-prem datacenter in a DR or migration scenario.

## About the Primary Authors

**Robert Castruita**
**VMware Solutions Engineer,**
**Center of Excellence - Google Cloud**

Robert Castruita has over 15 years of IT work experience, including extensive experience in a large enterprise network and systems design implementation for maintaining critical world-wide communications systems. Robert has had a colorful career supporting the Creative Artists Agency, Dell Inc., Booz Allen Hamilton, VMware and most recently Google. His overall architectural expertise translated directly into client success through the development of engineering and design implementation. He has managed and spearheaded efforts in the planning and development of Cloud IT architecture and network-based designs. His previous employment consisted of a role in Sr. Virtualization management and troubleshooting at Dell Inc., his role entailed the maintenance and proof of concept of both a client's physical and virtual servers' storage platforms. Robert currently possesses an MCSA Server 2012, VCP, F5 Administration and Security+ certification. Robert holds a Bachelor's Degree from Cal State Long Beach and Multi-Cloud Certifications.

**Albert Colas Prunera**
**Networking Specialist,**
**Google Cloud**

Albert Colas Prunera has over 14 years of high- tech experience in the areas of infrastructure, networking and security, application development, sales, training, authoring and public speaking. Albert has had a global career starting in Europe at Cisco and later on in the United States where he first joined a Cloud Networking startup, Hewlett Packard Enterprise, and most recently Google. Albert has been working closely with leading cloud service providers for the last decade to develop and deliver multi-cloud operations, including networking and security. He previously served as a Distinguished Engineer at Hewlett Packard Enterprise (HPE), his role entailed leading program strategy, roadmap development and planning as well as representing the company as the subject matter expert for the networking portfolio. Albert holds an MBA Degree from Santa Clara University, Leavey School of Business and a Master's and Bachelor's degree from La Salle (Ramon Llull University) Barcelona. Albert also holds multiple industry certifications from Cisco, CWNP, HP, Microsoft and Google, where he was part of a small team to launch the new Professional Cloud Network Engineer certification exam.

## About SPJ Solutions

SPJ Solutions Inc is the manufacturer of cITopus — a cloud automation, management, and migration tool — and a system integrator specializing in VMware, GCP, and other products. SPJ Solutions utilizes three components to deliver consulting services: 1) cITopus to perform automation, management, and migrations. 2) Experts specializing in VMware and Google Cloud. 3) Practice development and experience from previous engagements. SPJ Solutions is a partner of both VMware and Google Cloud.

## About the Secondary Authors

**Sudhansu Pati,** VCDX#255, CCIE#62902, OCM has over 25 years of experience in the areas of infrastructure, networking and security, application development, sales and delivery, training, authoring and public speaking. Sudhansu has been working closely with leading cloud providers for the last 15 years to develop and deliver multi-cloud operations, including automation, workload mobility, networking and security, and governance. Sudhansu is the Chief Technology Officer and Co-Founder of SPJ Solutions, a leading organization specializing in VMware, and Cloud technologies. Sudhansu is also the architect of cITopus, a software application, that allows automation and management of VMware products including NSX, VMC, vSphere, VCF, AVI, VRNI and migration of virtual machines between on-premises and cloud environments. Sudhansu is the primary author of the white paper.

**Jacob Bagwell,** VCIX-DTM, GCP, has tested the end-to-end connectivity between on-Premises and Google Cloud VMware Engine. Jacob also led performance benchmarking tasks — including gathering and analyzing the performance results from vCenter and Megaport — for each of the test cases presented on this white paper. Jacob works for SPJ Solutions.

**Sonny Pakdel,** VCIX-NV, VCI, has developed the test environments — including the virtual machines ,— and performed all of the test runs presented on this white paper. Sonny has over 20 years of experience in the IT field with a background in networking, security, Microsoft technology, and Linux administration. He has over 12 years of experience working with VMware products such as the vSphere, SRM, vROPS, and vCloud Director. Sonny works for SPJ Solutions.

**Matt Elliott,** VMware, has provided significant support to the performance benchmarking operation. Matt provided the VMware lab for the development of the white paper, installed and configured all VMware products needed for the benchmarking including vSphere, HCX and NSX, provided secure access to SPJ Solutions resources into the VMware lab, and provided support. Matt also architected and configured end-to-end connectivity between VMware lab and Google Cloud VMware Engine to be able to perform the virtual machine migration. Matt worked with Megaport to configure the 10GB port and connectivity.

**Paul McGuinness,** Megaport, has provided architectural guidance on the end-to-end connectivity between on-premises and Google Cloud using Megaport. Paul has developed the sections involving Megaport in this white paper. With more than 25 years of experience, Paul delivers an extensive range of expertise in solution design and technical consulting for customers and partners of all sizes.