

# Migrating On-Premise Workloads to Google Cloud Platform

## Virtual Machine Performance Benchmarks



WHITE PAPER

February 28, 2022

# Migrating On-Premises Workloads to Google Cloud Platform - Virtual Machine Migration Performance Benchmarks

---

## Table of Contents

1. Executive Summary
2. On-premises and Google Private Cloud Connectivity
  - a. Overview
  - b. End-to-End Connectivity using Megaport
  - c. VPN Connectivity
3. Virtual Machine Migration between on-premises and GCP
  - a. Virtual Machine Migration Tool
  - b. Virtual Machine Migration Test Procedure
    - i. Virtual Machine Types
    - ii. Virtual Machine Waves
    - iii. Virtual Machine Migration Test Scenarios
  - c. Virtual Machine Migration Test Results
4. Benefits of Google Cloud & Google Cloud VMware Engine Networking
  - a. Dynamic Routing Mode
  - b. Google VPC network peering
  - c. Multi-VPC connectivity
  - d. Cloud DNS
5. Conclusion

---

## Executive Summary

This white paper provides architecture of end-to-end connectivity between an on-premises data center and Google Cloud Platform. This white paper also produces performance benchmarks of virtual machine migration from on-premises to Google Cloud Platform. It is best suited for organizations that are looking to migrate workloads to Google Cloud Platform. The specific audience that will benefit from this white paper are cloud and data center engineers and technical managers.

There are **three** key objectives of this white paper.

- Analyze the performance of migration of virtual machines from an on-premises data center to Google Cloud Platform consisting of Google Cloud VMware Engine private cloud.
- Provide detailed technical guidance on establishing end-to-end connectivity between on-premises data center and Google Cloud Platform.
- Explain the benefits the Google Cloud Platform 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 Platform 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. This architecture presented in this white paper offers optimized workload migration from on-premises to Google Cloud Platform. 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 Platform. Third, organizations can learn the benefits that Google Cloud Platform offers to improve the service delivery of business applications.

Organizations either in the evaluation phase or in the execution phase of migration of workloads from on-premises to Google Cloud Platform 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 outage of business applications. The cost of migration includes the cost of establishing end-to-end connectivity between on-premises to Google Cloud Platform, and the cost of 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 outage of applications 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 Platform 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 Platform 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 doesn’t affect business operations.

VMware provided an on-premises lab environment to test workload migration to Google Cloud Platform. The lab environment was connected to Google Cloud Platform 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 Platform, customers can utilize the power of Google Cloud Platform, 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 Platform.

---

# On-Premises and Google Cloud Platform Connectivity

## Overview

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

- On-Premises CPU: Intel(R) Xeon(R) CPU E5-2630 v4
- On-Premise ESXi Host: PowerEdge R730xd

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**

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.

- Google Cloud VMware Engine CPU: Intel(R) Xeon(R) Gold 6240 CPU
- Google Cloud VMware Engine Host: PowerEdge R640

<b>Product</b>	<b>Version</b>	<b>Build</b>
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***

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

- Google Cloud VMware Engine CPU: Intel(R) Xeon(R) Gold 6240 CPU @ 2.60GHz
- Google Cloud VMware Engine Host: PowerEdge R640

<b>Product</b>	<b>Version</b>	<b>Build</b>
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***

## **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 VMware networks on-premises data center and Google Cloud Platform.

Megaport provides a portal to configure the connectivity between on-premises and Google Cloud Platform. 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.

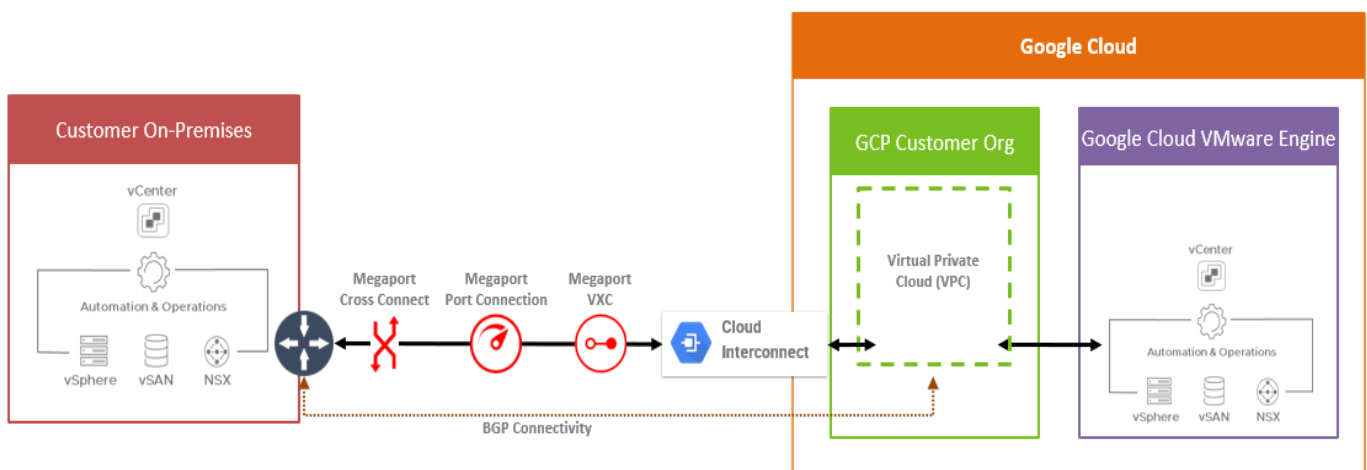
The next step is the creation of a cloud router using Google Cloud Console's Hybrid Connectivity link. Under Hybrid Connectivity, choose Interconnect. Once the Google Partner Interconnect is created, a pairing key is generated, which must be entered into Megaport's Google Connectivity page.

After that, a Virtual Cross Connect (VXC) must be created 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.

1. Either click on the Google Cloud tile or "Add Connection.:"
2. Select "Google Cloud"
3. Put in the pairing key from Google Cloud Partner Interconnect.
4. Select the country and on-ramp location where you want to connect to Google Cloud, and click "Next."
5. Add a name for the VXC
6. Select your rate limit
7. Select a VLAN corresponding to A-End on-premises Port VLAN.
8. Next, add VXC and click "Order." Now we see this is in a deployable state.
9. 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 that need to configure end-to-end connectivity between the on-premises data center and Google Cloud Platform.



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

## VPN Connectivity

Google offers Cloud VPN to connect on-premises data centers and Google Cloud Platform using internet-based IPSec Tunnels. Cloud VPN is useful for low volume data transmission between on-premises and Google Cloud Platform hence 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 Platform VMware networks via the internet. L3 connectivity is sufficient for the migration of VMs from on-premises to Google Cloud Platform. 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.

---

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

### Virtual Machine Migration Test Procedure

To perform benchmarking of virtual machine migration, a test plan, and a test procedure was 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:

- Windows
- Ubuntu
- TinyCore Linux

Though virtual machines containing RedHat and/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, 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 Windows 2019 Server Operating System installed utilizing 4vCPUs, 16G RAM, and 90G



Hard Disk. The virtual machines were cloned from a VMware template using 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 32G Hard Disk. The virtual machines were created from a VMware template using vCenter “Clone virtual machine from Template” wizard.

## ***TinyCore Virtual Machines***

The TinyCore virtual machine type was not chosen to represent real-world workload, rather to test the end-to-end connectivity and to test the volume of virtual machine migration from on-premises data center to Google Cloud Platform and vice versa. These virtual machines contained TinyCore 11.1 Operating System utilizing 0.1 vCPUs, 512MB RAM, and 256MB Hard Disk. The TinyCore virtual machines were created using cITopus. cITopus Utility Appliance has a pre-packaged and hardened VM containing TinyCore Operating System. cITopus 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 benchmark test of virtual machine migration. Whenever 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 migration 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 as follows 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 Platform. 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 Platform. 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 Platform 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 Platform:** This is the most common scenario hence applicable to all customers that are migrating workloads into Google Cloud Platform.

**Google Cloud Platform 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 Platform to on-premises data center 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 device in an on-premises data center, edge device on-premises to Megaport physical port (1G or 10G) and Megaport port to Google Cloud Platform 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 i.e. VPN was 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.

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 TinyCore VM	50 50	2:06*	2:12*	2183
2	HCX Bulk	GCP -> On-prem	Windows VM TinyCore VM	50 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 TinyCore VM	50 50	10:49	11:30	248
8	HCX vMotion	GCP -> On-prem	Windows VM TinyCore VM	50 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 Ubuntu VM	50 50	11:53	11:09	N/A
13	HCX vMotion	GCP -> On-prem	Windows VM	50	15:02	13:09	108

			Ubuntu VM	50			
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 TinyCore VM	50 50	2:02		N/A
17	vMotion	GCP -> On-prem	Windows VM TinyCoreVM	50 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 10G connectivity was established between on-premises and Google Cloud, only 2G could be utilized due to a limitation from HCX.

## Benefits of Google Cloud Networking

### Dynamic Routing Mode

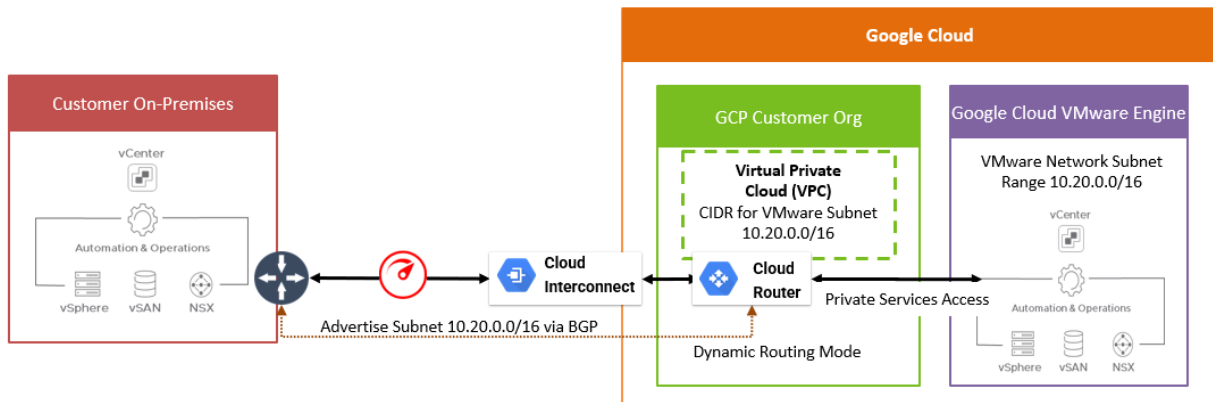
When configuring VPC networks, you can set a dynamic routing mode to match how you would like the Cloud Router of that VPC network to dynamically advertise and propagate routes throughout the region, or globally in all VPC networks. This feature tremendously helps organizations utilize Google Cloud Platform as it eliminates additional networking tasks i.e. 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 export the

VMware networks from the Google Cloud VMware Engine to the Google Cloud Router, details of which can be found from the following Google documentation.

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

These 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 Platform 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

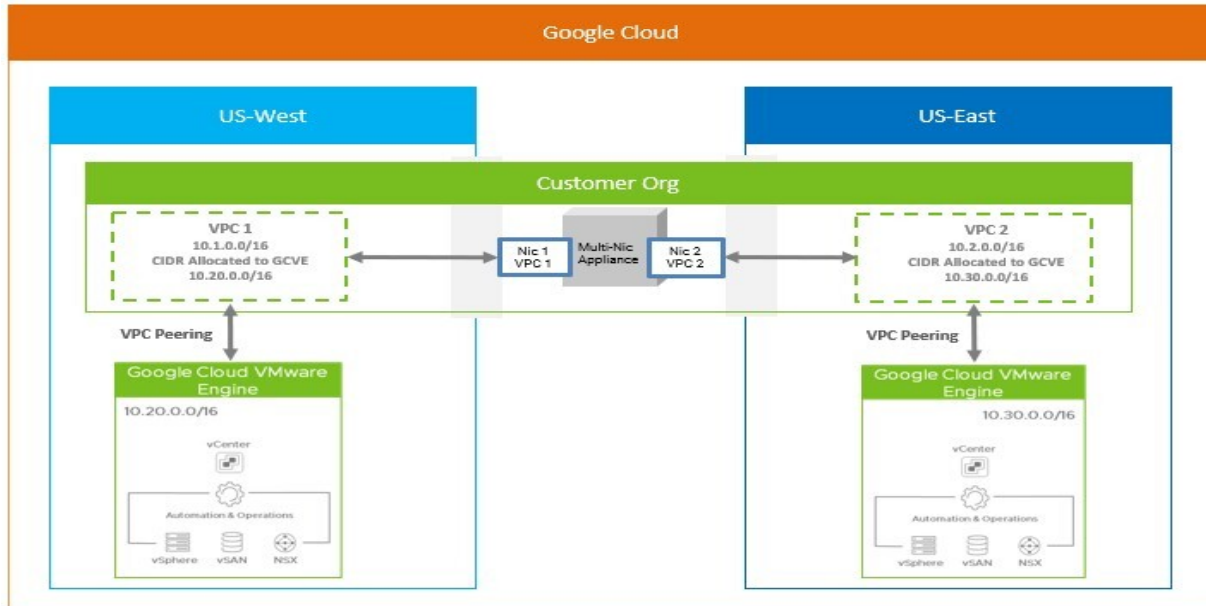
Google *VPC Network Peering* is a capability offered by Google, similar to other CSP offerings but with the added advantage of Global VPCs, which allows workloads within a VPC to communicate with workloads from other VPCs using internal private IP addresses thereby eliminating the need for using external or public IPs, even when crossing regions. Core advantages of Google VPC Network Peering is reduction of network latency and cost and increase of network security. Using VPC Network Peering, workloads (such as virtual machines and VMkernel adapters) residing on networks in one Google Cloud Platform private cloud can communicate with workloads of another Google Cloud Platform 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 other networks the other peer may be peered to. A VM appliance can be deployed to peer between networks, utilizing source or destination-based routing policies to allow

traffic to traverse across networks or even regions. 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/vpc-peering?hl=en#multiple\\_network\\_interfaces\\_per\\_instance](https://cloud.google.com/vpc/docs/vpc-peering?hl=en#multiple_network_interfaces_per_instance)

The following diagram shows VPC peering between two regions (US-West and US-East) where network connectivity is established between two separate Google Cloud Platform instances using Google VPC Network Peering and an additional appliance with connectivity to both VPC networks.



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

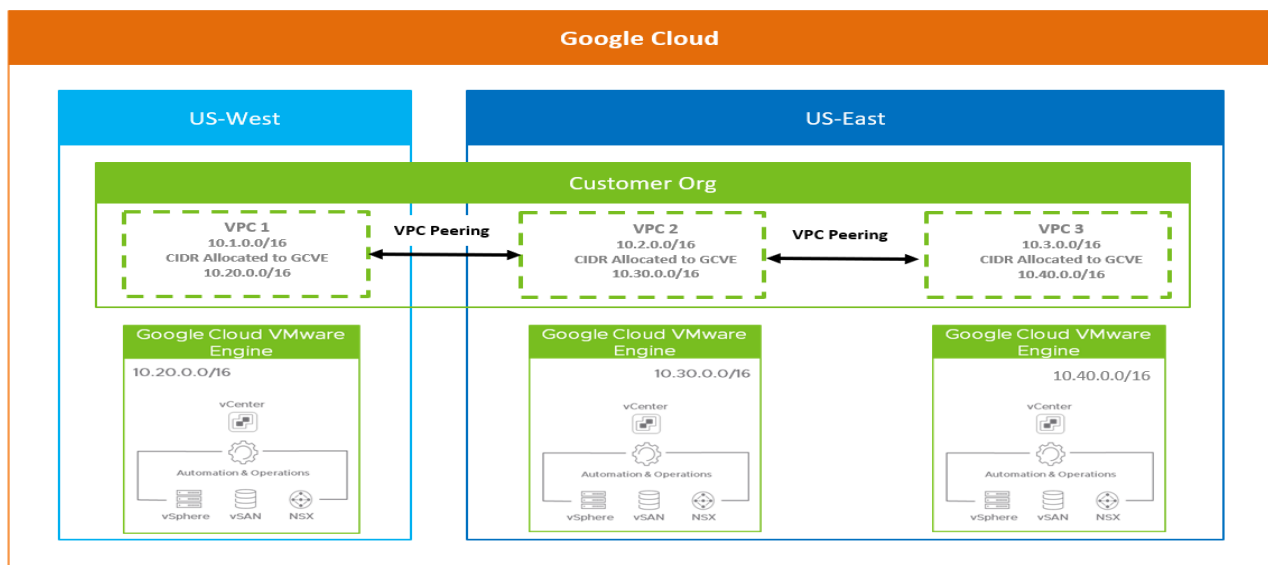
## Multi-VPC connectivity

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 of 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. The use of *Multi-VPC connectivity* allows direct private IP access from one Google Cloud Platform instance to another or more without added complexity and cost. 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 Platform. on-premises VMware infrastructure and workloads can communicate with Google Cloud Platform VMware infrastructure and workloads using private IPs without requiring any new network infrastructure.



**Figure 5: Multi-VPC connectivity**

An alternative to Multi-VPC connectivity is the use of Global VPC, where a VPC spans across multiple regions without communicating across the public internet. Global VPC allows on-premises data centers to use a shared connection to gain access to resources across all regions.

## 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 the Google Cloud Platform. 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 Platform from anywhere, including on-premises.

---

## 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 SPJ Solutions

SPJ Solutions INC, is the manufacturer of cITopus, a cloud automation, management and migration tool, as well as 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 GCP.

### About the Authors



**Sudhansu Pati**, VCDX#255, CCIE#62902, OCM has over 25 years of experience in infrastructure, networking and security, application development, sales/delivery, training, authoring, and public speaking. Sudhansu has worked 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. He is the CTO and Co-Founder of SPJ Solutions and 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 this 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 10G port and connectivity.

**Paul McGuinness**, Megaport, has provided architectural guidance on the end-to-end connectivity between on-premises and Google Cloud Platform 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.