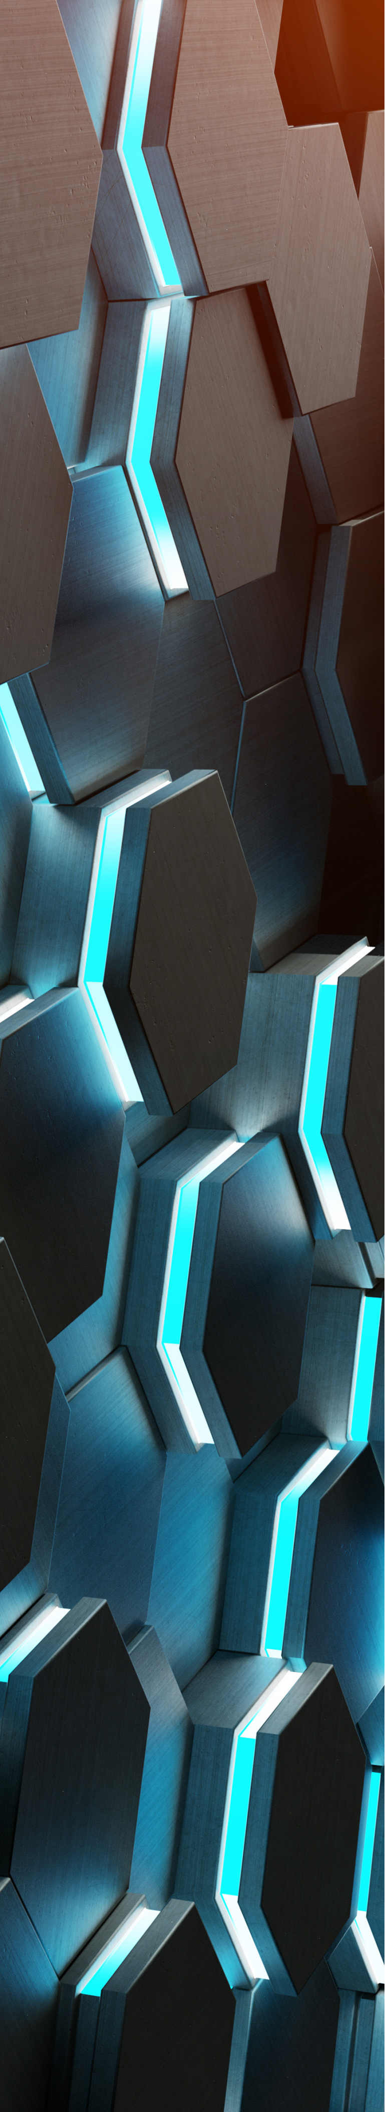

bluealtair
Driving Digital Success

Overcoming Challenges in Apigee Hybrid Installations

By Brajesh De

Global lead – API Management and Integration;
Digital Application Development, Blue Altair

Overview

Apigee is a leading API Management platform from Google. It enables organizations to develop, manage, and publish their products and services as APIs. Its rich analytics dashboard provides visibility into their API usage and adoption. To help businesses meet diverse requirements, Apigee provides two deployment models: Cloud or Software as a Services (SaaS) and Hybrid.

While the popularity of the Apigee SaaS model continues to grow, Apigee's Hybrid model is best suited for organizations that are bound by more restrictive business requirements and regulations as it allows for the installation of the Apigee runtime components on an infrastructure and network of their choice. However, given the multitude of available networks and infrastructure available in the marketplace, the nuances of installing a hybrid deployment model can be tricky and sometimes come with unforeseen challenges.

In this blog, we'll take a closer look at the deployment architecture for Apigee Hybrid in an on-premises environment; highlight common challenges that occur during the installation process; and explore best practices to ensure a seamless deployment.

Apigee SaaS vs Hybrid Comparison

First, let's explore why implementing an Apigee Hybrid deployment model over the SaaS model may be the best choice for your business. For many organizations, Apigee's SaaS deployment model delivers optimal results by utilizing a cloud environment which can reduce overall maintenance and operational cost. However, a full review of business and regulatory requirements should be conducted to determine if Apigee's Hybrid deployment model is required. Generally speaking, organizations that need to have all of their data within the enterprise firewall or within their geo-political boundary may select Apigee Hybrid as it allows businesses to install the Apigee runtime components on an infrastructure and network of their choice. With the Hybrid model, organizations can also leverage the SaaS instance of Apigee as needed to set up Apigee only on Google Cloud in geographic regions where Google has a presence. However, in this setup, customers have lesser control over data residency which could be problematic if they must comply with certain government and industry-specific compliance regulations pertaining to data residency, security, ad proximity to backend applications.

Installing Apigee Hybrid

Apigee Hybrid is a preferred option for businesses who must retain their API data to meet various regulatory requirements. With Apigee Hybrid, clients' can retain their data on their enterprise network or in a cloud platform and region of their choice. Platforms currently supported by Apigee Hybrid include Azure Kubernetes Service (AKS) on Azure, Elastic Kubernetes Service (EKS) on AWS, Google Kubernetes Engine (GKE) on GCP, and Anthos for Openshift and on-premise bare metal servers. The management plane components are hosted and managed by Google itself on Google Cloud.

Setting Up Apigee Hybrid

When installing Apigee Hybrid, start by setting up an Anthos cluster on bare metal servers inside your enterprise data center and then deploy Apigee Hybrid runtime components on Anthos. Anthos lets you run Kubernetes clusters for application deployment directly on your own hardware. This comes with the advantage of delivering the best performance and flexibility with compatible security. By leveraging the multi-cluster deployment model of Anthos, you can manage a fleet of clusters (user clusters) from a centralized cluster called the Admin Cluster, as shown in Figure 1.

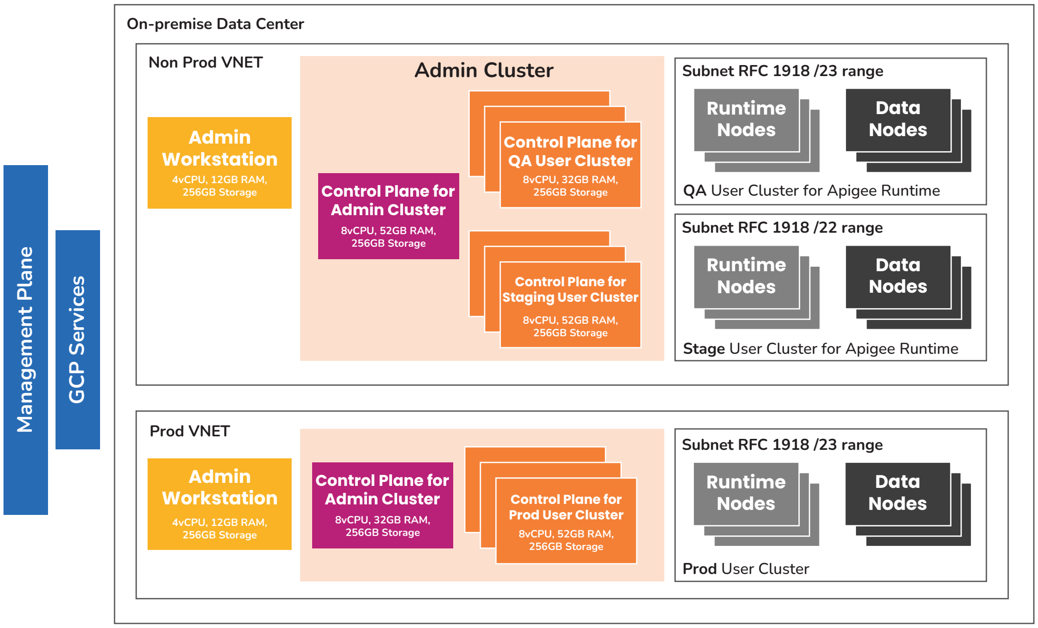
Management Plane
GCP Services
On-premise Data Center
Non Prod VNET
Admin Workstation
4vCPU, 12GB RAM, 256GB Storage
Admin Cluster
Control Plane for Admin Cluster
8vCPU, 52GB RAM, 256GB Storage
Control Plane for QA User Cluster
8vCPU, 32GB RAM, 256GB Storage
Control Plane for Staging User Cluster
8vCPU, 52GB RAM, 256GB Storage
Subnet RFC 1918 /23 range
Runtime Nodes
Data Nodes
QA User Cluster for Apigee Runtime
Subnet RFC 1918 /22 range
Runtime Nodes
Data Nodes
Stage User Cluster for Apigee Runtime
Prod VNET
Admin Workstation
4vCPU, 12GB RAM, 256GB Storage
Control Plane for Admin Cluster
8vCPU, 32GB RAM, 256GB Storage
Control Plane for Prod User Cluster
8vCPU, 52GB RAM, 256GB Storage
Subnet RFC 1918 /23 range
Runtime Nodes
Data Nodes
Prod User Cluster

Figure 1: On-premise Data Center Components for Apigee Hybrid

The Apigee application is deployed as pods and services that run on worker nodes of Anthos user cluster. It is recommended that you have at least three nodes for Apigee runtime and three nodes for Cassandra datastore. The workloads in each user cluster will be created and managed by its own control plane that will be part of the Admin Cluster of the Anthos setup. Since Kubernetes control plane uses etcd for managing its configuration data, it is recommended that you have a minimum of three nodes for the control plane setup to meet the resiliency and quorum requirements. The Admin Workstation node hosts Command Line Interface (CLI) tools and configuration files for easy provision of clusters during installation and interaction with them after installation.

Installation Challenges and Solutions

The installation process of Anthos requires connectivity to Google servers to download the required container images and run pre-flight checks. Many organizations route outbound internet traffic through their proxy servers that mask the original IPs for security reasons, making it seem like the request originated from the proxy server.

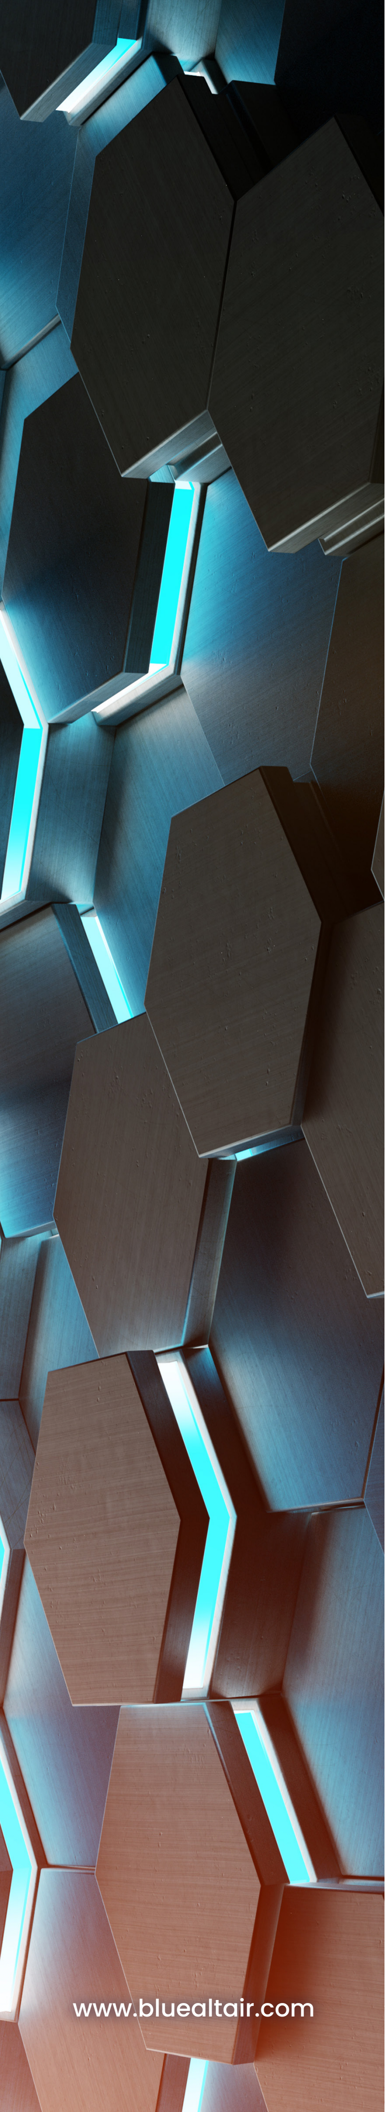

www.bluealtair.com

A man in the middle (MITM) proxy provides higher levels of security by terminating all SSL requests at the proxy server. Countries like China have additional restrictions imposed by their Great Firewall which require special network setup to route traffic for meta data information to Google Cloud.

When you set up the Anthos cluster behind MITM proxy on bare metal servers for a client, you may encounter some of these challenges:

1. Cluster creation failed with certificate validation errors while trying to pull image from gcr.io via MITM proxy.
2. Conflict between the containerization software and versions supported by Red Hat's Enterprise Linux (RHEL) operating system of the Admin Workstation and the software and version required by ansible scripts for the pre-flight check. RHEL supports Podman while the pre-flight check needs Docker > 19.03 installed.

The challenge with MITM proxy could be addressed by updating the configurations on the proxy server such that it bypasses the SSL checks while connecting to the Google container registry. Alternatively, you can set up a local repo of container images required for the installation. The choice of solution would, however, depend on an organization's security policies.

The problem of container software version mismatch can be addressed by updating the operating system on the Admin Workstation to the latest version of RHEL and then installing docker. This will help resolve some of the challenges being faced, allowing you to proceed with the installation.

Other Factors for Successful Apigee Hybrid Installation

The on-premise setup of Apigee Hybrid platform is non-trivial. Therefore, it needs proper planning to ensure that you meet all the recommended pre-requisites for hardware, network connectivity and for dependency on software component versions.

You will also need to have seamless coordination between and across several teams spanning infrastructure, network, firewall, security and your Apigee Platform setup team along with the support from senior executives.

About the Author

Brajesh has 25 years of experience in delivering large-scale digital transformation projects. He has architected and designed several distributed, cloud native, highly scalable and secure applications leveraging API-led architecture on cloud. He holds two patents across many countries for his work in API Assessment and Data Intelligence

Brajesh is the author of **API Management: An Architect's Guide to Developing and Managing APIs for Your Organization.**

bluealtair
Driving Digital Success