



Reference architecture:

# Anthos hybrid environment (part 1) - design prerequisites and considerations

[Part 2 - Implementation details](#)

August 2023

---



## Table of contents

<b>Overview</b>	<b>4</b>
<b>Architecture</b>	<b>6</b>
<b>Anthos components</b>	<b>8</b>
<b>Reference deployments</b>	<b>9</b>
Global view	9
Regional view	12
Clusters and environments	15
Observability	18
Configuration management	18
Security	19
<b>Design prerequisites</b>	<b>21</b>
Regions and sites	21
Google Cloud setup	22
On-premises setup	23
Site services	23
Compute requirements	24
Operating system	25
Storage	25
Networking	26
IP addresses	26
Internet connectivity	26
Intercluster connectivity	26
Intracluster connectivity	27
Basic network	27
Load balancer	28
On-premises to Google Cloud connectivity	28
<b>Design considerations</b>	<b>29</b>
Hybrid deployment	29
Availability	30
Network security and management	32
Operations	34
Scale and limits	35
Cluster and node limits	35
Load balancing limits	36
Anthos Config Management limits	36
Observability limits	36
Multi-tenancy	37
Observability	39



Multi-cluster management	39
Gitops-based configuration management	39
Application availability	40
<b>References</b>	<b>41</b>



## Overview

Organizations that embrace cloud-first technologies like containers, container orchestration, and service meshes, often reach a point where they need more than a single Kubernetes cluster. Many organizations that use Google Cloud also want to run workloads in their own data centers, factory floors, retail stores, or even in other public clouds.

However, operating multiple Kubernetes clusters has its own difficulty and overhead in terms of consistent configuration, security, and management. For example, manually configuring one Kubernetes cluster at a time creates risks, and it can be challenging to see exactly where errors are happening.

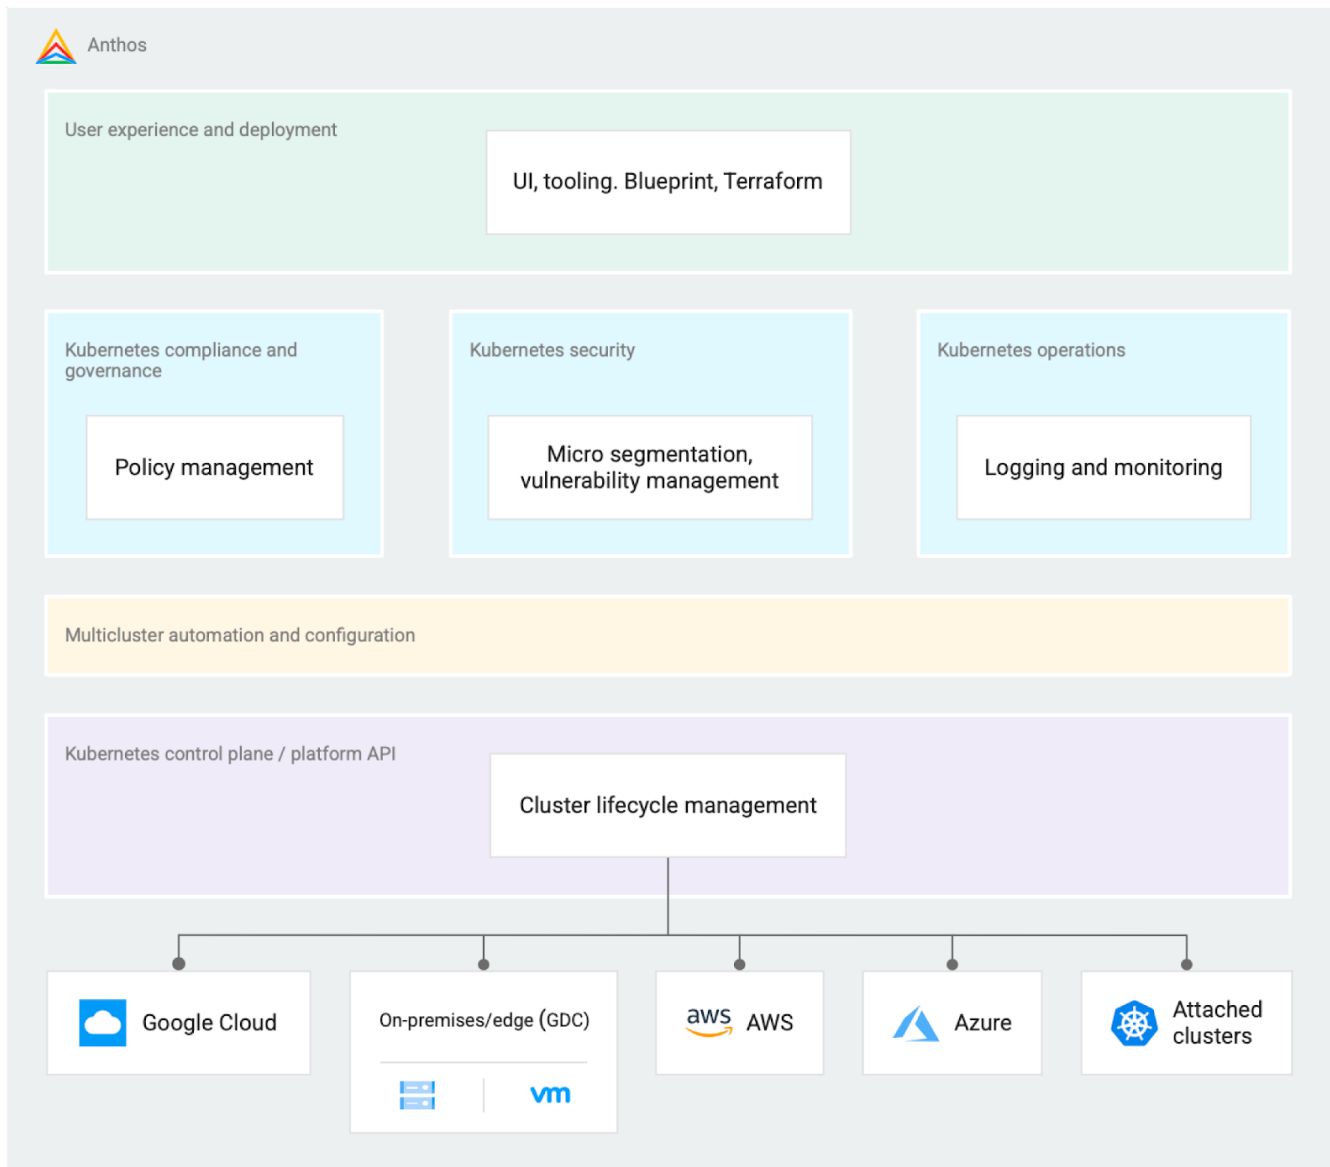
Anthos is Google's cloud-centric container platform for running modern apps anywhere consistently at scale. Anthos can help organizations by providing a consistent platform that lets them:

- Modernize applications and infrastructure in-place.
- Create a unified cloud operating model (single pane of glass) to create, update, and optimize container clusters wherever they are.
- Scale large multi-cluster applications as *fleets* - logical groupings of similar environments - with consistent security, configuration, and service management.
- Enforce consistent governance and security from a unified control plane.

Anthos helps you increase operational consistency in governance and security and developer velocity while reducing cost, deployment risk, and operational complexity. Specifically, Anthos helps with the following areas:

- Customers who want cloud-like experience on-premises or are looking for a unified solution while migrating their applications to cloud (***Anthos hybrid environment***).
- Google Cloud customers who want to better manage their containerized applications (***Anthos on Google Cloud***).
- Customers who want to solve multicloud complexity with a consistent governance, operations, and security posture (***Anthos Multi-Cloud***).

The following diagram shows a high-level overview of Anthos in a hybrid environment. Anthos can help with Kubernetes compliance and governance, security, and operations, along with multi-cluster automation and configuration. Kubernetes clusters managed by Anthos can then run on-premises, in Google Cloud, or in another cloud provider:



This reference architecture provides opinionated guidance to deploy Anthos in a hybrid environment to address some common challenges that might face.

In this reference architecture, the term *cluster* means a Kubernetes cluster managed by Anthos unless stated otherwise. For example, some sections discuss *VMware vSphere clusters* composed of ESXi servers that pool compute resources.

This reference architecture is separated into two parts. Make sure you read both parts carefully to learn how to plan, design, and implement your own Anthos hybrid environment:

- Part 1 (this document) - Architecture, Anthos components, reference deployments, design prerequisites, and design considerations.
- [Part 2 - Implementation details.](#)

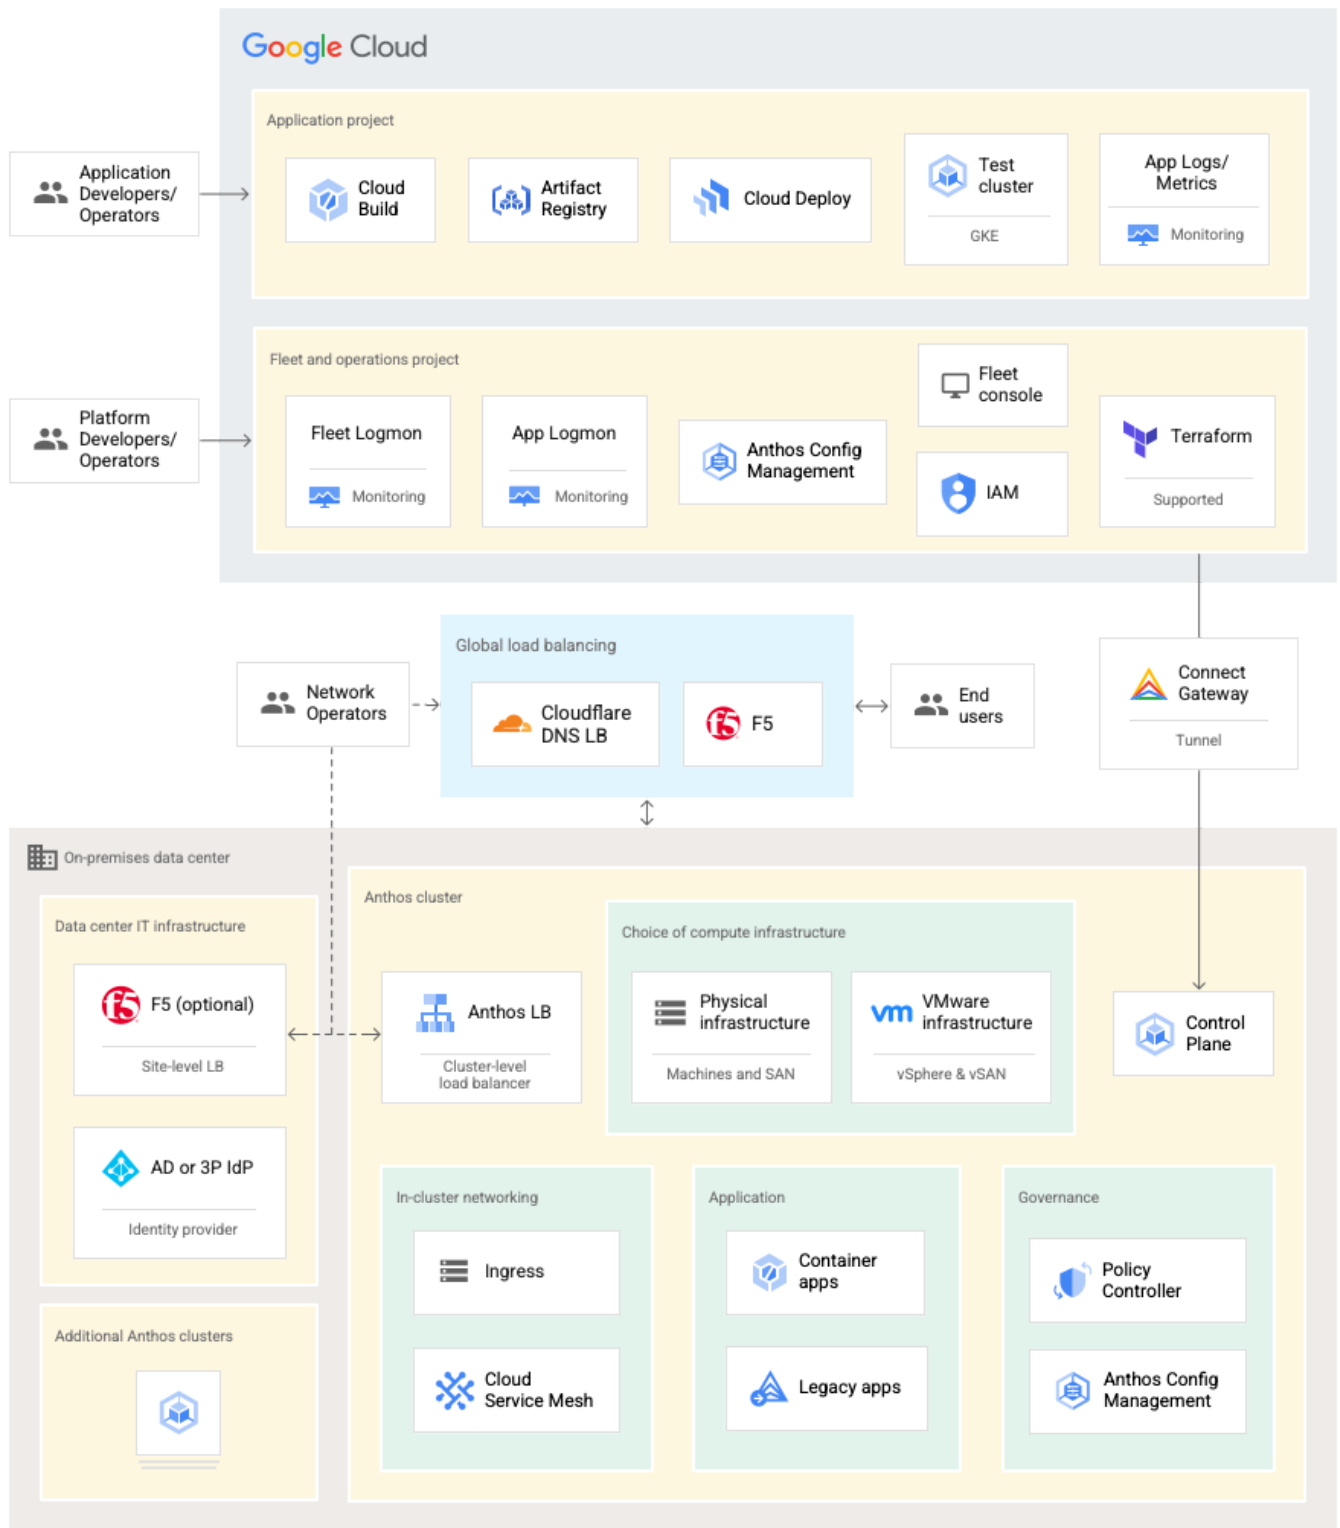
## Architecture

The following architecture diagram provides an overview of a complete Anthos deployment in a hybrid environment. This reference architecture shows you how to appropriately plan, deploy, and configure these components:

- Google Cloud-based services help you manage logging and monitoring data, store container images, and provide configuration management.
- Global load balancing solutions manage traffic flows between the cloud-based services and your on-premises environment.
- The Connect gateway provides secure communication for Anthos management between environments.
- On-premises components that run in your own data center like physical servers and clusters, identity solutions, and load balancers complete the hybrid approach.

Although you might not use all of these services in your own deployments, this reference architecture explains how all of these components can work together. As your business needs change, you can add additional features and services into your environment to better run and manage your workloads.

The diagram also shows some different user personas who interact with the services, such as application developers, application operators, platform developers and operators, and network operators. Each of the personas has access to the resources that they need.



## Anthos components

Anthos consists of services that run in Google Cloud, and various functional components that run on-premises. This reference architecture uses the following Google Cloud-based services:

- **The Google Cloud console** provides a consolidated view of Anthos clusters that span across sites, and lets platform administrators manage and observe their setup across fleets.
- **The Connect gateway** lets people and automation tools connect to clusters in the cloud using Identity and Access Management (IAM).
- **Hosted cluster lifecycle management services** let you create and manage Anthos clusters on VMware and Anthos clusters on bare metal using the Google Cloud console.
- **Fleets** let you group together multiple Kubernetes clusters from on-premises Anthos clusters, Anthos on Google Cloud, and Anthos Multi-Cloud to enable and configure multi-cluster functionalities. Clusters are automatically added to a fleet when they're created.
- **Anthos Config Management** is a set of Google Cloud-hosted and in-cluster components that let you automatically deploy shared environment configurations and enforce approved security policies across fleets.
- **Anthos Service Mesh** is a suite of tools that helps you monitor and manage a reliable service mesh on-premises or on Google Cloud. Anthos Service Mesh is powered by Istio, a highly configurable and powerful open source service mesh platform.

Anthos provides conformant Kubernetes releases with opinionated provisioning and enhancements to more deploy and manage cluster lifecycle. In a hybrid deployment, Anthos lets you deploy Kubernetes clusters with GKE version consistency on bare metal or VMware vSphere environments with support from Google.

**Anthos clusters on VMware (GKE on-prem)** and **Anthos clusters on bare metal** provide upstream Kubernetes releases with proprietary enhancements and support. All Kubernetes cluster components, like the control plane and worker nodes, are hosted in your on-premises data center. Clusters can be managed locally or through the cloud.

- Anthos clusters on VMware integrate with VMware vCenter to provision Kubernetes clusters and provide lifecycle management of the virtual machines and operating systems.
- Anthos clusters on bare metal use your own physical or virtual machines, and you manage the base operating system.



## Reference deployments

This section of the reference architecture provides high-level examples for some important components in an Anthos hybrid environment. These reference deployments show you how to plan, design, and implement the components that make up a complete Anthos hybrid environment. The [Design prerequisites](#) and [Design considerations](#) sections in the rest of this reference architecture provide additional context to plan and design your Anthos hybrid environment.

### Global view

This reference architecture consists of two or more on-premises computing customer sites and two or more Google Cloud regions, as shown in the following diagram.

There are two sites in this global view diagram, with a total of eight clusters. Multiple clusters are used for several reasons, such as:

- Two sites, ABC01 and XYZ01, are used for disaster recovery.
- Two environments, production and staging, to test infrastructure changes.
- Two types of clusters are used in each environment: admin clusters and user clusters. This approach separates administrative resources, which is a security best-practice.

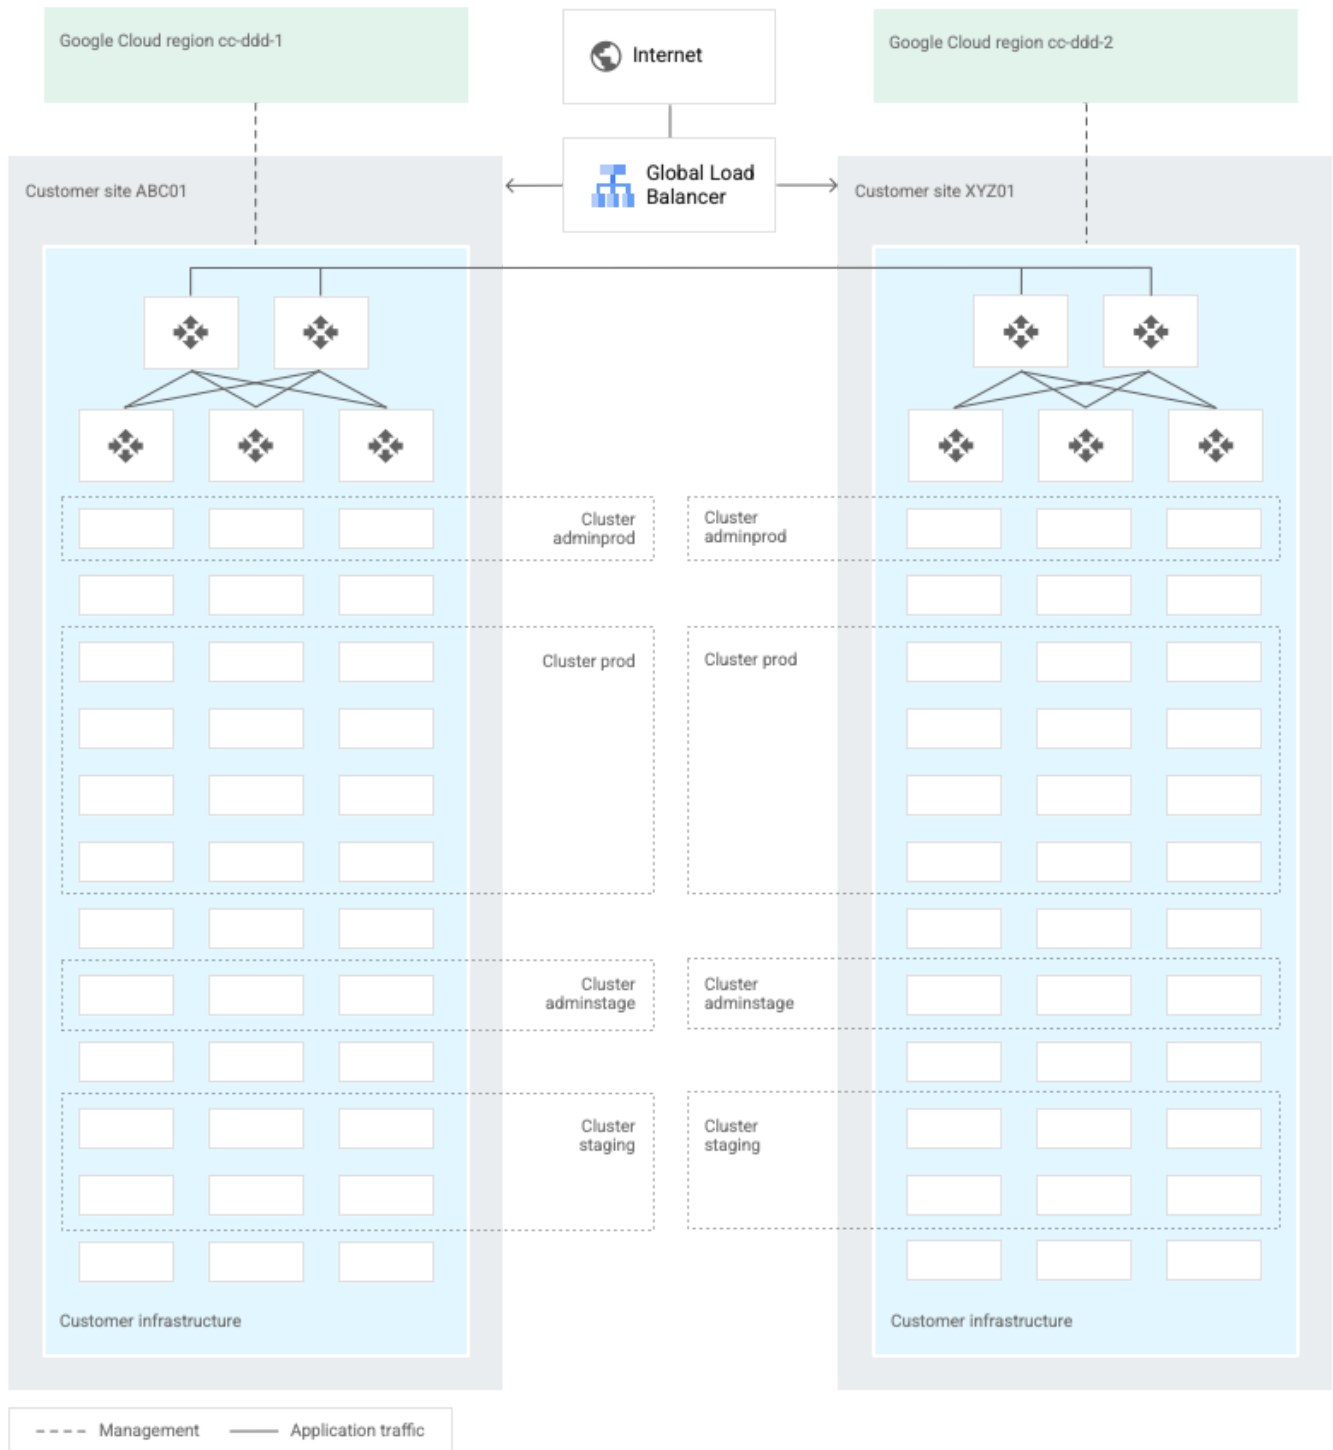
Anthos clusters have two types of clusters - *admin* clusters and *user* clusters. Admin clusters manage the lifecycle of user clusters - they are administrative. User clusters are where you deploy your applications. Separate admin clusters for production and staging environments isolates changes and allows progressive updates.

Using two sites provides for disaster recovery in the following ways:

- Allows critical applications to run in active-active configuration across both sites, with DNS load balancing to distribute traffic.
- Databases and backups can be replicated across sites.

Each customer site in the following diagram is also paired with a different Google Cloud region. Multiple Google Cloud regions serve several purposes, such as:

- If there's a configuration error at the cloud level, the scope of problems is limited. Changes made in one region at a time shouldn't affect applications that are architected for multi-region / multi-site. This behavior is an extra line of defense against misconfiguration, along with having multiple environments.
- In the unlikely event of a Google Cloud service outage, the outage is usually isolated to a single cloud region. A following section on availability provides details about the effect of regional outages.
- Network paths to cloud services can be optimized on a per-site basis.



This reference architecture uses separate production and staging environments. Separate staging environments let you safely test the following types of changes:

- New versions of software that you want to deploy to production.
- Validate new Anthos software versions before you upgrade production clusters.
- Validate changes to cloud-based settings which affect multiple clusters at once.



Staging environments allow infrastructure changes with potentially broad effects to be tested before deployment to production. A staging environment contains realistic versions of your applications, and might also be used to test application or infrastructure changes. Not every site needs a staging environment, particularly if you have more than two sites. However, having at least two staging clusters allows changes to multi-cluster management features of Anthos to be better tested than with only one.

Additional clusters to support other environments<sup>1</sup> can be created on-premises. For example, an on-premises sandbox cluster can be used to test new Anthos features and keep the staging environment as close as possible to production. You can also create test and development environments with cloud-based clusters using Google Kubernetes Engine (GKE).

## Regional view

In the previous global view, inbound traffic passes through a global load balancer which selects a site to receive the traffic. The site traffic then either passes through a local traffic manager inside a perimeter network and is routed to an appropriate cluster to handle the request, or arrives at VIPs to be balanced using Anthos bundled load balancing.

To tolerate limited failures within a site, and to aggregate sufficient compute resources, clusters can be spread across racks. However, clusters don't stretch across sites. Each site is the boundary of a cluster.

The following diagrams illustrate a single site with the four clusters spread across three racks, in the case where virtualization isn't used.

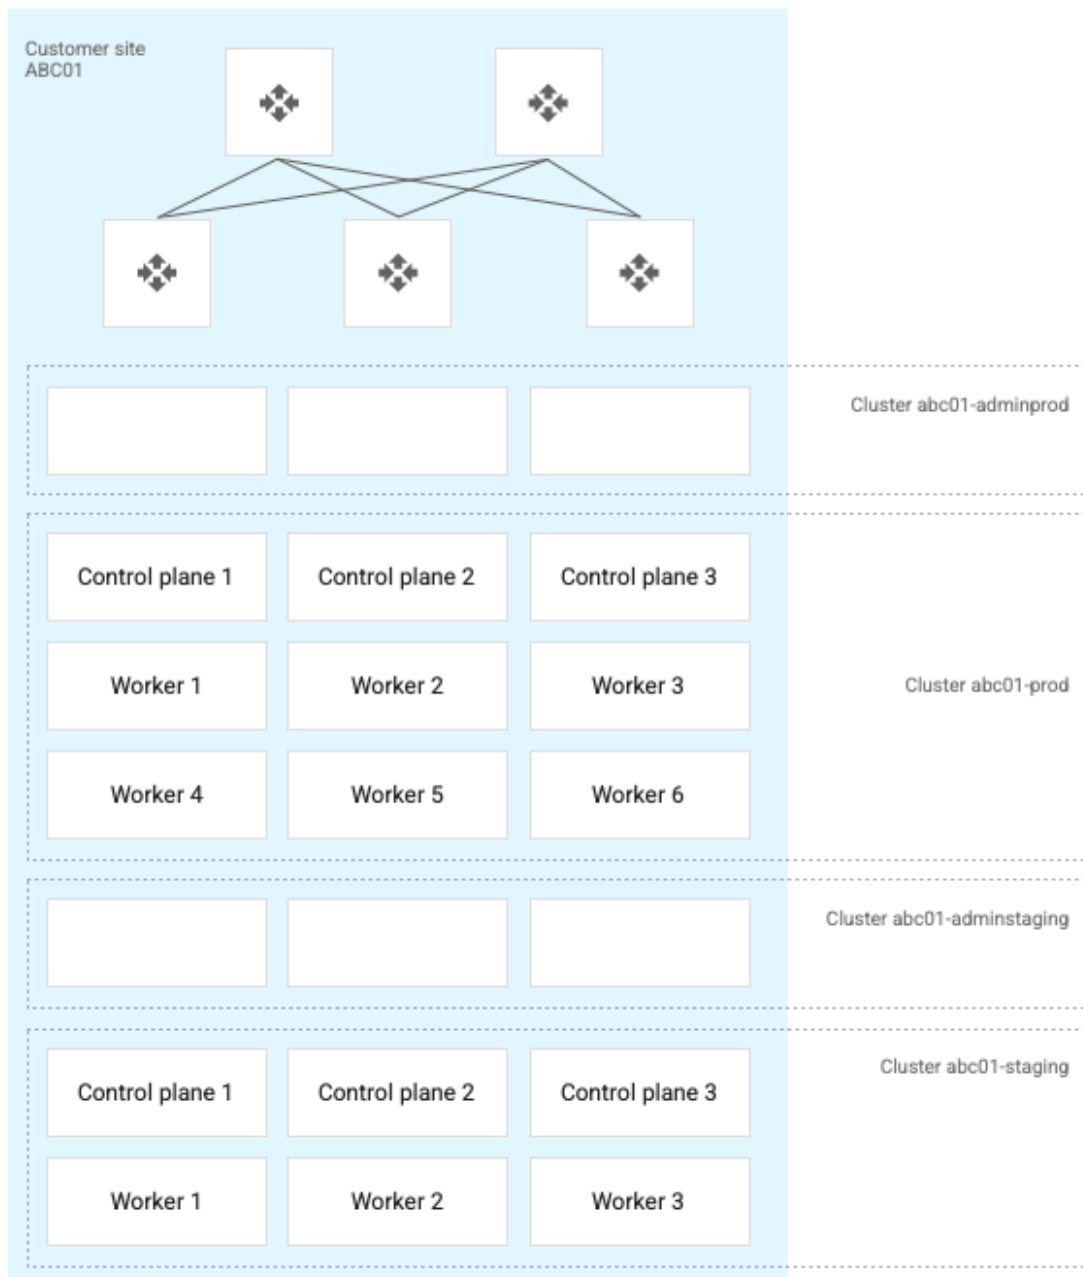
In both of the following example deployments, the clusters are configured in the following way:

- Each cluster has a highly available (HA) control plane with three members. This configuration allows for continued control plane availability concurrently with operating system upgrades, control plane software updates, or single-machine hardware or kernel failures.
- User clusters have separate control planes and worker nodes.
  - The separate control plane nodes reduce the likelihood of an application security breakout being escalated to the control plane.
  - Your applications run on the user cluster worker nodes.
- The admin clusters run local management functions for the user clusters.
  - There are two admin clusters so that admin cluster configuration changes and updates can be tested in the staging environment first.

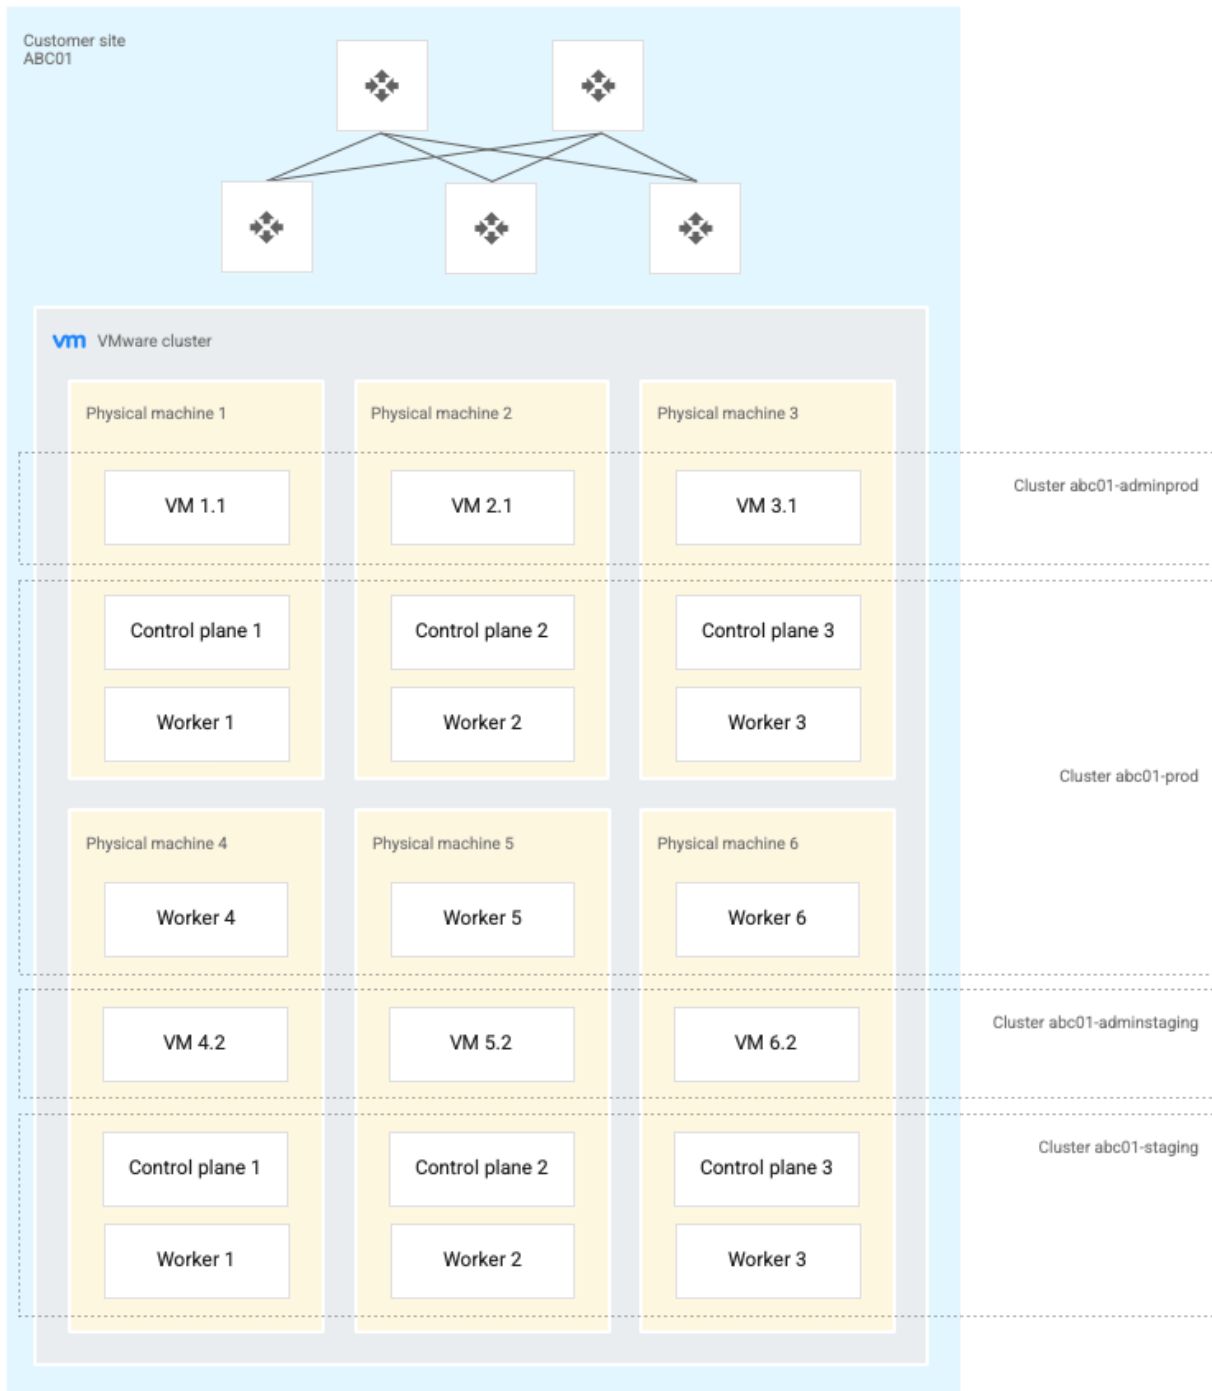
For simplicity of presentation, only one production cluster is shown per site. This design is a good starting point to bring a few large applications to production. As you bring more applications and lines of business to Anthos, it's useful to have several production clusters on a given site. Tradeoffs in the number of clusters per site are described in more detail in the [section on Multi-tenancy](#).

Each cluster typically resides on a separate subnet and VLAN from other clusters. We suggest that you put all of a cluster's nodes on a single subnet and VLAN, although other configurations are possible.

The following diagram shows an example of an Anthos clusters on bare metal deployment:



When VMware or other virtualization is used, the ESXi hosts or equivalent can be spread across several racks. This approach enables the VMs to be spread across hosts, as shown in the following diagram:



For Anthos clusters on VMware, vSphere controls the VM placement. This management control might result in a less organized arrangement than illustrated. If you want a stricter deployment pattern, you can use VM host group affinity to force vSphere VM placement to follow the rack layout. With this approach, an ESXi host group can be created that reflects a rack. These groups can also be thought of as “partitions” within a vSphere cluster. When you create node pools, you can assign the host groups to a specific node pool. This behavior forces vCenter to follow the suggestions and place all VMs from a node pool onto the referenced ESXi host group.



## Clusters and environments

Anthos works together with Google Cloud's resource hierarchy<sup>2</sup> to help apply policies, organize resources, and control access. Organization level policies and tags can then be applied to resources and users of the Anthos-related projects. This approach reinforces consistent behavior with how cloud-only resources are controlled.

On-premises clusters are represented by cloud resources. Kubernetes-specific resources, such as Pods or Deployments, are typically not directly governed by Google Cloud resource hierarchy, but by Anthos Config Management and Policy Controller.

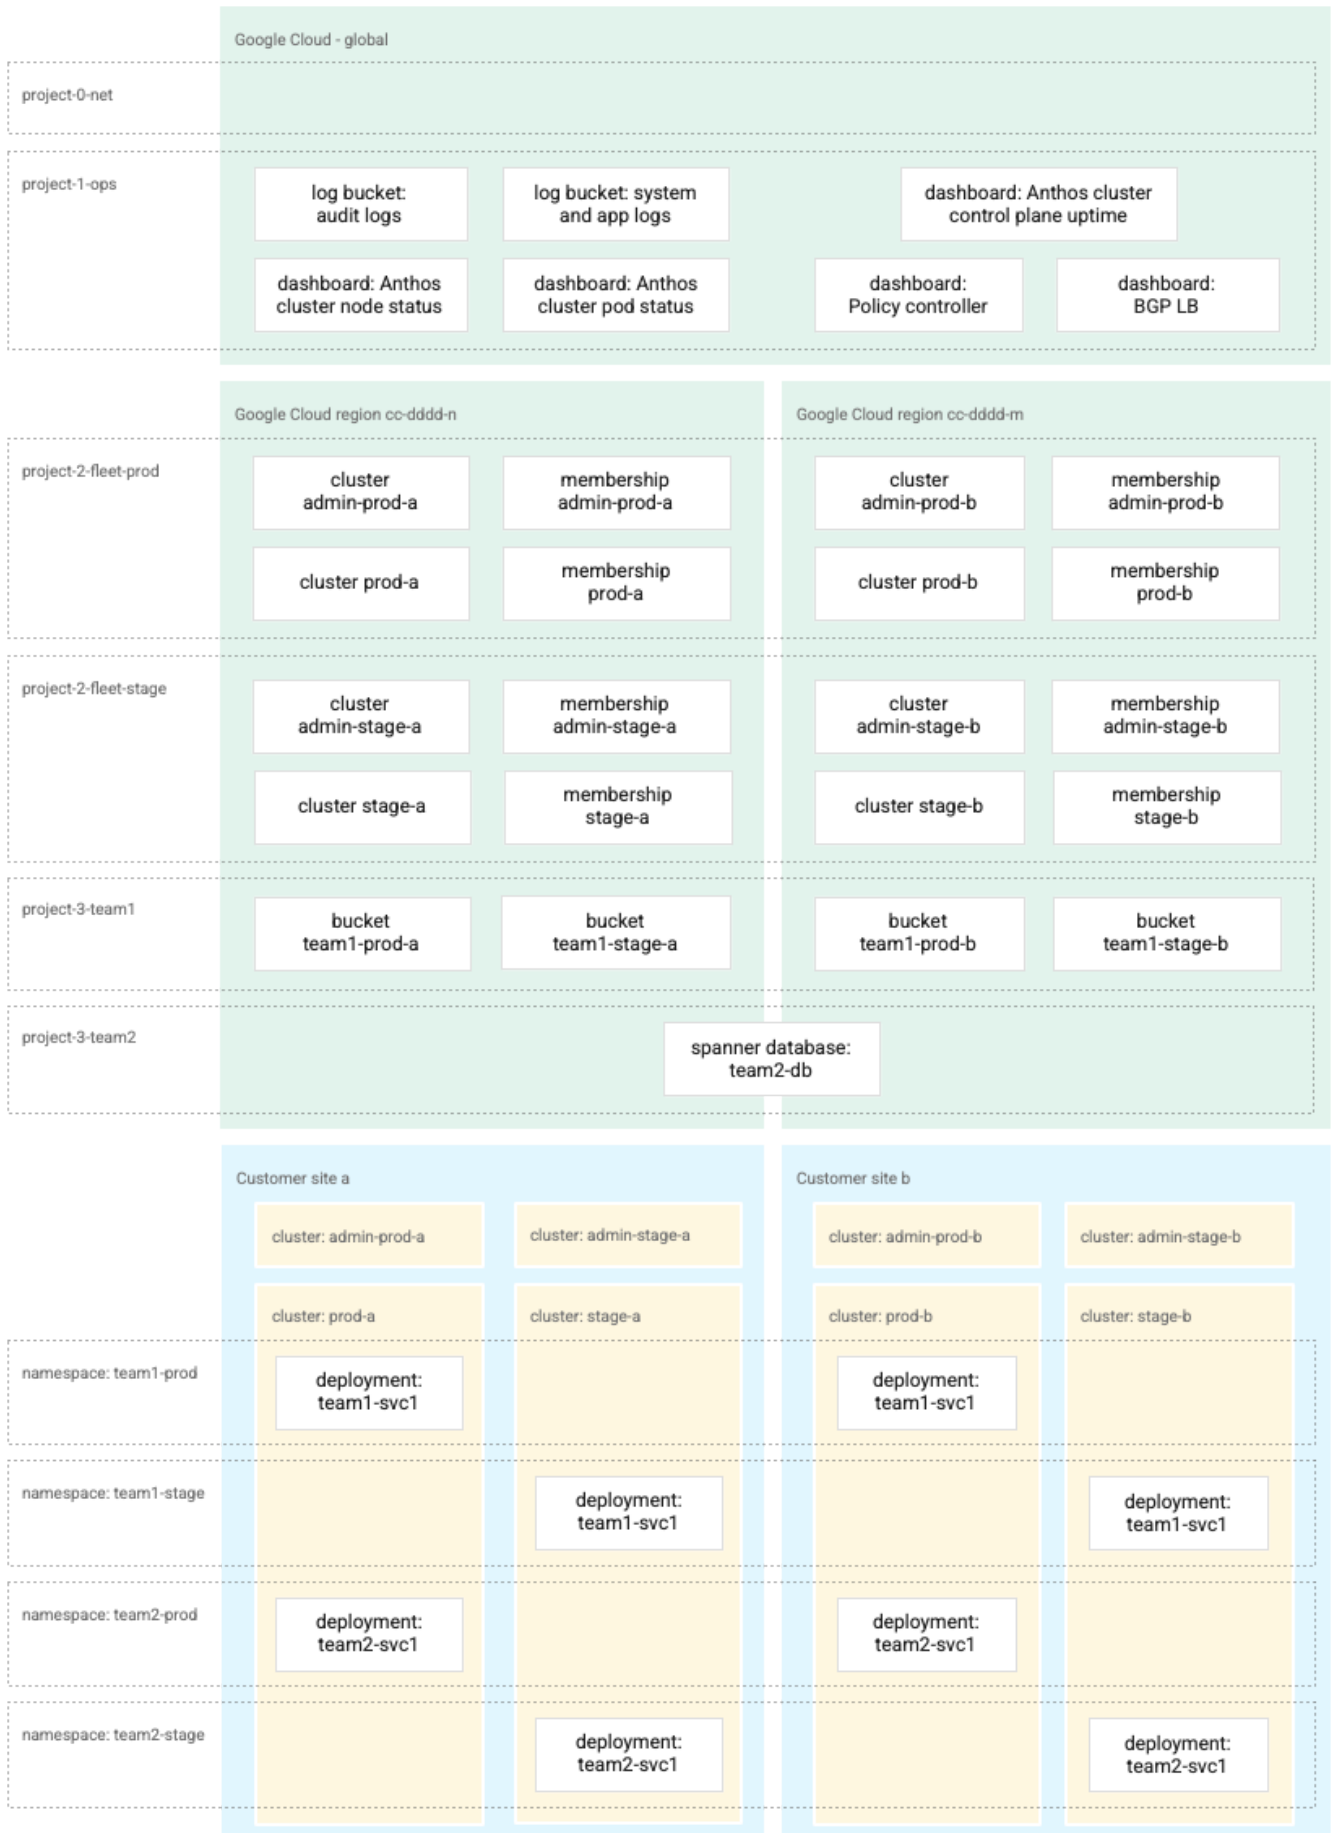
The following diagram gives a global view of logical cloud and on-premises resources that exist in this reference architecture. Resources are spread across several projects with appropriate identity and access management (IAM) controls applied. Also shown are examples of applications deployed across sites, namespaces, and clusters. By separating resources, you can provide more granular access to management capabilities. Resources spread across regions and locations help provide redundancy and provide resources close to where your users are.

In this diagram, resources in Google Cloud are organized into projects and locations as follows:

- `project-1-ops` holds operational data like audit logs, application logs, and dashboards. This project is a global project in Google Cloud that spans all regions.
- `project-2-fleet-prod` holds production clusters and identity information. This project spans two regions for redundancy or geo proximity to where you want applications to run.
- `project-2-fleet-stage` holds staging clusters and identity information. This project spans the same two regions as the production clusters and data.
- `project-3-team1` and `project-3-team2` are where individual teams can store data and run applications. Some teams might split between production and staging buckets and across regions, or some teams might use services with built-in replication across regions like Cloud Spanner.

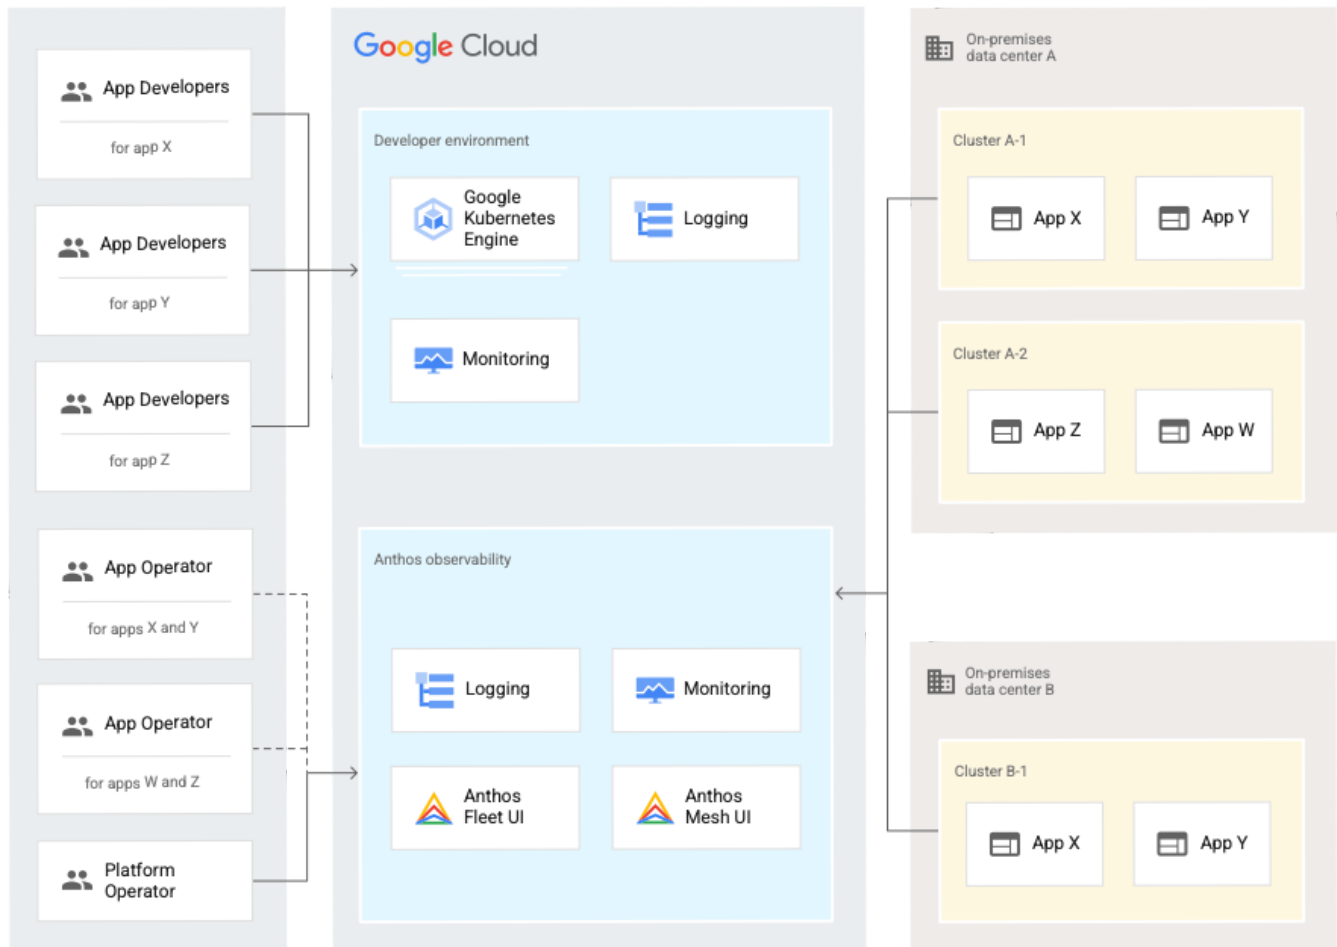
The diagram also shows resources deployed across namespaces in two customer sites as follows:

- Each site has admin production and staging clusters, like `abc01-adminprod` and `xyz01-adminstage`.
- Each site also has two user production and staging clusters, like `abc01-prod` and `xyz01-staging`. These clusters are where your application workloads run.
- Multiple namespaces like `team1-prod` and `team2-stage` exist across user clusters and sites for individual teams to deploy and run applications in both staging and production environments.



## Observability

The following diagram shows how application developers and application or platform operators can view logging and monitoring data in Google Cloud. The on-premises clusters and applications send this logging and monitoring data back to Google Cloud for analysis and review:



In this approach, different personas are granted access to only the environments they need. Applications that run in on-premises clusters direct logging and monitoring data back into Google Cloud for analysis and review.

A following section in this reference architecture on implementation details provides more context and considerations to deploy this observability piece.


## Configuration management

Anthos Config Management can be used to manage Kubernetes objects in all the clusters. Anthos Config Management is a GitOps-style tool<sup>3</sup> that uses a Git repository or Open Container Initiative (OCI)



image as its storage mechanism and source of truth. Git provider workflows allow multiple stakeholders to participate in review of changes.

As shown in the following diagram, a common Anthos Config Management deployment uses one folder containing configuration for all clusters. Separate additional folders each hold configuration data, one for application:



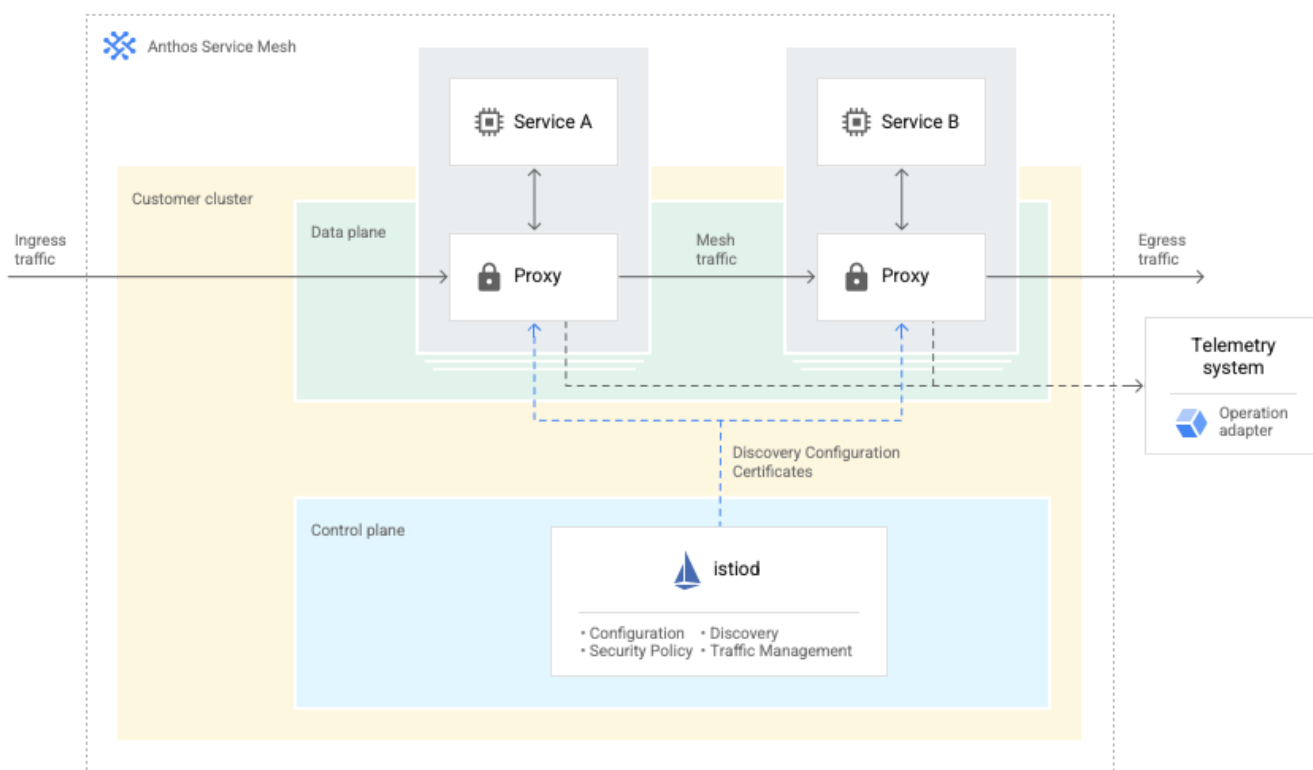
A following section in this reference architecture on implementation details provides more context and considerations to deploy this configuration management piece.

## Security

The following services can help secure traffic as part of a deployment of this reference architecture. These services help with authentication, connectivity, and communication in a cluster:

- Anthos Identity Service connects clusters to on-site identity providers to authenticate local access.
- Connect gateway and workforce identity federation can provide secure cloud-mediated access to mobile workforce clusters without using a VPN.
- Workload Identity provides on-premises workloads with managed short-lifetime credentials for access to cloud resources.
- Anthos Service Mesh encrypts and controls communication between services in the same cluster.

The following diagram shows how Anthos Service Mesh can control the flow of traffic between services within a cluster:



A following section in this reference architecture on implementation details provides more context and considerations to deploy this service mesh piece.

## Design prerequisites

To deploy Anthos in a hybrid environment that follows the guidance in this reference architecture, there are several design prerequisites that must be met. Prerequisites include deciding on Google Cloud regions and local sites to use, planning for the required on-premises hardware, and understanding the network and interconnectivity requirements.

This section of the reference architecture discusses the design prerequisites to deploy Anthos in a hybrid environment. Identify the key areas in this section that align with your own needs and goals, then verify that your environment meets the defined requirements.

### Regions and sites

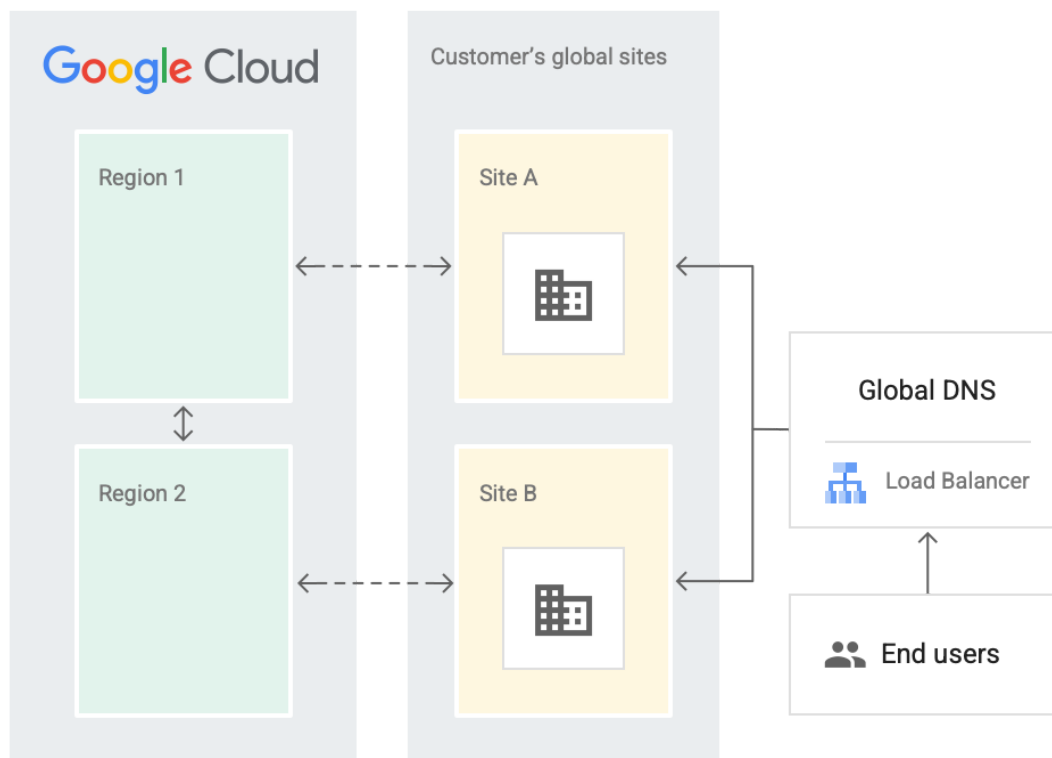
**Identify several on-premises data centers, or sites, to run Anthos clusters.** Sites may be data centers that you own, colocation facilities, or major offices or facilities with secure machine rooms. Each site typically has a dozen to hundreds or more physical machines.

To build and manage highly available applications, choose at least two separate geographic sites. You can choose more than two sites to support your required workloads and availability demands. Anthos hybrid environments typically have fewer than 50 sites.

Each site should have a short name. This reference architecture recommends a naming convention for sites of a city code plus two digits, like nyc01.

**Choose at least two different Google Cloud regions which are proximate to the sites of your Anthos clusters.** These sites run the associated Google Cloud services for the Anthos hybrid environment in this reference architecture.

The following diagram shows how you can connect using Cloud DNS and Cloud Load Balancing to Anthos clusters that run in two different on-premises sites. Two proximate Google Cloud regions provide additional Anthos services:



## Google Cloud setup

**Install the latest version of the Google Cloud CLI<sup>4</sup> and related tools**, such as the latest available version of `kubectl` and `anthos-auth`. If you're behind a corporate proxy or firewall, you can reconfigure `gcloud` CLI<sup>5</sup> so it can successfully access the internet. You manage resources using different tools:

- Google Cloud resources are managed using the Google Cloud console, `gcloud` CLI, or Terraform.
- In-cluster resources are managed with Anthos Config Management. The use and configuration of Anthos Config Management is discussed later in this reference architecture.

**Plan the Google Cloud organization and structure.** Before you deploy Anthos, set up a Google Cloud organization and any folders that you want to use. Anthos projects should be placed within this folder hierarchy.

**Plan the projects, naming conventions, and permissions to hold your Google Cloud resources.** This reference architecture deploys resources across several projects. This approach lets different teams access and own the resources appropriate to them:

1. A host network project, such as `project-0-net`, that contains Google Cloud networking resources. A network operations team typically owns this project.
2. Two fleet projects, such as `project-2-fleet-prod` and `project-3-fleet-staging`, that contain Google Cloud resources which represent your on-premises clusters. These projects also include logging, monitoring metrics, dashboards, and alerts for your clusters and



applications. Depending on their permissions, users with access to this project can view, create, and update your on-premises clusters, and view logging and monitoring data.

3. Several application projects, one for each application or small set of related applications, such as `project-n-app-xyz`. The development team and the operations team for that application are granted access to this project. A central platform team can own and allocate the projects. This project also contains images and cloud CI/CD pipelines for the applications. As new applications are onboarded, you need a process to create these projects on-demand.

**Plan your cluster names.** Every cluster must have a unique name within the project in which it's registered. A consistent format for cluster names is recommended. This reference architecture uses the site code and an environment code in the cluster name.

## On-premises setup

The following guidance helps you understand the prerequisites to design the on-premises parts of an Anthos hybrid environment.

### Site services

**Run a local container image registry in each site.** This local registry ensures that images can still be pulled during network outages, and reduces image network traffic when applications scale up rapidly. Application configuration is also stored in an image registry. Use an OCI-compliant registry. Many popular image registries support the Open Container Initiative (OCI) Distribution Specification including Artifact Registry and JFrog Artifactory<sup>6</sup>.

**Provide Git access at each site.** Some configuration is stored in Git. Anthos Config Management needs read-only access to your Git repository to read cluster configuration. Google Cloud doesn't provide or support a Git repository. A cloud-based Git provider may be used as the primary storage and for reviews. Use either a GitHub or GitLab for Enterprise account to ensure sufficient fine-grained access controls and API request rates. A Git mirror at each site can be provided if it's necessary to continue operations when there's an internet or Git provider outage.

**Provide an identity provider.** Federate your corporate ID provider to Cloud Identity, including the ability to create new groups in your identity provider as needed. Local and cloud-based identity providers are supported. Management operations may be limited if there's an internet outage or provider outage affecting a cloud-based identity provider.

**Optionally run VMware.** Anthos supports creating clusters based on the following types of infrastructure:

1. Physical machines, managed by you.
2. Virtual machines, on VMware and managed by you.
3. Virtual machines, by using integration with VMware.



When using the third option with this reference architecture, the VMware requirements are as follows:

<b>Version</b>	vSphere 7.0 Update 3 or later
<b>License</b>	vSphere Enterprise Plus

## Compute requirements

Anthos offers two ways to run clusters in an on-premises environment - Anthos clusters on bare metal and Anthos clusters on VMware. Review the following compute requirements for each option. For advice on selecting one of these options, see the Design Considerations section.

The following sizes are suggested for deploying the eight clusters of the reference architecture described in a previous section. Other Anthos documentation might mention other node sizes. The sizes in the following table are for this reference architecture.

The following sizes and counts apply to both Anthos clusters on bare metal and Anthos clusters on VMware.

Recommended node sizes for this reference architecture:

<b>Node config</b>	<b>vCPUs per node</b>	<b>RAM per node (GB)</b>	<b>Storage space per node (GiB)</b>	<b>Storage throughput per node (sequential IOPS)</b>
<b>Admin cluster node</b>	4	16	128	50
<b>Staging control plane node</b>	8	32	256	50
<b>Production control plane node</b>	32	128	256	500
<b>Staging and production worker node</b>	4	16	128	Varies

Anthos clusters on VMware also require 380 GiB of storage for vCenter for VM templates, log and metrics buffering, and etcd object data.



Recommended node quantities for this reference architecture:

Node config	Clusters used in	Number needed	Total
<b>Admin cluster nodes</b>	abc01-adminprod	3	12
	abc01-adminstage	3	
	xyz01-adminprod	3	
	xyz01-adminstage	3	
<b>Staging control plane nodes</b>	abc01-staging	3	6
	xyz01-staging	3	
<b>Production control plane nodes</b>	abc01-production	3	6
	xyz01-production	3	
<b>Staging and production worker nodes</b>	abc01-staging	3 - 10	26 - 1020
	xyz01-staging	3 - 10	
	abc01-production	10 - 500	
	xyz01-production	10 - 500	
<b>All</b>			50 - 1044

The number of worker nodes varies with application needs. Production clusters usually have ten or more nodes. Up to 500 nodes per production cluster are supported in this configuration. Staging cluster size depends on how application workloads are scaled down in the staging environment and how traffic is generated or directed to the staging environment. The requirements might be one tenth of the production requirements, but should be at least three nodes.

## Operating system

Select a primary operating system for your Anthos clusters from one of the following:

- For Anthos clusters on bare metal:
  - Public Ubuntu (20.04 or 22.04)
  - RHEL (8.6, 8.7, 8.8)
- For Anthos clusters on VMware:
  - Container-Optimized OS<sup>7</sup> from Google
  - Ubuntu from Google

Anthos clusters on VMware also support worker nodes with Windows Server 2019 10.0 and later.

## Storage

For Anthos clusters on VMware, we recommend using the storage policy in the cluster configuration. With VMware Storage Policy Based Management (SPBM), Anthos clusters on VMware places the VMs on datastores that belong to the same storage policy. This behavior ensures that VMs are allocated storage resources based on the administrator's specific requirements. Additionally, Anthos clusters on VMware automatically adjusts its metadata at the cluster level when storage vMotion happens, which prevents the clusters from breaking. Daily operations on the storage, such as adding or removing a datastore, or migrating Anthos cluster on VMware from one datastore to another, are simpler and easier under SPBM. Anthos clusters on VMware also support VMware-supported storage such as VMFS or vSAN.

For Anthos clusters on bare metal, use NAS/SAN storage from one of the following supported storage providers:

- Dell EMC
- Hitachi
- NetApp

## Networking

The following network requirements should be considered as part of your own deployments that follow this reference architecture.

### *IP addresses*

- IPv4 network addresses are used for machines and for service and pod IP ranges.
- For Anthos clusters on VMware, make sure you plan your IP address needs<sup>8</sup>.

### *Internet connectivity*

- An internet connection is needed for operational purposes.
  - Firewalls must allow outbound connections to certain Google Cloud services<sup>9</sup>. Google Cloud won't make direct inbound connections.
  - Anthos uses the following two patterns to communicate with Google Cloud:
    - TLS connections initiated to a Google Cloud service, such as logs that are sent to a logging service.
    - Some Anthos management services connect to Anthos on-premises clusters by using a tunnel (the Connect Agent). This tunnel also appears to customer firewalls as an outbound TLS connection.
  - Your employees and automation solutions may connect to clusters over this tunnel.
  - The Anthos install and update processes pull images from a Google Cloud service.
- Use of a perimeter network for ingress traffic might not be necessary.





- Traffic only reaches applications when explicitly configured to do so. This traffic includes containerized workloads (Pods) and VMs workloads (using Anthos for VMs on bare metal clusters).
- Use of an outbound HTTP proxy is common.
  - You can run an HTTP CONNECT-based proxy, like Squid, between your cluster and the rest of the internet to limit connections to known domains.
  - Size the proxy to handle peak outbound traffic. This traffic includes container image pull during applications scale up, if there is no registry behind the proxy.

### *Intercluster connectivity*

- All clusters need outbound connectivity to Google Cloud such as for the connect agent or logging and monitoring.
- For Anthos clusters on bare metal:
  - Machines in a cluster need to be routable to each other (L3 connected).
  - Machines can span VLANs to span multiple power and networking failure domains.
- For Anthos clusters on VMware, VMs must be on the same L2 domain.
- We recommend running in island-mode for IPv4.
  - When using dual-stack networking, flat-mode is used for IPv6.

### *Intracluster connectivity*

- All the machines in a single cluster need to have L2 or L3 connectivity with each other.
  - For Anthos clusters on bare metal, ensure L2 or L3 connectivity between machines, and from the admin cluster to the machines.
  - For Anthos clusters on VMware, the VMware Distributed vSwitch or standard vSwitch must be configured to provide L2 connectivity between dynamically provisioned VMs.
    - Set your precreated port groups to the correct VLAN ID<sup>10</sup>.
    - Make sure that the L2 subnet (VLAN) is correctly configured and accessible by all ESXi hosts within the chosen vSphere cluster. Also check that the core routers are interconnected.
- It's not recommended to create network security boundaries between node pools in a single cluster.
  - You might have regulations that require network isolation between applications, such as in-scope for PCI vs not in scope for PCI. If so, consider using multiple clusters if the applications don't require direct connectivity.
  - Alternatively, consider using Anthos NetworkPolicy and Anthos Service Mesh to control traffic between applications within a cluster.
- All the VMs that are part of your Anthos clusters on VMware infrastructure must use the same Network Time Protocol (NTP) server.

### *Basic network*

You can choose software-defined networking (SDN) solutions so long as the core networking configurations are correct. For example, consider the following basic network functionality:

- The IP address, netmask, and gateway all work correctly.



- ARP or NDP messages are honored.
- DHCP, DNS, and NTP servers work correctly.
- Internet gateway or proxies also work correctly.

Depending on different organization structures, all those basic requirements might require cross-team collaboration.

In Anthos clusters on bare metal, you manage the machines. These machines can be in different VLANs (L2 network) so long as they're L3 connected. There can't be any NAT or proxy between L2 networks.

In Anthos clusters on VMware, the software stack manages the VM lifecycle. All the nodes in the same cluster connect to the same virtual portgroup, so they're L2 connected. As a result, Anthos clusters on VMware can enable direct routing mode for pod-to-pod connectivity, while Anthos clusters on bare metal must use tunnel mode.

The following table summarizes the basic network requirements:

	<b>Anthos clusters on bare metal</b>	<b>Anthos clusters on VMware</b>
<b>Customer-managed network</b>	L2 and L3	L3
<b>L2 / switch</b>	Your choice	VMware Distributed Switch
<b>Node placement per cluster</b>	L3 connected	L2 connected
<b>Pod connectivity mode</b>	Geneve tunnel	Direct routing (tunnel disabled)

### *Load balancer*

Global load balancing is also called a DNS load balancer or Global Traffic Manager. If you have an existing global load balancing solution, continue using it. Anthos doesn't integrate with the global load balancer.

Site-level load balancing requirements:

- If you have an existing F5 Local Traffic Manager or other load balancers such as HAProxy or Citrix AD, you can continue using it. Use the manual load balancing mode during setup of your Anthos deployment.
- If you don't have an existing Local Traffic Manager, use the MetalLB bundled load balancing. MetalLB is available for both Anthos clusters on bare metal and Anthos clusters on VMware.

The following table provides an overview for load balancing modes for ingress traffic to your Anthos clusters. The HA control plane load balancing is automatically selected at installation:



<b>Goal</b>	<b>Cluster type</b>	<b>Mode to use</b>
<b>Use Anthos-provided load balancer</b>	Either	MetalLB
<b>Use an existing load balancer</b>	Either	Manual

#### *On-premises to Google Cloud connectivity*

All on-premises clusters connect to a Google Cloud region through the connect gateway, a special-purpose network tunnel service. The session is initiated from within the on-premises network. Connect gateway uses a gRPC / TLS tunnel, and only supports HTTP access to the Kubernetes cluster API. The tunnel doesn't provide general access to the on-premises network.



## Design considerations

This section helps as you design your own implementation of this reference architecture. These design considerations help you implement the most appropriate Anthos hybrid environment for your needs, and understand areas such as security and compliance, operations, and observability.

### Hybrid deployment

Anthos offers two ways to run clusters in an on-premises environment - Anthos clusters on bare metal and Anthos clusters on VMware. Which one to choose is based on your product requirements. The following table details support of some functionality across the two types of clusters:

Functionality	Anthos clusters on bare metal	Anthos clusters on VMware
Performance in general	Yes	Yes
Hypervisor overhead	No	Yes
Customize a supported node OS	Yes	No
Automated node lifecycle	No	Yes
Automated OS lifecycle	No	Yes
Cluster autoscaling	No	Yes
Cluster auto repair	No	Yes
High availability	Yes	Yes
Legacy VM workloads	Yes (A4VM)	No
Automated volume allocation for stateful container workloads	No <sup>11</sup>	Yes
Large scale deployment	Yes, with removing additional hypervisor cost	Yes, with on-demand scale up and down
Highly performance-sensitive workloads	Yes	No
Storage options	External storage	External storage or VMware vSAN or VMFS



The choice between Anthos clusters on VMware on Anthos clusters on bare metal is up to you. Existing deployments, investments, and requirements often help you determine which one is the most appropriate choice.

When you might use Anthos clusters on VMware:

- **vSphere investments:** If you have considerable vSphere investments and want to continue with your investment on vSphere. This guidance includes when you have deep operational knowledge of vSphere and plan to continue that investment.
- **vSphere operations:** You might have operationalized vSphere to provide virtualized compute, networking, and storage technologies to your customers, either as self-service or IT standard. If you want to leverage and continue to use this operational model, Anthos clusters on VMware would be ideal.
- **Node management:** You might want to automate all aspects of node management and to provide those capabilities as self-service to their customers. Anthos clusters on VMware are an ideal choice as it provides an automated lifecycle for nodes and OS, autoscaling and repair, and high availability.
- **Scalability and density:** If you have densely packed Kubernetes nodes and autoscaling of clusters, Anthos clusters on VMware are a better option. This approach provides better bin packing, autoscaling of nodes, and hitless planned downtime.
- **Workload characteristics:** Anthos clusters on VMware are a better fit for workloads that are highly stateful, need higher availability against planned and unplanned downtimes, and can run on hardware platform-independent virtualization layers. They're also ideal for workloads that use software defined storage (SDS) as their block store.

When you might use Anthos clusters on bare metal:

- **vSphere license costs:** If you want to optimize on cost, use Anthos clusters on bare metal as you don't have to incur the cost of vSphere licenses.
- **Operating system support:** If you want to bring your own OS due to company or regulatory requirements. This approach provides more options for OS supportability.
- **Workload characteristics:** If you run high performance and latency sensitive applications that can't tolerate hypervisor overhead. This solution is also ideal for workloads which need hardware dependent features like **SR-IOV** or CPU pinning.
- **VM workload support:** If you want to run workloads running in virtual machines together with containers. This approach supports both containers and VMs managed through a single Kubernetes stack.

## Availability

We recommend that you have multiple environments such as development, staging, and production. These environments let you adequately develop, test, and deploy applications in a more controlled way than a single deployment to production. Production clusters should use highly available (HA) clusters.



We recommended that staging clusters are HA to let you test application availability under simulated infrastructure failures. Development environments and similar can be non-HA.

Anthos workloads automatically move within a cluster in response to hardware faults. For this behavior to be effective, consider the following:

1. If one machine fails, is there enough spare capacity in the cluster to allow all workloads to fit on the remaining machines?
2. If there isn't enough spare capacity, how long will it take to get a replacement machine?
3. For Anthos clusters on VMware, vSphere handles physical machine failure. Anthos clusters on VMware use VM anti-affinity rules to ensure VMs, like the control plane or three-node node pool, are spread across three different ESXi hosts. To give you control of VM distribution across hosts, you can also define VM host group affinity rules in Anthos clusters on VMware version 1.16 and later. Make sure that your vSphere instance can tolerate the loss of one node.
4. For Anthos clusters on bare metal, you need a monitoring solution to detect and react to a failed physical machine.

To mitigate single-points of failure at a rack level, machines from several racks can be used to form a cluster:

- For Anthos clusters on bare metal, a cluster can use machines spanning multiple racks.
- For Anthos clusters on VMware, a cluster can use ESXi hosts placed across racks.
  - VMware will typically spread VMs across hosts.
- Neither cluster type should span sites.
  - For Anthos clusters on VMware, don't use VMware Stretched Clusters.

Anthos clusters on VMware version 1.16 and later can help to increase resiliency and support smarter availability designs by using VM host group affinity rules. These rules let vSphere be more granular on where to place specific VMs. This functionality allows the creation of so-called ESXi host groups within an existing vSphere cluster. ESXi host groups can match the racks where the selected hosts are placed, or could match hosts within a specific datacenter room or closeby location (same campus DC). This ability gives Anthos awareness of the physical infrastructure. By mapping the ESXi host groups to nodepools, it's possible to create availability zones within the same vSphere cluster. Anthos and vSphere make sure that VMs are placed in the corresponding host group, which means applications can be deployed across node pools to ensure a higher resiliency and availability.

Anthos is designed to ensure that applications continue to operate despite a temporary disconnection from Google Cloud<sup>12</sup>, such as a cloud service or network connectivity outage. Applications continue to run on the same nodes, and cluster networking and load balancing operate normally. Although unlikely, it's possible a physical machine might fail during a temporary disconnection event.

Consider the behavior of replacement Pods that start on other machines when a physical machine fails, concurrently with a temporary disconnection from Google Cloud:



1. Application container images might be stored in another site or in a cloud service such as Artifact Registry. Machines would be unable to pull the image to start the replacement process. For this reason, run a local image repository in each site that mirrors your primary image registry.
2. An HA control plane can tolerate the loss of one node when the physical machine is part of the control plane. The loss of additional nodes on several independent physical machines is unlikely to occur within the timeframe to replace the first failure. This potential problem can be further mitigated by running a local registry with a mirror of the Anthos system container images<sup>13</sup>.
  - a. For Anthos clusters on VMware, use *thick images* that contain system container images. Anthos clusters on VMware load these images into your local registry.
  - b. To replace a failed machine in Anthos clusters on bare metal, Anthos needs to install certain additional OS packages. Consider running a local OS package repository<sup>14</sup>. This approach can also be useful to remove a machine for troubleshooting.
3. If it's necessary to add worker nodes to the cluster, they may not be able to start without pulling certain container images. Running a local registry mirror ensures they can start.

Consider the availability of authentication services. Anthos Identity Service bridges to third-party authentication providers for login over local networks. Anthos Identity Service doesn't have a cloud dependency. However, Anthos Identity Service can be configured to use a cloud-based identity provider, such as Azure AD. If there's a network outage between your site and the internet, you might not be able to authenticate to your clusters.

## Network security and management

Consider when to use customary network security practices like VLANs and firewalls, and when to use Anthos networking features.

Customary network security practices are recommended for separating clusters from other clusters or from other parts of the network, and include the following suggestions:

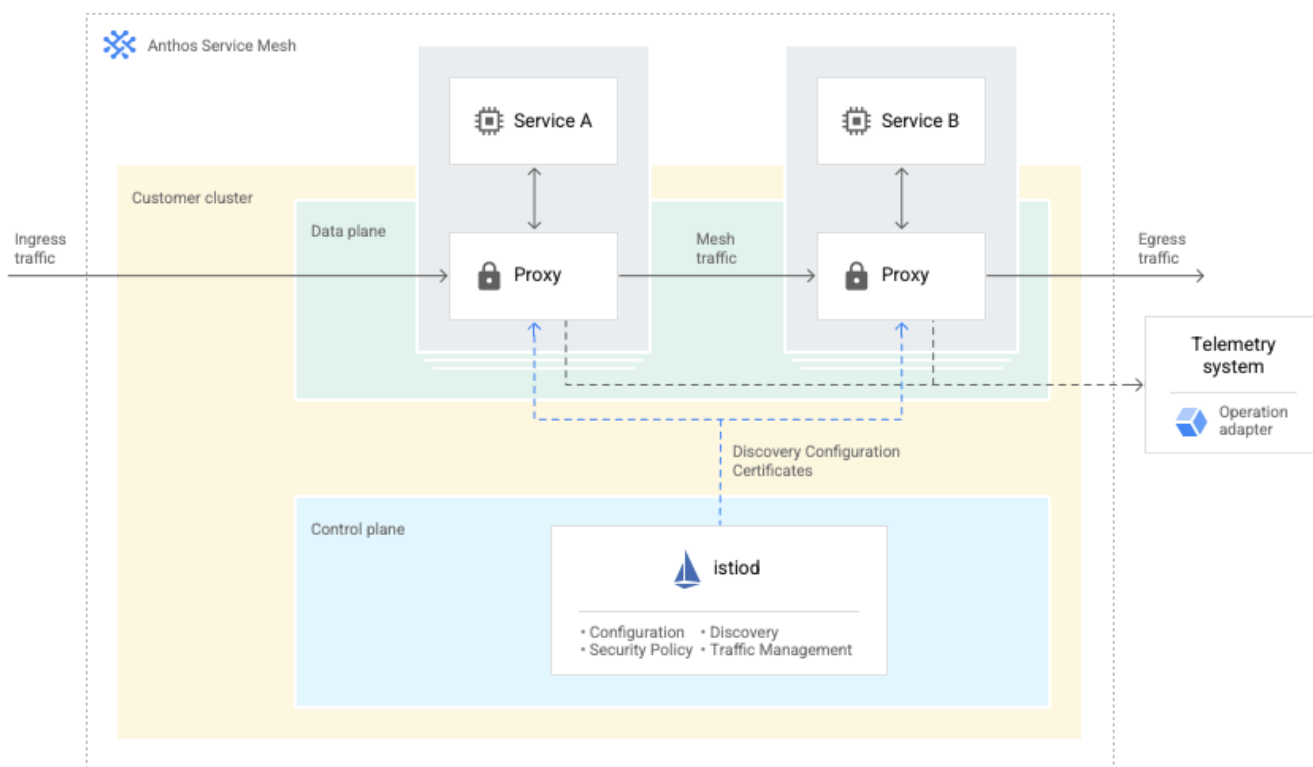
- Place the admin clusters on an administrative subnet and VLAN.
  - Although this separation isn't required by Anthos, separating management traffic from non-management traffic is a commonly accepted practice.
  - Place all nodes of an admin cluster on the same VLAN.
- Place user clusters from different environments or with different security requirements on different subnets and VLANs from each other.
  - Although this separation isn't required by Anthos, separating production and non-production environments is a common practice.

To help mitigate the risk of data exfiltration with Anthos clusters on bare metal, you can use VPC Service Controls. VPC Service Controls work with either Cloud Interconnect or Cloud VPN to provide additional security for your clusters. Using VPC Service Controls, you can add projects to service perimeters that protect resources and services from requests that originate outside the perimeter.

In a user cluster, you can use Anthos networking features to replace the removed boundaries when previously separated applications are consolidated onto a single cluster. Some factors to consider including the following:

- Containers in a cluster can communicate to each other by default. Communication within a cluster can be restricted using Kubernetes Network Policies, or with Anthos Service Mesh.
- Several user clusters can be placed on the same VLAN.
  - In this case, isolation can be achieved with separate namespaces and the use of Kubernetes Network Policies.
  - This approach can make reallocation of machines between uses easier for Anthos clusters on bare metal.

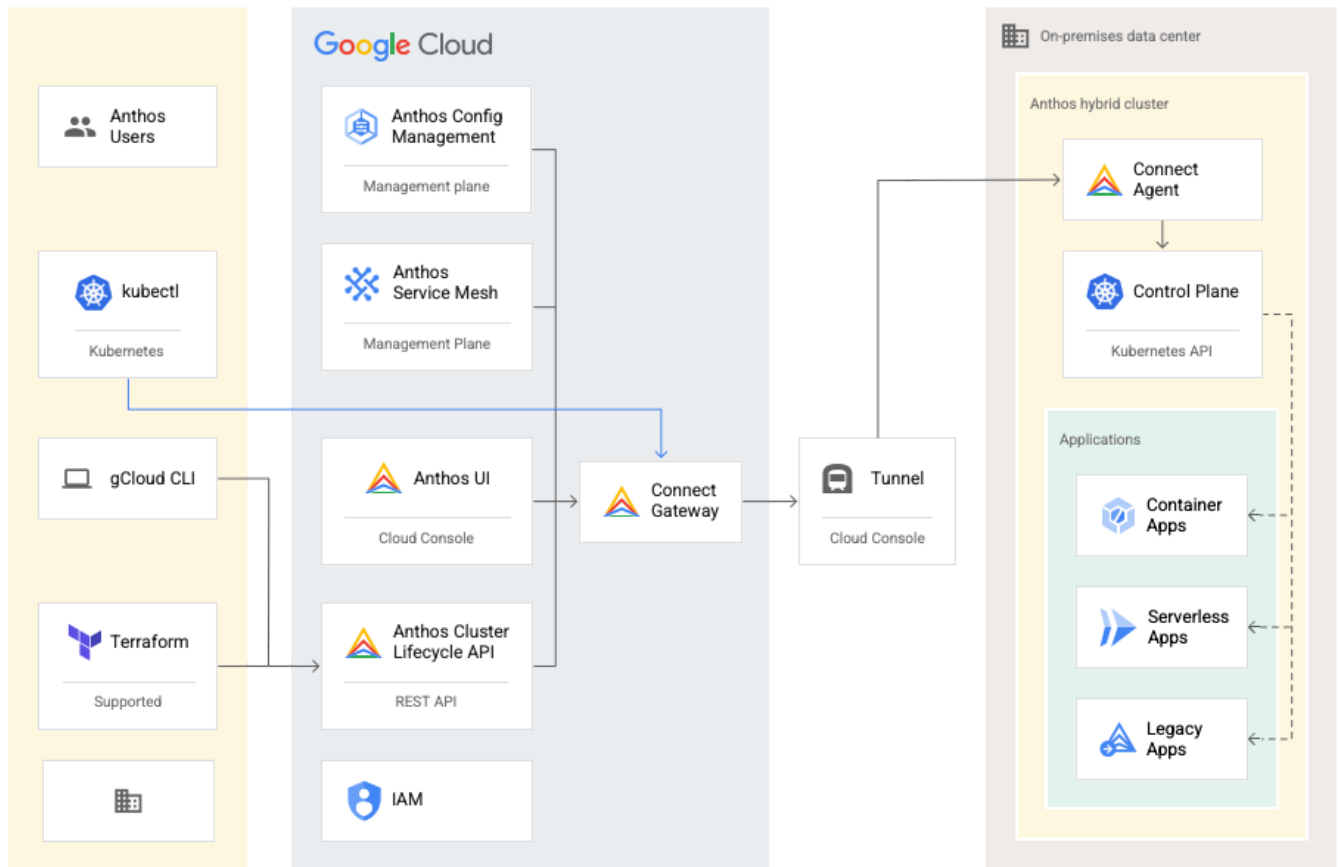
The following diagram shows how Anthos Service Mesh can control the flow of traffic between services in the same cluster:



When using Anthos Service Mesh to secure service traffic within a cluster, you should also secure external traffic. Don't allow direct connections to clusters for management operations by users or automated processes. Instead, connect to clusters by using the private on-premises network or through the connect gateway.

The following diagram shows how all operations can use the connect gateway to connect to clusters using a secure tunnel:





## Scale and limits

The following scale and limits apply to deployments of this reference architecture. Take these values into consideration for your own design.

### Cluster and node limits

The following scalability limits apply when all the recommendations in this reference architecture are followed. These limits apply to Anthos clusters on bare metal and Anthos clusters on VMware. Other documentation may mention lower limits for more general use cases, or higher limits for more specific use cases.

<b>Clusters</b>	Total per fleet	250
	Admin clusters per site	1-2 <i>recommended</i>
	User clusters per admin cluster	100
<b>Pods</b>	Per node	110
	Per cluster	15,000
	Changed per second	20 <i>created / modified / deleted</i>



<b>Nodes</b>	Per cluster	500
	Per node pool	200
<b>Services</b>	Per cluster	10,000
	Per cluster, type nodePort OR LoadBalancer	250 <i>total of both types, average of ≤ 60 endpoints per service</i>
	In a single namespace	5,000

### Load balancing limits

Anthos clusters configured in accordance with this reference architecture and that use MetalLB load balancing support clusters with the following maximums:

<b>Maximum flows per Service</b>	Up to 30,000
<b>Maximum VIPs per cluster</b>	1,000
<b>Failover time (min-max)</b>	2 seconds - 15 seconds

Manual load balancer performance varies, but typically meets or exceeds the preceding limits.

### Anthos Config Management limits

Use multiple repositories, at least one root repository and one repository for each namespace. This approach isolates faults to a smaller surface, provides better use of repository access controls, supports more objects, and reduces sync latency. The following additional scalability limitations apply:

<b>Total resources in any one root or namespace repository</b>  <i>Mean resource size = 3 kB</i>	4,000
<b>Number of clusters managed by Anthos Config Management</b>	250
<b>Number of namespaces repositories</b>	2,500



## Observability limits

The following observability limits apply to deployments of this reference architecture:

<b>Logging throughput per node</b>	100 KB per second per node
<b>Number of teams with restricted logging access</b>  <i>Limit applies to the number of logViews in the operations project, which provides access to a group for a subset of logs.</i>  <i>No limit on number of teams with no or full logs access.</i>	25

## Multi-tenancy

When you deploy a new application, consider whether to place the application in a new namespace in an existing cluster, or to create a new cluster. The former mode is called *multi-tenant*, and the latter are *single-tenant clusters*. *Tenant* means an application developer team. This term is not the same as a customer of a SaaS application provider.

Advantages of multi-tenant clusters include the following:

- Reduced resource fragmentation.
- No need to wait for cluster creation for new tenants.
- Easier to configure communication between services in the same cluster.
- Cost-effective to provide high-availability for small workloads.
- Lower administrative effort with fewer clusters.

Advantages of single-tenant clusters include the following:

- Chargeback can be based more directly on hardware costs.
- Delegates management of cluster capacity and application priorities to an application team, such as a single-tenant cluster that runs many batch jobs throughout the day.
- Allows application teams to manage their own API extensions, or *operators*, without affecting other tenants.
- Allows application teams to manage namespaces when dynamically creating applications, such as in a SaaS application.
- Separates noisy-neighbor applications from other applications.
- Provides an extra security boundary for applications that run risky code.
- Placing an entire cluster within a network boundary to meet compliance requirements.

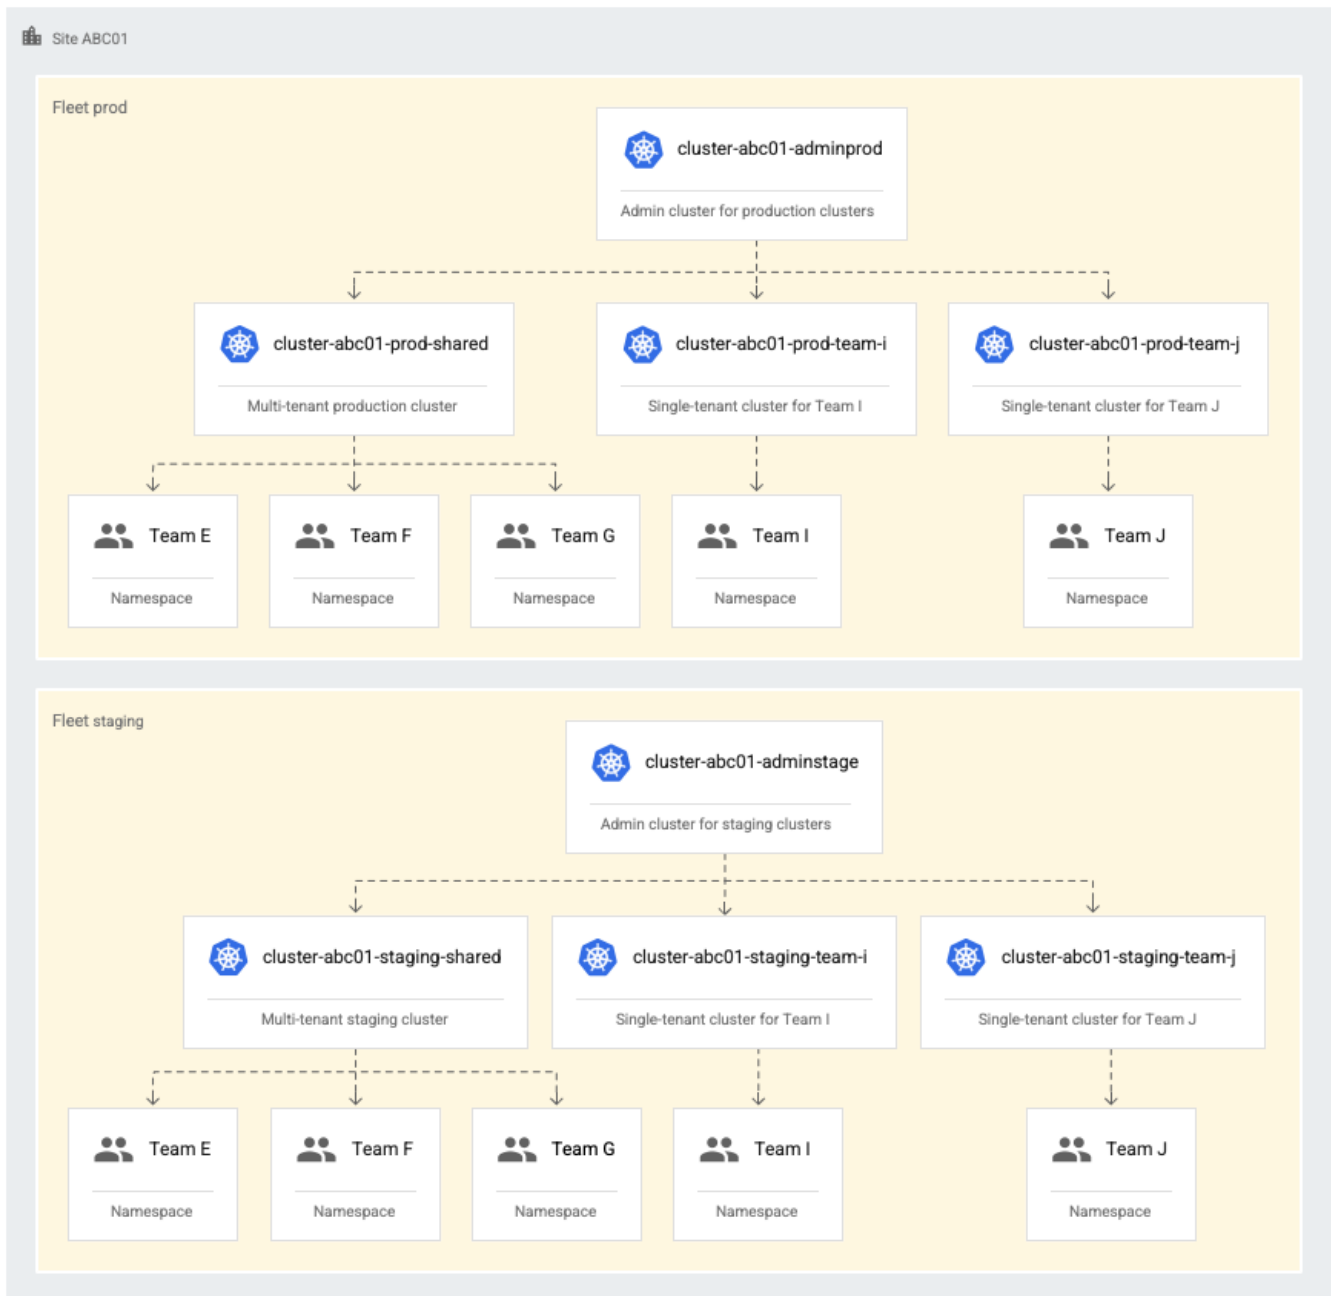


The following are examples of common deployment patterns:

- **One multi-tenant cluster for small workloads and a few single-tenant clusters for critical workloads.** It's common to have a few large applications, and many small applications. The operational overhead of a single-tenant cluster is justified for large workloads, but might not be for small workloads. For example, a large, highly replicated search application with hundreds of Pods might run in a single-tenant cluster per site, such as `cluster-abc01-prod-search`. Tens or hundreds of miscellaneous small services might then run in a multi-tenant cluster per site, such as `cluster-abc01-prod-shared`.
- **Separation of workloads to meet compliance requirements.** For example, site `abc01` might have two production clusters: `cluster-abc01-prod-pci` for PCI-DSS-compliant workloads, and `cluster-abc01-prod-nonpci` for other workloads. These clusters use different network configurations for the two sets of machines.
- **Separation of workloads by business unit.** Business units might own specific machine configurations, or want to be charged back at the granularity of machines. In this example, each business unit can have its own production clusters.

In the preceding example deployment patterns, a single admin cluster is sufficient for all production workloads. It's still recommended to have a corresponding set of staging clusters before you deploy workloads to the production clusters. The same trade offs apply to teams using Google Cloud clusters<sup>15</sup>.

The following diagram illustrates how deployments look with the addition of some single-tenant clusters:



As more user clusters are added on a site, the number of admin clusters remains fixed at two.

## Observability

Anthos distinguishes between system and application observability data. System observability data includes logging and monitoring of Anthos system components. System observability data is sent to the site's corresponding Google Cloud region.

Cloud Monitoring automatically creates a multi-project metrics scope inside the fleet host project. Fleet deployments that contain cluster members from different projects can see a unified view of cluster resource metrics coming from any of their projects.



This reference architecture recommends, and assumes, that application observability data is sent to the nearest Google Cloud region. Consider the following when deciding whether to send application observability to Google Cloud:

- Sending logging and monitoring to a 3rd party logging solution might be a smaller change to current practices.
- Sending logging data to a local logging store might offer a more familiar path to meeting compliance requirements. However, Cloud Logging supports several features for meeting compliance requirements<sup>16</sup>. These features include encrypted logs storage with Customer Managed Encryption Keys (CMEK), External Key Managers<sup>17</sup>, Access Transparency, and Access Approval.
- Preconfigured logging dashboards<sup>18,19</sup> are available for Anthos on Google Cloud.
- Prometheus compatibility for Cloud Monitoring.
- There's a cost for observability data that's ingested and stored in Google Cloud<sup>20</sup>.

## Multi-cluster management

Anthos can help increase developer velocity, reduce cost, reduce deployment risk, and increase reliability. To fully achieve these benefits, you typically have multiple clusters at a single site.

Clusters can be resized as needs change, or the hardware can be used to create short-lived clusters for specific tasks. In addition to having a large production cluster, multiple small clusters are useful for staging and development environments, and separating regulated workloads.

Staging environments, which are similar to the production environment, reduce deployment risk through representative testing and QA. Access to developer environments increases deployment velocity through earlier identification of system integration issues.

You might need to request a quota increase<sup>21</sup> if you plan to have more than 50 clusters in one fleet<sup>22</sup>.

## Gitops-based configuration management

When you first explore Anthos operations for a proof of concept, you can create and update clusters using the Google Cloud console. This approach provides guidance and validation of parameters at each step.

As you gain experience, switch to using declarative configuration to create and update clusters. The Google Cloud Terraform provider<sup>23</sup> supports creating on-premises clusters, and clusters created by the Google Cloud console can be imported into Terraform. Terraform communicates with a Google Cloud service to validate the commands and securely communicate with your admin clusters.



You can also use Anthos Config Management to declaratively configure the resources inside a cluster. These resources include Namespaces, CRDs needed for operators to be installed in the cluster, and Kubernetes RBAC roles and bindings.

When a declarative configuration is adopted, the steps to create a new cluster are as follows:

- Define, plan, and apply a new cluster in Terraform.
- Define in-cluster files for Anthos Config Management. You can define these by copying and customizing from a base configuration, or generation from a template with substitution.

Store both the Terraform `.tf` files and Anthos Config Management `.yaml` files in one or more Git repositories. If stored in a single repository, `.tf` files cluster creation and `.yaml` files for Anthos Config Management can be reviewed as a single pull request.

## Application availability

Unlike some virtual machines, containers don't use live migration if there's an infrastructure failure. Therefore, applications should be built to have redundant containers within the same cluster. This approach allows the application to tolerate failure of or maintenance operations on a single machine. Applications that need the highest level of availability should be built spanning multiple sites as well.

The following table outlines the minimum recommended application replication:

Availability requirement	Application type	Redundancy model	In-cluster redundancy	Cross-site redundancy
Low	Stateless	None	Single pod	Single site
	Stateful	None	Single pod	Single site
Medium	Stateless	Active-active	Three pods	Single site
	Stateful	Active-passive	Two pods	Single site
High	Stateless	Active-active	Three pods minimum, typically more	Multi-site
	Stateful	Active-active	Three pods minimum, typically more	Multi-site



Examples of applications with different availability requirements:

<b>Availability requirement</b>	<b>Examples</b>
Low	A developer testing environment
Medium	Asynchronous data processing, or internal services
High	Websites and services used by customers



## References

1. <https://cloud.google.com/kubernetes-engine/docs/best-practices/upgrading-clusters#multiple-environments>
2. <https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy>
3. <https://www.weave.works/blog/gitops-operations-by-pull-request>
4. <https://cloud.google.com/sdk/docs/install>
5. <https://cloud.google.com/sdk/docs/proxy-settings>
6. <https://www.jfrog.com/confluence/display/JFROG/Docker+Registry>
7. <https://cloud.google.com/container-optimized-os/docs/concepts/features-and-benefits>
8. <https://cloud.google.com/anthos/clusters/docs/on-prem/latest/how-to/plan-ip-addresses>
9. <https://cloud.google.com/anthos/clusters/docs/on-prem/latest/how-to/firewall-rules>
10. <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.networking.doc/GUID-004E2D69-1EE8-453E-A287-E9597A80C7DD.html>
11. <https://cloud.google.com/anthos/docs/resources/partner-storage>
12. <https://cloud.google.com/anthos/docs/concepts/anthos-connectivity>
13. [https://cloud.google.com/anthos/clusters/docs/bare-metal/latest/installing/registry-mirror#create\\_clusters\\_from\\_the\\_registry\\_mirror](https://cloud.google.com/anthos/clusters/docs/bare-metal/latest/installing/registry-mirror#create_clusters_from_the_registry_mirror)
14. <https://cloud.google.com/anthos/clusters/docs/bare-metal/latest/installing/package-server>
15. <https://cloud.google.com/kubernetes-engine/docs/concepts/multitenancy-overview>
16. <https://cloud.google.com/blog/products/identity-security/5-must-know-security-and-compliance-features-in-cloud-logging>
17. [https://cloud.google.com/logging/docs/routing/managed-encryption#external\\_key\\_considerations](https://cloud.google.com/logging/docs/routing/managed-encryption#external_key_considerations)
18. <https://cloud.google.com/anthos/clusters/docs/on-prem/latest/how-to/logging-and-monitoring>
19. <https://cloud.google.com/service-mesh/docs/observability/explore-dashboard>
20. <https://cloud.google.com/stackdriver/pricing>
21. <https://cloud.google.com/anthos/fleet-management/docs/quotas>
22. [https://cloud.google.com/anthos/clusters/docs/on-prem/latest/concepts/scalability#gke\\_hub](https://cloud.google.com/anthos/clusters/docs/on-prem/latest/concepts/scalability#gke_hub)
23. <https://registry.terraform.io/providers/hashicorp/google/latest/docs>