

# *C3 AI Installation Guide – Google Cloud Platform*

*Version 8.9*

*29 January 2026*

## Legal Notice

C3.ai products and services are sold subject to the C3.ai terms and conditions agreed at the time of purchase. Except as expressly permitted in that agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means the C3.ai products, services, or documentation.

The information contained herein is subject to change without notice, and is not warranted to be error-free. The information is provided by C3.ai "as-is" for informational purposes only, without representation or warranty of any kind, and C3.ai or its affiliated companies will not be liable for errors or omissions with respect to the information. The only warranties for C3.ai products and services are those that are set forth in the express warranty statements, if any, accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. If you find any errors, please report them to us in writing.

If this software or documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then use the following notice: "U.S. GOVERNMENT END USERS: C3.ai programs, including any integrated software, any programs installed on any hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government."

C3.ai materials are not intended for use in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control systems, life support machines or other equipment in which the failure of the C3.ai materials could lead to death, personal injury, or severe physical or environmental damage. C3.ai disclaims any and all liability arising out of, or related to, any such use of the C3.ai materials.

Information contained in this document regarding third party product or services does not constitute a license from C3.ai to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party. C3.ai is not responsible for and expressly disclaims all warranties of any kind with respect to third-party content, products, and services. C3.ai is not responsible for any loss, costs, or damages incurred due to the access to or use of third-party content, products, or services, except as set forth in a written agreement between you and C3.ai.

Any software coding samples included in this documentation are examples only and are not intended to be used in a production environment. The code is provided "as-is" and use of any code is at your own risk. C3.ai does not warrant the correctness or completeness of the code given herein, and C3.ai is not liable for errors or damages caused by usage of the code.

The business names used in this documentation are fictitious and are not intended to identify any real companies currently or previously in existence.

C3 AI, C3.ai, and the C3.ai logos are trademarks or registered trademarks of C3.ai, Inc. in the United States and/or other countries. All other product names, trademarks, and registered trademarks are the property of their respective owners.

## Table of Contents

<b><i>C3 AI Deployment Options: Guidance for Enterprise Customers</i></b> .....	<b>3</b>
1. C3 AI SaaS / PaaS Subscription (Preferred Standard Option) .....	4
2. Customer-Hosted, C3 AI-Managed Deployment.....	4
Summary Table .....	5
<b><i>C3 AI Installation Requirements for Google Cloud Platform</i></b> .....	<b>7</b>
Required Google cloud services.....	8
Google Cloud access requirements .....	9
Network configuration .....	11
<b><i>HashiCorp Terraform Configuration</i></b> .....	<b>15</b>
Getting started.....	15
<b><i>Installation Steps</i></b> .....	<b>16</b>
1. Enable the VPC and required GCP services .....	18
2. Validate and provide access to the cluster .....	22
3. Complete Installation.....	24

# C3 AI Deployment Options: Guidance for Enterprise Customers

C3 AI offers flexible deployment models to meet the diverse needs of enterprise customers. Selecting the appropriate deployment option is a critical decision that impacts project timelines, service-level agreements (SLAs), and roles and responsibilities (RACI). This document outlines each option, highlights key considerations, and underscores the benefits of the C3 AI SaaS/PaaS Subscription, which is the recommended approach for most organizations.

## 1. C3 AI SaaS / PaaS Subscription (Preferred Standard Option)

The standard C3 AI SaaS/PaaS (Software as a Service / Platform as a Service) subscription is the most typical deployment option to leverage the C3 AI Platform and Applications. It is a fully hosted and managed service by C3 AI in Google Cloud Platform (GCP). Customers may select their preferred GCP region for data residency.

### Key Features and Benefits:

- **Lower Total Cost of Ownership (TCO):**  
Standardized technologies and processes enable rapid deployment, streamlined support, and efficient issue resolution. C3 AI maintains specific enterprise SLAs to deliver an industry-leading service with lower TCO.
- **Reduced Operational Burden:**  
Internal teams can focus on leveraging AI applications for business value, rather than managing infrastructure setup and maintenance.
- **Scalability:**  
The SaaS/PaaS model supports seamless scaling as business needs evolve. C3 AI manages all scaling needs and capacity planning required to ensure consistently available platform and applications.
- **Security and Compliance:**  
C3 AI employs industry standard cybersecurity and access control practices to safeguard customer applications and data. C3 AI holds and maintains critical compliance attestations like SOC2, ISO27001, and FedRAMP.

### Why Choose SaaS/PaaS?

This model is the fastest, most cost-effective way to realize value from C3 AI products and generate AI-driven insights. It is recommended for organizations seeking minimal operational overhead and maximum agility.

## 2. Customer-Hosted, C3 AI-Managed Deployment

For organizations with non-standard data residency, security, or governance requirements, C3 AI supports deployments within a customer's own Google Cloud Virtual Private Cloud (VPC). C3 AI Operations manages the deployment, maintenance, and support within the customer's environment. Your organization will have responsibility for portions of the infrastructure to ensure C3 AI Operations can successfully deploy and manage C3 AI Products. Coordination with C3 AI Operations will be required for future upgrades, change and incident management activities. Additional charges may apply to support a customer-hosted deployment.

### Key Considerations:

- **Customer Responsibilities:**
  - Provision and secure the Google Cloud Virtual Private Cloud (VPC) according to internal policies.
  - Provide timely and required access to C3 AI Operations for installation and ongoing support.
  - Manage and troubleshoot infrastructure changes outside C3 AI's control that may affect availability or performance.
  - Assume all infrastructure hosting costs within the customer's cloud account.
- **Control and Access:**  
Customers retain greater control and thus greater responsibility over their GCP subscription and can limit permissions granted to C3 AI.

### When to Choose This Option:

This model is suitable for organizations with:

- Internal processes requiring direct control over cloud resources.
- Policies with non-standard local data residency, security, or governance requirements.

## Summary Table

Deployment Model	Managed By	Hosted In	Customer Responsibilities	Recommended For
<b>SaaS/PaaS Subscription (Preferred)</b>	C3 AI	C3 AI GCP Cloud	Minimal	Most organizations
<b>Customer-Hosted, C3 AI-Managed</b>	C3 AI	Customer Google Cloud Virtual Private Cloud (VPC)	VPC provisioning, access, infra costs	Regulated/controlled industries

Selecting the right deployment option is essential for project success. C3 AI strongly recommends the SaaS/PaaS Subscription for most enterprises, as it maximizes value, reduces risk, and accelerates time-to-insight.

If you have questions or require a tailored recommendation, please reach out to your C3 AI representative.

## Customer Hosted Install Requirements: Checklist

### Summary

For C3 AI to operate in customer-hosted GCP Cloud accounts, your organization must meet the following requirements consistently throughout the contract term. Deviations from the installation requirements incur additional operational fees.

You agree that your organization will allow C3 AI Operations to deploy all infrastructure required to support the C3 AI applications and platform per this specification and utilizes C3 AI deployment automation. This checklist only applies to customer hosted installations.

### Installation and Operational Management Checklist

For C3 AI Operations to deploy a cluster in a customer-hosted deployment, you must provide the following access, network setup, and infrastructure to C3 AI:


1. An IAM role called C3.AdminOps with temporary administrator privileges to a dedicated Azure subscription, AWS account, or GCP project so that C3 AI can perform tasks to set up your deployment.

C3 AI requires administrator privileges to set up an IAM policy and create a role that allows C3 AI Operations to perform installation, setup, and deployment tasks. You can remove administrator access after initial setup and may need to provide administrator access again to support new product releases.

2. An IAM role called C3.Ops with access to your dedicated GCP project, Azure subscription, or AWS account so that C3 AI can manage infrastructure.
3. If your firewall settings prevent C3 AI Operations from configuring access to required endpoints, you must allow access to the following:
  - a. Routing to and from the C3 AI network so C3 AI can receive metrics, logs, and observability data to operate the deployment.
  - b. Connectivity to endpoints that allow C3 AI product functionality.
4. New VPC and subnet with at least 1024 IP address (/22) available for C3 deployment. Each subnet must be across three availability zones.
5. Static DNS entry for C3 AI URL (for example, c3project.customer.com).
6. Private and public key (and the certificate chain if necessary) from the x509 certificate. C3 AI does not support self-signed certificates or certificates signed by an internal Certificate Authority. You must provide a certificate from a public Certificate Authority.

## C3 AI Installation Requirements for Google Cloud Platform

The C3 Agentic AI Platform integrates with core Google Cloud services such as Google Compute Engine, Virtual Private Network (VPC), and IAM, enabling cohesive security and infrastructure management.



*Figure 1. GCP Architecture with C3 Agentic AI Platform Deployment*

The C3 Agentic AI Platform requires specific Google Cloud Platform (GCP) cloud services and infrastructure for successful deployment, as well as specific access requirements for C3 AI Operations to install, administer, and upgrade the C3 Agentic AI Platform and C3 AI Applications.

Install and upgrade requirements including Bill of Materials (BOM) details can be reviewed on the documentation site <https://docs.c3.ai/versions-and-compatibility/upgrade-requirements/8.9>.

The following sections describe the specific services and access needs, including network configurations and subnet requirements, security group egress and ingress rules, and subnet-level access control lists (ACLs).

## Required Google cloud services

The table below describes the GCP Cloud infrastructure services required by the C3 Agentic AI Platform. You are required to provide the services below configured to C3 AI specifications as documented in the HashiCorp Terraform scripts.

Google Cloud Platform Service	Version	Description
Kubernetes Engine (GKE)	1.34	Operating environment responsible for the deployment, scaling, and management of the C3 Agentic AI Platform and C3 AI Applications.
Cloud Key Management Service	Current version	Scalable, centralized, fast cloud key management.
Cloud SQL for PostgreSQL	15	Relational data required for internal operations of the C3 Agentic AI Platform. See the <a href="#">GCP Cloud SQL database version policies</a> website for specific minor version supported.
Cloud Storage	Current version	Reliable and secure object storage used for the management of application and platform configuration, and other ancillary tasks.
Identity and Access Management (IAM)	Current version	Fine-grained access control and visibility for centrally managing cloud service account resources.  You must create a role called C3.Ops in your GCP deployment. Add the GCP roles defined in the HashiCorp Terraform scripts so that the C3.Ops role inherits their permissions.
Virtual Private Cloud (VPC)	Current version	Dedicated, isolated network for inter-C3Cluster communication.

## Google Cloud access requirements

The table below describes the access requirements for C3 AI Operations to install, administer, and upgrade the C3 AI Applications and C3 Agentic AI Platform.

To ensure security, operational excellence, and customer trust, C3 AI leverages an industry standard privileged access management (PAM) for infrastructure access and

authentication. All access by our C3 AI Operations team to customer hosted deployments will be managed exclusively through this PAM solution. This approach ensures every interaction with the systems is secure, fully auditible, and aligned with industry best practices.

Why PAM and what this means for you:

- **Centralized Access Control:** All infrastructure access is managed through a single, secure platform, reducing complexity and risk.
- **Enhanced Security:** PAM enforces strong authentication and zero-trust principles, ensuring only authorized personnel can access your systems.
- **Full Auditability:** Every session and command is logged, providing complete visibility for compliance and security reviews.
- **Rapid Access Revocation:** Access can be granted or revoked instantly, minimizing exposure during personnel changes or incident response.
- **Operational Efficiency:** Our Operations team benefits from streamlined workflows, reducing time-to-resolution for support and maintenance tasks.
- **Industry Best Practices:** PAM aligns with leading security frameworks and compliance standards, reinforcing trust and reliability.

Requirements of the PAM solution:

- One time use of Service Account role to deploy the cluster and for select infrastructure upgrades
- Service account created in your SSO / Active Directory used to manage the C3 AI Platform

C3 AI Operations will deploy and install the PAM solution. You can request an audit log of user access by contacting C3 AI Customer Support.

Access Requirements	Description
A dedicated GCP project	When creating the project, C3 AI requires: (1) Subscription identifier, (2) Cloud region.
Secure internet access to the GCP project	Secure, remote access via internet (VPN access is acceptable) to a bastion host from which C3 AI Operations personnel can administer cloud infrastructure and C3 AI services.
GCP service accounts	<p>Provide C3 AI Operations access to the GCP service accounts CLUSTERNAME-rsgp-c3-01 and c3-privileged</p> <p>Map the c3-privileged GCP service account to the C3 AI Kubernetes service account c3-privileged. <a href="#">See Map Cloud Provider Identity to Kubernetes Service account.</a></p>
A bastion host accessible by C3 AI Operations to manage the cluster	The bastion host will be used by C3 AI Operations to administer the C3 AI Applications and C3 Agentic AI Platform. Software utilities required to be installed on the bastion host must include: RedHat 8, Google Cloud Command-Line Interface ( <a href="#">gcloud CLI</a> ) latest version or greater, <a href="#">kubectl</a> v1.34, <a href="#">Helm</a> v3.12+, HashiCorp Terraform (>= 1.13.0), Python 3.12, and Docker.
Access to C3 AI and third-party library and image repositories	Access to C3 AI and third-party repositories for the container images, Python libraries, NodeJS libraries, and runtime billing data collection. If connecting to remote C3 AI, Python, and NodeJS artifact repositories violates security standards, the C3 Agentic AI Platform can be configured to connect to local artifact repositories such as GCP Artifact Registry, JFrog, and Anaconda Enterprise.
X.509 certificate for terminating network encryption	<p>A fully qualified domain name for C3 AI Cluster ingress configuration (for example, c3project.customer.com). You are responsible for providing the private and public key (and the certificate chain if necessary) from the x509 certificate to C3 AI. These are placed in a Kubernetes secret and used by C3 AI cluster ingress controller.</p> <p>C3 AI does not support self-signed certificates or certificates signed by an internal Certificate Authority. You must provide a certificate from a public Certificate Authority.</p>

## Network configuration

To deploy the C3 Agentic AI Platform in your own VPC, you must create the VPC following the requirements enumerated in VPC requirements section below.

### VPC requirements

Your VPC must meet the requirements described in this section to host a C3 AI cluster.

### VPC region

The GCP region where the deployment will occur. Refer to [GCP documentation](#) for a list of available regions.

### VPC sizing

The C3 Agentic AI Platform requires five (5) CIDR blocks.

IP Address Range	Association
10.0.0.0/22	Private IPs that are routable to a public-facing Internet Gateway; used by Cloud SQL (Postgres), GKE Cluster, and node pools.
172.16.0.0/18	Used for GKE services
172.20.0.0/14	Used for GKE pods
10.0.6.16/28	Used for the GKE control plane
10.0.5.0/24	Used as a proxy subnet

The VPC must have DNS hostnames and DNS resolution enabled.

### Subnets

C3 AI must have access to at least two (2) subnets for each cluster, with each subnet in a different availability region. The subnets are:

- A subnet for the GKE cluster and node pools, which has secondary CIDR ranges for GKE services (size /18) and pods (size /14)
- A proxy subnet

Each subnet must have a netmask between /15 and /22.

For GKE deployment, C3 AI configures two non-overlapping network address spaces. The primary virtual network (VNet) uses the address space 10.0.0.0/22. For the Kubernetes pod and service networks, we use secondary CIDR ranges of 172.18.0.0/15 and 172.16.0.0/18, respectively.

## VPC Networks

The VPC should have a custom static route to apply a specific 0.0.0.0/0 route and network tags to desired subnets.

## Additional networking requirements

- Workloads/Resources must be private.
- Subnets must have outbound access to the public network using a NAT gateway and internet gateway, or other similar customer-managed appliance infrastructure.
- The NAT gateway must be set up in its own subnet that routes quad-zero (0.0.0.0/0) traffic to an internet gateway or other customer-managed appliance infrastructure.

## Firewall Policies

C3 AI must have access to at least one GCP firewall policy and no more than five policies. You can reuse existing policies rather than creating new ones.

Firewall policies must follow the following rules.

### Endpoint access

If your firewall settings prevent C3 AI Operations from configuring endpoint access, you must allow outbound access to the following endpoints to allow C3 AI product functionality:

- api.github.com
- codeload.github.com
- Cloudfront.net
- conda.anaconda.org
- docker.io
- files.pythonhosted.org
- github.com
- grafana.com
- Huggingface.co
- jfrog.c3.ai
- nodejs
- npmjs.com
- prdgkemis.c3.ai (if MIS access is needed)
- pypi.org

- pypi.python.org
- quay.io
- registry.c3.ai
- registry.npmjs.org
- repo.anaconda.com
- repo.continuum.io
- telemetry.c3.ai
- vault.c3iot.io

You must allow outbound access to the following endpoints for C3 AI Monitoring. These IP addresses collect standard operational metrics:

- 44.230.42.147/32
- 54.187.151.165/32

You must allow inbound access to the following C3 AI Operations endpoints to operate the deployment:

- 12.226.154.130/32
- 13.214.249.29/32
- 18.136.19.189/32
- 34.231.113.223/32
- 34.232.23.54/32
- 34.238.215.224/32
- 34.82.144.175/32
- 52.48.79.190/32
- 54.76.64.220/32
- 70.35.33.244/32

### Egress (outbound)

- Allow all TCP and UDP access to the workspace security group (for internal traffic)
- Allow TCP access to 0.0.0.0/0 for these ports:
  - 443: for C3 AI infrastructure, cloud data sources, and library repositories

### Ingress (inbound)

- Allow all TCP and UDP access to the workspace firewall policy (for internal traffic)
- 443: for C3 AI application access
- 22: for SSH access to a bastion host

## Subnet-level network ACLs

Subnet-level network ACLs must not deny ingress or egress to any traffic.

- ALLOW ALL from Source 0.0.0.0/0. This rule must be prioritized.
- Egress:
  - Allow all traffic to the C3 AI cluster VPC CIDR, for internal traffic.
  - Allow TCP access to 0.0.0.0/0 for these ports:
    - 443: for C3 AI infrastructure, cloud data sources, and library repositories.

## HashiCorp Terraform Configuration

HashiCorp Terraform is a popular open-source tool for creating safe and predictable cloud infrastructure across several cloud providers. Terraform scripts are used to create the cloud infrastructure required by the C3 Agentic AI Platform and automate the deployment of the C3 Agentic AI Platform in your Google Cloud Platform (GCP).

**NOTE:** For C3 AI customer-managed deployments, any customization performed on the Terraform scripts must be reapplied with each version of the Terraform scripts from C3 AI.

## Getting started

In this section, you install and configure requirements to use Terraform. You then configure Terraform authentication. Following completion of this section, go to “Installation Steps” section below to deploy and configure the cloud infrastructure required by the C3 Agentic AI Platform.

### Requirements

To use Terraform to create cloud infrastructure resources required by the C3 Agentic AI Platform in your GCP organization / project, you must have the following:

- A Google Cloud organization. A Google Cloud project in the account. The GCP project name must adhere to the following restrictions:
  - For Dev and QA clusters: <stg><cloud><customerabbreviation>; in which <cloud> is gke. For example, stggkecust
  - For Production clusters: <prd><cloud><customerabbreviation>; in which <cloud> is gke. For example, prdgkecust
  - The project name should not include a hyphen and must start with a letter; only lowercase letters and numbers are allowed with no other special characters or diacritics (accented letters); and should be less than 15 characters total.

- Privileges to deploy, operate, and delete the infrastructure services. See the README.md file in the downloaded Registry folder for the most up-to-date information.
- On your local development machine, you must have:
  - The HashiCorp Terraform CLI. See [Install Terraform](#) on the Terraform website to download the binary of the required Terraform version specified in the main.tf file example in the “Installation Steps” section below. Select AMD64 or ARM64 depending on the which matches the client hardware from which you will run the Terraform scripts.
  - The gcloud CLI. See [Install the gcloud CLI](#) on the Google Cloud website.
  - The gcloud CLI, signed in through the [gcloud auth application-default login](#) command to obtain user access credentials via a web flow and put them in the well-known location for Application Default Credentials (ADC).

```
gcloud --project="" auth application-default login
```

**NOTE:** Replace:

- <project-id> with the specific Google Cloud project ID to use for this deployment. If omitted, then the current project is assumed.

For more details, see [Installing Google Cloud SDK](#) and [Authorize the gcloud CLI](#) on the Google Cloud website.

## Installation Steps

Installation of the C3 Agentic AI Platform on the Google Cloud Platform (GCP) is a multi-step process due to limitations of HashiCorp Terraform and GCP-specific configuration requirements. The installation process is the following:

1. Enable the VPC and required GCP services.
2. Validate the configuration of the VPC and required GCP services and provide C3 AI Operations access to the cluster.
3. C3 AI Operations completes the installation of the C3 Agentic AI Platform.

To create a VPC, C3 AI requires the use of HashiCorp Terraform and will provide a set of Terraform scripts to assist you in the creation of the VPC and required GCP Services.

**NOTE:** See the “HashiCorp Terraform Requirements” section above to ensure all requirements are met prior to completing the installation steps below.

A description of the Terraform modules is below. See the README.md file in the downloaded Registry folder for the most up-to-date information.

Terraform Module	Description
bootstrap	Configures the necessary Identity and Access Management (IAM) roles and policies to allow a Terraform orchestrator to deploy all services required for C3 Agentic AI Platform on GCP.
c3cluster	Coordinates the execution of all other Terraform modules.
gke-cluster	Configures GCP Kubernetes Engine (GKE), including VPC configuration, endpoint access, authorized IP addresses, and the version of Kubernetes used by the cluster.
gke-nodepool	Configures the GKE node groups, including default instance size, required subnet, and permissions assigned to each node.
Firewall policies	Configures ingress and egress security rules.
iam	Configures the required IAM roles and policies.
kms	Configures the GCP Key Management service.
network	Configures the VPC, including public and private subnets, internet gateway, CIDR blocks, DHCP, and NAT.

Terraform Module	Description
postgres	Creates a GCP Cloud SQL database (PostgreSQL) and assigns the database to the database subnet.
gke-sa	Configures the workflow identity to be used by the C3 AI cluster.
gcs	This module configures the GCP cloud storage resources to be used with the C3 AI cluster.

In addition to the required tools listed in the “HashiCorp Terraform Requirements” section, install TFSwitch, which is a tool used to switch easily between Terraform versions.

See [Install TFSwitch](#) and [TFSwitch Quick Start](#) on the TFSwitch website for more information.

## 1. Enable the VPC and required GCP services

This guide shows you how to create the cloud infrastructure services required by the C3 Agentic AI Platform using HashiCorp Terraform on GCP.

### 1.1 Run the bootstrap module

This module creates the necessary IAM roles and policies to configure the VPC and required GCP services. Configure a new `main.tf` file below, replacing the CAPITALIZED variable names with your values.

**NOTE:** The project name must adhere to the following restrictions:

- For Dev and QA clusters: <stg><cloud><customerabbreviation>; in which <cloud> is gke. For example, stggkecust
- For Production clusters: <prd><cloud><customerabbreviation>; in which <cloud> is gke. For example, prdgkecust
- The project name should not include a hyphen and must start with a letter; only lowercase letters and numbers are allowed with no other special characters or diacritics (accented letters); and should be less than 15 characters total.

**NOTE:** For more configuration options, download the Terraform modules from the C3 AI Registry folder provided by C3 AI, and view the “Inputs” section in the main README.md file.

```

module "bootstrap" {
  source      = "<c3_url>/tf-registry_c3/gcp/c3//modules/bootstrap"
  version     = ">=VERSION_NUMBER"
  project_id  = "GCP_PROJECT_ID"
  project_name = "GCP_PROJECT_NAME" # Optional : Only when project name is
  different than project id

  ## PROVIDE BELOW VARIABLES ONLY WHEN YOU ARE CREATING THE PROJECT AND
  LINKING IT TO A BILLING ACCOUNT ##

  folder_id    = "GCP_PROJECT_FOLDER_ID"
  billing_account = "GCP_BILLING_ACCOUNT"
  is_org_admin  = true
}

provider "google" {}

terraform {
  required_version = ">= 1.13.0"
  required_providers {
    google = {
      source  = "hashicorp/google"
      version = "7.15.0"
    }
  }
}

```

**NOTE:** Replace:

- **GPC\_PROJECT\_NAME** with the name of the C3 AI cluster. The cluster name must adhere to the following restrictions:
  - For Dev and QA clusters: <stg><cloud><customerabbreviation>; in which <cloud> is gke. For example, stggkecust
  - For Production clusters: <prd><cloud><customerabbreviation>; in which <cloud> is gke. For example, prdgkecust
  - The cluster name should not include a hyphen and must start with a letter; only lowercase letters and numbers are allowed with no other special characters or diacritics (accented letters); and should be less than 15 characters total.
- **VERSION\_NUMBER** with the version of the **bootstrap** module listed on the C3 AI BOM for the release version.

## 1.2 Run Terraform commands

After configuring the main.tf file, run the following Terraform commands from the same directory

```
tfswitch
terraform init
terraform plan --out out.plan
terraform apply "out.plan"
```

NOTE: If you receive a “Command not found: Terraform” after running the commands above, the `terraform` binary might not be in your path. See the [Get Started in GCP – Install Terraform](#) page on the HashiCorp Terraform website for more information.

## 1.3 Run the `c3cluster` module

This module coordinates execution of all other Terraform modules. Configure a new `main.tf` in a separate directory from the bootstrap module, replacing the CAPITALIZED variable names with your values. Note that you must assume the role created by bootstrap module.

Contact your account manager for the list of IP addresses required by C3 AI. These values will be used to update the `ip_allowlist` section below.

**NOTE:** C3 AI requires a set of CIDR blocks to be whitelisted for C3 AI Operations to deploy the C3 Agentic AI Platform.

**NOTE:** C3 AI infrastructure Terraform modules create a new VPC and subnets. If your organization must separately create these network artifacts, the Terraform module can be modified to utilize them rather than create new ones. For details, refer to the examples/existing\_network/README.md file contained in the Terraform module documentation.

```

module "c3cluster" {
  source      = "<c3_url>/tf-registry__c3/gcp/c3"
  version     = ">VERSION_NUMBER"
  c3_region   = "us-west1"
  project_id  = "GCP_PROJECT_ID"
  project_name = "GCP_PROJECT_NAME" # Only use it when project name is
  different than project id

  # Please reach out to C3 CoE to obtain C3 control Ips
  ip_allowlist = [
    {
      cidr_block    = "CIDR_TO_WHITELIST",
      display_name = "WHITELISTED_CIDR_NAME"
    },
  ]

  # Please reach out to c3 CoE to obtain the list of Domains to whitelist
  for CORS
  gcs_cors_domains = ["http://*.DOMAIN_NAME"]
}

provider "google" {}

terraform {
  required_version = ">= 1.13.0"
  required_providers {
    google = {
      source  = "hashicorp/google"
      version = "7.15.0"
    }
  }
}

```

- **GPC\_PROJECT\_NAME** with the name of the C3 AI cluster. The cluster name must adhere to the following restrictions:
  - For Dev and QA clusters: <stg><cloud><customerabbreviation>; in which <cloud> is gke. For example, stggkecust
  - For Production clusters: <prd><cloud><customerabbreviation>; in which <cloud> is gke. For example, prdgkecust
  - The cluster name should not include a hyphen and must start with a letter; only lowercase letters and numbers are allowed with no other special characters or diacritics (accented letters); and should be less than 15 characters total.
- **VERSION\_NUMBER** with the version of the `c3cluster` module listed on the C3 AI BOM for the release version.

### 1.3.1 Implement CORS policy for C3 AI Ex Machina

If the installation of the C3 Agentic AI Platform includes C3 AI Ex Machina, setting the C3 AI CORS domain is all that is necessary. The CORS policy facilitates file uploads for C3 AI Ex Machina.

See `storage_cors_domains` in the `main.tf` example above.

Also, see the `cors_rules.tf` template example in the Terraform modules for more configuration details.

After configuring the **main.tf** file, run the example below from the same directory as the new **main.tf** file

```
tfswitch
terraform init
terraform plan --out out.plan
terraform apply out.plan
```

## 2. Validate and provide access to the cluster

Validate the configuration of the VPC and required GCP services and provide C3 AI Operations access to the cluster.

After the VPC and required cloud services are configured, you are required to execute the C3 AI Cluster Validation Utility and provide the results to C3 AI. If all checks performed by the C3 AI Cluster Validation Utility pass, the VPC is suitable for C3 AI Operations to deploy the C3 Agentic AI Platform on the Kubernetes cluster.

**NOTE:** If the cluster validation utility fails, you must remediate all exceptions. All checks must pass for C3 AI Operations to be able to deploy the C3 Agentic AI Platform on your Kubernetes cluster. See the next section for details.

Once the checks are successfully completed, provide C3 AI Operations access to the cluster. Refer to the subsequent section for more information.

### 2.1 C3 AI Cluster Validation Utility

Run the C3 AI Cluster Validation Utility and provide results to C3 AI Operations

Contact the C3 AI Center of Excellence (CoE) for more information and to obtain the C3 AI Cluster Validation Utility.

Run the C3 AI Cluster Validation Utility to determine whether the infrastructure requirements are fulfilled to allow the C3 AI Operations to deploy the C3 Agentic AI Platform.

If the C3 AI Cluster Validation Utility indicates the VPC is ready for C3 AI Operations to deploy the C3 Agentic AI Platform on the Kubernetes cluster, provide the output to C3 AI Operations.

If the output indicates the VPC is not ready, remediate all exceptions and rerun the C3 AI Cluster Validation Utility.

## 2.2 Provide C3 AI Operations access to the cluster

In addition to the output of the C3 AI Cluster Validation Utility, you must provide C3 AI Operations with the following.

Title	Description
C3 AI Operations credentials	Credentials for C3 AI Operations team members.
GKE cluster name	<p>The name of the GKE cluster where the C3 Agentic AI Platform will be installed.</p> <p><b>NOTE:</b> The cluster name must adhere to the following restrictions:</p> <ul style="list-style-type: none"> <li>For Dev and QA clusters: &lt;stg&gt;&lt;cloud&gt;&lt;customerabbreviation&gt;; in which &lt;cloud&gt; is gke. For example, stggkecust</li> <li>For Production clusters: &lt;prd&gt;&lt;cloud&gt;&lt;customerabbreviation&gt;; in which &lt;cloud&gt; is gke. For example, prdgkecust</li> </ul> <p>The cluster name should not include a hyphen and must start with a letter; only lowercase letters and numbers are allowed with no other special characters or diacritics (accented letters); and should be less than 15 characters total.</p>
Region	The GCP region of the GKE cluster.
Cloud SQL Postgres endpoint	Endpoint value for the instance. From GCP Console, select your project, go to SQL and look for the SQL resource <C3_CLUSTERNAME>-pg-shared. Get the “Private IP address” and share it with C3 AI Operations.
Cloud SQL Postgres credentials	Credentials required for the C3 Agentic AI Platform to connect to PostgreSQL. From GCP Console, select your project, go to SQL and look for SQL resource <C3_CLUSTERNAME>-pg-shared. Then, click Users and click three (3) dots (...) next to Postgres user and change the password. Share this with C3 AI Operations securely.
Service Account Names	From GCP Console, go to the desired GCP project, navigate to Cloud IAM dashboard, then select Service Accounts. Get the principal of the service account starting with c3aiops and c3server and share it with C3 AI

Title	Description
	Operations.
GCP Bucket Name	Within the created project in the GCP Console, go to the <a href="#">Cloud Storage Buckets</a> page. There should only be one bucket listed.
Domain name	A fully qualified domain name for C3 AI cluster ingress configuration (for example, c3project.customer.com).
Public and private key	The private and public key (and the certificate chain if necessary) from the x509 certificate. This will be required for ingress configuration.

It is recommended that the sharing of Postgres credential and certificates occur using GCP Vault.

To grant C3 AI Operations GKE cluster administration permissions to the GCP project, run the following code snippet.

```
gcloud auth login

gcloud projects add-iam-policy-binding PROJECT_ID \
    --member=user:USER \
    --role=roles/c3dopsrole-01
```

**NOTE:** Replace:

- PROJECT\_ID with the ID of the project or fully qualified identifier for the project
- USER with the principal to add the binding for. Should be of the form user@email.

### 3. Complete Installation

C3 AI Operations completes the installation of the C3 Agentic AI Platform. With the infrastructure properly configured and GKE node pool configuration updated, C3 AI Operations will continue with the installation of the C3 Agentic AI Platform.

At the conclusion of the VPC creation and the deployment of the C3 Agentic AI Platform, the Google Cloud Platform (GCP) environment will resemble the following.




Figure 2. GCP Architecture with C3 Agentic AI Platform Deployment