C3 AI Installation Guide – Amazon Web Services

Version 8.10

3 June 2026

Legal Notices

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. Reverse engineering, disassembly, or decompilation of this C3 AI software

The information contained herein is subject to change without notice. THE INFORMATION AND DOCUMENTATION ARE PROVIDED "AS IS" AND "AS AVAILABLE," WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, ACCURACY, COMPLETENESS, AND NON-INFRINGEMENT. The information is provided by C3.ai 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 or any commitment by C3.ai to deliver any product, code, functionality, or service. 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 the following notice is applicable: "U.S. GOVERNMENT END USERS: C3.ai programs (including any integrated software, any programs embedded, installed, or activated on hardware, and modifications of such programs) and C3.ai computer documentation or other C3.ai data delivered to or accessed by U.S. Government end users are "commercial computer software," or "commercial computer software documentation," pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations, including FAR 12.212, FAR 27.405-3, and, for Department of Defense acquisitions, DFARS 227.7202-1 through 227.7202-4. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of (i) C3.ai programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), (ii) C3.ai computer documentation, and/or (iii) other C3.ai data, is subject to the rights and limitations specified in the license or subscription contained in the applicable contract. The terms governing the U.S. Government's use of C3.ai cloud services are defined by the applicable contract for such services. 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, other equipment or any other inherently dangerous applications in which the failure or malfunction of the C3.ai materials could lead to death, personal injury, or severe physical or environmental damage. If you use the C3.ai materials in any such application, you are responsible for taking all appropriate fail-safe, backup, redundancy, and other measures to ensure their safe use. 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 products 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.

Table of Contents

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)

Overview:

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 Amazon Web Services (AWS). Customers may select their preferred AWS 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 Amazon 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:**
  - Provide a dedicated AWS account for the deployment.
  - C3 AI Operations provisions the VPC via Terraform by default.
  - 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 AWS 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
SaaS/PaaS Subscription (Preferred)	C3 AI	C3 AI AWS Cloud	Minimal	Most organizations
Customer-Hosted, C3 AI-Managed	C3 AI	Customer Amazon Virtual Private Cloud (VPC)	AWS account, 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.

Pre-Installation Preparation

IMPORTANT: Before beginning the installation process, you must gather the following items from C3 AI. Failure to obtain these in advance will cause delays during infrastructure deployment.

Contact your C3 AI account manager and the C3 AI Center of Excellence (CoE) to obtain:

1. **IP Address Allowlist** -- The set of C3 AI Operations IP addresses that must be whitelisted in your firewall and security group rules. These are required for C3 AI to access and operate your deployment.
2. **CORS Domain List** -- If your installation includes C3 AI Ex Machina, you will need a list of domains to whitelist for CORS. This enables file upload functionality.
3. **C3 AI Cluster Validation Utility** -- A validation tool you will run after infrastructure deployment to confirm your environment meets all C3 AI requirements. Obtain this from the CoE before beginning installation.
4. **Terraform Module Version** -- Confirm the correct version number from the C3 AI Bill of Materials (BOM) for your release.

You must also determine the following before starting:

5. **Cluster Name** -- Your cluster name must follow a specific naming convention. See Appendix A: Cluster Naming Convention for full details.
    - Dev/QA clusters: stgaws{customerabbreviation} (e.g., stgawscust)
    - Production clusters: prdaws{customerabbreviation} (e.g., prdawscust)
    - Lowercase only, no hyphens, no special characters, must start with a letter, fewer than 15 characters total.
6. **AWS Region** -- The AWS region where the deployment will occur.
7. **AWS Account ID** -- The dedicated AWS account for the C3 AI deployment.

Customer Hosted Install Requirements: Checklist

Summary

For C3 AI to operate in customer-hosted AWS 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 AWS account 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. You must grant administrator access again for new product releases and for any subsequent infrastructure changes, because the administrator role is required to read current infrastructure state.

2. An IAM role called C3.Ops with access to your dedicated AWS account so that C3 AI can manage infrastructure. This role must remain active throughout the contract term.

3. Allow C3 AI Operations to deploy and manage a Privileged Access Management (PAM) solution. C3 AI leverages an industry-standard PAM for all infrastructure access to customer-hosted deployments. This requires:

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

4. If your firewall settings prevent C3 AI Operations from configuring access to required endpoints, you must allow access to the following:

- Routing to and from the C3 AI network so C3 AI can receive metrics, logs, and observability data to operate the deployment.
- Connectivity to endpoints that allow C3 AI product functionality. See the full endpoint list in the Network configuration section.
- CORS domain whitelisting if using C3 AI Ex Machina (obtain list from C3 AI CoE).

5. New VPC and subnet with at least 1024 IP addresses (/22) available for C3 deployment. Each subnet must be across three availability zones. A secondary CIDR of 172.0.0.0/16 is also required for EKS pod networking.

6. Access to C3 AI and third-party container image and library repositories (or local alternatives such as AWS ECR, JFrog, or Anaconda Enterprise).

7. Static DNS entry for C3 AI URL (for example, c3project.customer.com).

8. Private and public key (and the certificate chain if necessary) from the x509 certificate. Certificates issued by a public or an internal Certificate Authority are supported. Coordinate with the C3 AI Center of Excellence if you plan to use self-signed or internal-CA certificates so that the required trust chain can be pre-staged in the cluster.

9. EKS version 1.31 or newer. The current default in the Terraform module is 1.34.

10. After infrastructure deployment, execute the C3 AI Cluster Validation Utility and provide the results to C3 AI Operations. All checks must pass before C3 AI Operations can proceed with platform installation.

11. After infrastructure deployment, provide C3 AI Operations with all required cluster information. See Appendix B: Post-Deployment Information Handoff for the complete list.

NOTE: Installation uses a two-stage Terraform deployment: first the bootstrap module, then the c3cluster module.

C3 AI Installation Requirements for Amazon Web Services

The C3 Agentic AI Platform integrates with core AWS services like Amazon EC2, VPC, and IAM, enabling cohesive security and infrastructure management. The platform also supports AWS-native tools such as Amazon S3 bucket for durable backups.

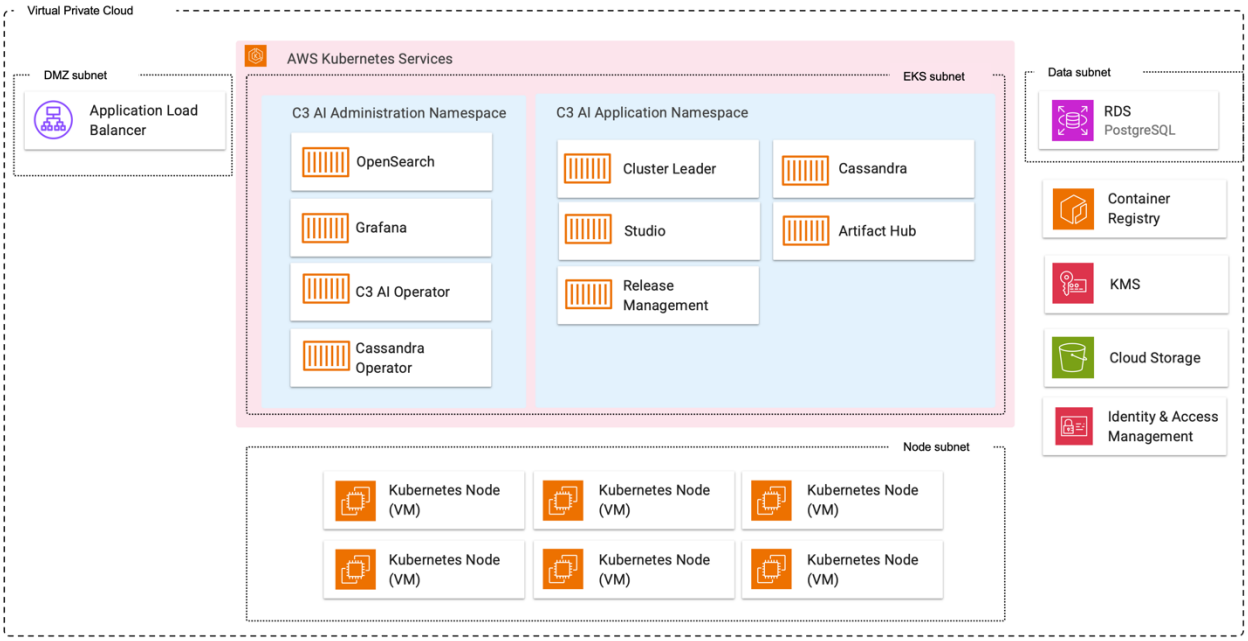

The C3 Agentic AI Platform requires specific AWS cloud services and infrastructure for successful deployment, as well as specific access requirements for C3 AI Operations to install, administer, and upgrade the C3 Agentic AI Platform and C3 AI Applications.

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/8.10.

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 Amazon cloud services

The table below describes the Amazon 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.

Amazon Cloud Service	Version	Description
Elastic Kubernetes Engine (EKS)	1.34	Operating environment responsible for the deployment, scaling, and management of the C3 Agentic AI Platform and C3 AI Applications.

Secrets Manager	Current version	Scalable, centralized, fast cloud key management.
RDS (PostgreSQL)	15	Relational database service (RDS) required for internal operations of the C3 Agentic AI Platform.
S3	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.
Virtual Private Cloud (VPC)	Current version	Dedicated, isolated network for inter-C3Cluster communication.
EC2	Current version	Compute instances required by Amazon EKS.

Amazon 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 auditable, 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 an IAM 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
- An additional **/28 CIDR IP block** for the C3 Software-Defined Perimeter/Management (SDM) peered VPC

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 Amazon sub account	When creating the project, C3 AI requires: (1) Account identifier, (2) Cloud region.
IAM user account for each C3 AI Operations member	C3 AI Operations personnel must be granted IAMReadOnlyAccess and IAMUserChangePassword managed policies. You must create a role called C3.Ops in your AWS deployment. Add the IAM policy specifications defined in the HashiCorp Terraform scripts.
AWS roles	Provide C3 AI Operations access to the AWS roles CLUSTERNAME-rsgp-c3-01 and CLUSTERNAME-c3-privileged. Map the CLUSTERNAME-c3-privileged AWS role to the C3 AI Kubernetes service account c3-privileged. C3 AI Operations requires access to the IAM roles and Kubernetes service accounts provisioned by Terraform during cluster deployment. See Section 2 (Create and attach policy to C3 AI Operations IAM users) for IAM policy setup, and Appendix B (Post-Deployment Information Handoff) for the roles and ARNs handed off after deployment.
Secure internet access to the Amazon sub account	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.
A bastion host accessible by C3 AI Operations to manage the cluster	The bastion host is used by C3 AI Operations to administer the C3 AI Applications and C3 Agentic AI Platform. Deploy the bastion host in a private subnet of the C3 AI cluster VPC; C3 AI Operations connects to it through the PAM solution, so no public IP is required. Software utilities required on the bastion host must include: RedHat 8, AWS Command Line Interface (CLI) (latest version), kubectl v1.34, Helm v3.12+, Helmfile plugin, HashiCorp Terraform (>=1.13.0), TFSwitch, Python 3.12, Docker, yq, and jq.
Access to C3 AI 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, AWS Container registry, JFrog, and Anaconda Enterprise).
X.509 certificate for terminating network encryption	A fully qualified domain name for C3 Cluster ingress configuration (for example, c3project.customer.com). You are responsible for providing the public certificate with the complete chain and the private key to C3 AI. These are placed in a Kubernetes secret and used by C3 AI cluster ingress controller. Certificates issued by a public or an internal Certificate Authority are supported; coordinate with the C3 AI Center of Excellence if you plan to use self-signed or internal-CA certificates so that the required trust chain can be pre-staged in the cluster.

Network configuration

To deploy the C3 Agentic AI Platform in your own VPC, you must create the VPC following the requirements enumerated in the 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 Amazon region where the deployment will occur. Refer to AWS documentation for a list of available regions.

VPC sizing

The C3 Agentic AI Platform requires two (2) CIDR blocks.

VPC IP address ranges

IP Address Range	Description
/22 range	Private IPs that are routable to a public-facing Internet Gateway; used by RDS (Postgres), EKS Cluster, and node pools.
172.0.0.0/16	Non-routable space, used by EKS pods.

DNS

The VPC must have DNS hostnames and DNS resolution enabled.

Subnets

The Terraform module creates the following subnets within the VPC for each cluster (one subnet per availability zone, for three availability zones):

Subnet Type	Count	Default Prefix	Purpose
DMZ	3 (one per AZ)	/25	Public load balancing
Data	3 (one per AZ)	/25	RDS PostgreSQL (private)
EKS	3 (one per AZ)	/24	EKS cluster and node pools (private)
EKS Pod	3 (one per AZ)	/17 27	Pods running in the EKS cluster (private, non-routable)

For subnet sizing options, see the "Inputs" section of the main README.md file in the downloaded C3 AI Registry folder.

- For EKS 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.

Subnet route table

The route table for workspace subnets must have quad-zero (0.0.0.0/0) traffic that targets the appropriate network device.

Additional subnet requirements

- Subnets must have outbound access to the public network using a cloud native NAT gateway and internet gateway.
- 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.

Security groups

C3 AI must have access to at least one AWS security group and no more than five security groups. You can reuse existing security groups rather than create new ones.

Security groups must include the rules described in the following subsections: Endpoint access, Egress (outbound), Ingress (inbound), and Subnet-level network ACLs.

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. These endpoints provide access to container registries, language-runtime package repositories (Python, NodeJS, Anaconda), C3 AI artifact servers, and the vault that secures platform credentials. Blocking any of them prevents platform installation, upgrades, or runtime operations. Contact the C3 AI Center of Excellence for a per-endpoint justification if required for security review.

- conda.anaconda.org – Package repository for Conda environments and dependencies
- files.pythonhosted.org – File hosting service for Python packages distributed via PyPI
- github.com – Source code hosting and version control platform
- c3ai.grafana.net – C3 AI's Grafana-hosted monitoring and observability dashboards
- huggingface.co – Repository for pre-trained machine learning models and datasets
- jfrog.c3.ai – C3 AI's internal JFrog Artifactory instance for artifact and package management
- nodejs.org – Official Node.js runtime downloads and documentation
- npmjs.org – Package registry for Node.js/JavaScript dependencies
- prdgkemis.c3.ai – C3 AI endpoint for MIS (Management Information System) access (if required)
- pypi.org – Primary Python package index for installing Python libraries
- pypi.python.org – Legacy Python package index mirror, an alias for PyPI
- registry.c3.ai – C3 AI's private container and artifact registry
- repo.anaconda.com – Anaconda's repository for curated data science packages
- repo.continuum.io – Legacy Continuum Analytics (now Anaconda) package repository
- Cloudfront.net - Amazon CloudFront's CDN domain, used to serve cached content (files, packages, assets) from AWS edge locations worldwide.
- telemetry.c3.ai – C3 AI endpoint for collecting platform telemetry and usage data

- vault.c3iot.io – C3 AI's HashiCorp Vault instance for secrets and credentials management

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 443 outbound to the endpoints listed in the Endpoint access section.

Ingress (inbound)

- Allow all TCP and UDP access to the workspace security group (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 443 outbound to the endpoints listed in the Endpoint access section.

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 AWS account.

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, you go to the "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 AWS account, you must have the following:

- An AWS account.
- 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 AWS CLI
    - The eksctl command-line tool. See Install eksctl on the Amazon EKS website.
- 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.
- The following environment variables:
    - AWS_ACCESS_KEY_ID, set to the value of your AWS user's access key ID. See Programmatic access in the AWS General Reference.
    - AWS_SECRET_ACCESS_KEY, set to the value of your AWS user's secret access key. See Programmatic access in the AWS General Reference.
    - AWS_REGION, set to the value of the AWS Region code for your AWS account. See Regional endpoints in the AWS General Reference.

NOTE: Terraform state file placement. Store the Terraform state file in a secure remote backend rather than on a local workstation. Typical placements include an S3 bucket (with versioning enabled) plus a DynamoDB table for state locking. For custom, FieldOps, or Fed deployments, coordinate with C3 AI FieldOps on state-file location and backend configuration. Losing or corrupting the state file complicates subsequent upgrades and cluster teardown.

Installation Steps

Installation of the C3 Agentic AI Platform on AWS is a multi-step process due to limitations of Terraform and AWS-specific configuration requirements. The installation process is the following:

1. Create the VPC and required AWS services.
2. Grant C3 AI Operations access to EKS as a Kubernetes administrator.
3. C3 AI Operations installs and configures required EKS options.
4. Update your EKS node pool configuration.
5. Validate the VPC and configuration of required AWS services.
6. C3 AI Operations completes the installation of the C3 Agentic AI Platform.

See Appendix C: Installation Timeline for estimated durations and responsibilities for each stage.

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 AWS Services. If you are unfamiliar with Terraform, review their Get Started -- AWS documentation.

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

NOTE: See the Pre-Installation Preparation section to ensure you have obtained all required items from C3 AI before beginning.

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 IAM roles and policies to allow a Terraform orchestrator to deploy all services required by the C3 Agentic AI Platform on AWS.
c3cluster	Coordinates the execution of all other Terraform modules.
eks-addons	This module will deploy and configure EKS managed add-ons.
eks-cluster	Configures AWS Elastic Kubernetes Service (EKS), including VPC configuration, endpoint access, authorized IP addresses, and the version of Kubernetes used by the cluster.
eks-nodepool	Configures the AWS EKS node groups, including default instance size, required subnet, and permissions assigned to each node.
firewall	Configures ingress and egress security rules.
iam	Configures the required IAM roles and policies.
kms	Configures the AWS Key Management service.
network	Configures the VPC, including public and private subnets, internet gateway, CIDR blocks, DHCP, and NAT.
postgres	Creates an AWS RDS database and assigns the database to the database subnet.
vault	Seeds the application role and secret in the AWS Secrets Manager.
S3	This module configures the AWS S3 buckets to be used with the C3 AI cluster.
vpc-endpoint	Configures AWS VPC endpoints for private connectivity to AWS services (such as S3 and Secrets Manager) without traversing the public internet.

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. Create the VPC and required AWS services

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

1.1 Run the bootstrap module

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

NOTE: The cluster name must adhere to the naming convention described in Appendix A: Cluster Naming Convention.

NOTE: For more configuration options, download the Terraform module 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/aws/c3//modules/bootstrap"
  version      = "VERSION_NUMBER"
  cluster_name = "CLUSTER_NAME"  # Replace with name of c3 deployment
  account_id   = "AWS_ACCOUNT_ID"
  region       = "REGION"

  # Trusted identities ARN who will be allowed to assume the infrastructure c
reation role
  trusted_identifier_arns = ["TRUSTED_IDENTITY_ARN"]
}

provider "aws" {}

terraform {
  required_version = ">=1.13.0"
  required_providers {
    aws = {
      source  = "hashicorp/aws"
      version = "6.28.0"
    }
  }
}
```

NOTE: Replace:

- CLUSTER_NAME with the name of the C3 AI cluster. See Appendix A for naming rules.

- VERSION_NUMBER with the version of the bootstrap module listed on the C3 AI BOM for the release version.
- TRUSTED_IDENTITY_ARN -- List of AWS accounts and IAM users who are authorized to manage C3 infrastructure.

1.2 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 AWS -- 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 IAM_ROLE_NAME role created by the bootstrap module, which is typically formatted as follows: ${CLUSTER_NAME}_c3icrole-01 in the default Terraform scripts. For more details see https://repost.aws/knowledge-center/iam-assume-role-cli.

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. You should have obtained these during the Pre-Installation Preparation step.

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/aws/c3"

  version      = "VERSION_NUMBER"
  cluster_name = "CLUSTER_NAME"  # Replace with name of c3 deployment
  c3_region    = "REGION"
  ip_allowlist = [
    {
      cidr_blocks  = [
        "CIDR_TO_WHITELIST",
      ],
      display_name = "WHITELISTED_CIDR_NAME"
    }
```

```
  ]

  # Please reach out to C3 CoE to obtain the list of Domains to whitelist for
CORS
  s3_cors_domains = ["http://*.DOMAIN_NAME"]
}

provider "aws" {}

terraform {
  required_version = ">=1.13.0"
  required_providers {
    aws = {
      source  = "hashicorp/aws"
      version = "6.28.0"
    }
  }
}
```

NOTE: Replace:

- CLUSTER_NAME with the name of the C3 AI cluster. See Appendix A for naming rules.
- 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 s3_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 creating the main.tf file, run the example below from the same directory as the main.tf file:

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

2. Grant C3 AI Operations access to EKS as a Kubernetes administrator

2.1 Create and attach policy to C3 AI Operations IAM users giving permissions to assume the Infrastructure Management IAM role and generate kubeconfig

Create the policy by doing the following:

1. In the AWS Identity and Access Management (IAM) dashboard, go to IAM > Policies.
2. In the Permissions tab, select Create Policy to enter the following JSON.

NOTE: See the iam_policies.tf template example in the Terraform modules for more configuration details.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "sts:AssumeRole"
      ],
      "Resource": [
        "arn:aws:iam::ACCOUNTID:role/CLUSTER-c3icrole-01"
      ]
    },
    {
      "Sid": "ManageOwnAccessKeys",
      "Effect": "Allow",
      "Action": [
        "iam:CreateAccessKey",
        "iam:DeleteAccessKey",
        "iam:GetAccessKeyLastUsed",
        "iam:GetUser",
        "iam:ListAccessKeys",
        "iam:UpdateAccessKey",
        "iam:TagUser"
      ],
      "Resource": "arn:aws:iam::ACCOUNTID:RESOURCE_TYPE"
    }
  ]
}
```

NOTE: Replace:

- CLUSTER with the name of the EKS cluster. See Appendix A for naming rules.
- REGION with AWS region code associated with the EKS cluster.

- ACCOUNTID with ID of the AWS account that owns the resource, without the hyphens. For example, 123456789012.
- RESOURCE_TYPE with user, group, or the resource type to be assigned to the user. Use the value set in Step 2.1.

Attach the JSON policy to the C3 AI Operations IAM user by doing the following:

NOTE: Contact the CoE for the specific C3 AI Operations IAM users that need to be added.

1. In the Entities Attached tab, select IAM Users from the Entity Type drop-down menu, and enter the applicable Entity Names.
2. Then, click Attach Policy.

Once creation of the EKS identity mapping is complete, notify C3 AI Operations and they will complete the configuration of the EBS CNI driver, Kubernetes Autoscaler, Kubernetes Metrics Server, and the Calico network policy engine for EKS.

3. Validate the VPC and configuration of required AWS services

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 Common Issues and Troubleshooting section for guidance.

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

3.1 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. You should have obtained this during the Pre-Installation Preparation step.

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 VNet is not ready, remediate all exceptions and rerun the C3 AI Cluster Validation Utility.

3.2 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 information listed in Appendix B: Post-Deployment Information Handoff.

Title	Description
C3 AI Operations credentials	Credentials for C3 AI Operations team members.
EKS cluster name	By default, CLUSTER-kube-01. Confirm this in the AWS console by going to "Elastic Kubernetes Service" and identifying the newly created cluster. See Appendix A for cluster naming rules.
Region	The AWS region associated with the EKS cluster.
Role-Arn to get Kubeconfig	The ARN of the role created in Step 3; likely, arn:aws:iam::ACCOUNTID:role/C3_CLUSTER_ID-kubeconfig-01. Copy it from IAM -> Roles -> find the role you created previously.
AWS Postgres endpoint	From AWS console, go to RDS Services, locate <C3_CLUSTER_ID>-pg-shared and provide the Endpoint value for the instance.
AWS Postgres Admin password (after changing it)	From RDS services, locate <C3_CLUSTER_ID>-pg-shared and hit the Modify button at the top right. Configure a new password by pressing "Auto generate a password" or by filling in new password and take note of it. Or, if you do not have permissions to do so, run aws rds modify-db-instance --db-instance-identifier --master-user-password <pwd> as the role C3_CLUSTER_ID-c3icrole-01.
EKS Pod Subnets with Availability Zones	Go to VPC -> Subnets. Search for subnets with name having a prefix of C3_CLUSTER_ID-sn-ekspod, taking note of the Subnet ID and Availability Zone of each one.
EKS security group ID	Go to VPC -> Security Groups. Search for security groups with a prefix of C3_CLUSTER_ID-sg-eks -- this should show one security group. Take note of its Security group ID.
S3 Bucket name	By default, ACCOUNTID--CLUSTERNAME. Confirm this in the AWS console by going to S3 and identifying this bucket.
Domain name	A fully qualified domain name for C3 AI Cluster ingress configuration (for example, c3project.customer.com).
Public and private key	The public certificate with the complete chain and the private key. This will be required for ingress configuration.

4. Complete Installation

C3 AI Operations completes the installation of the C3 Agentic AI Platform.

With the infrastructure properly configured, 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 AWS Cloud environment will resemble the architecture shown in Figure 1 above.

Common Issues and Troubleshooting

The following issues are commonly encountered during infrastructure deployment. Review these before contacting C3 AI support.

EKS cluster cannot access the internet

Cause: NAT Gateway or Internet Gateway is misconfigured, or route tables are not pointing to the correct targets.

Solution: Verify that the NAT Gateway is deployed in its own dedicated subnet with a route to the Internet Gateway. Verify that private subnet route tables have a 0.0.0.0/0 route pointing to the NAT Gateway.

C3 AI Operations cannot access the bastion host

Cause: C3 AI Operations IP addresses are not whitelisted in security groups or network ACLs.

Solution: Verify that all C3 AI Operations IP addresses (obtained during pre-installation preparation) are included in the security group ingress rules for port 22 (SSH) and port 443. Verify that subnet-level network ACLs are not denying traffic.

Certificate validation fails

Cause: The issuing Certificate Authority's trust chain is not present in the cluster.

Solution: If using a certificate from a public Certificate Authority, verify the certificate chain file is included alongside the certificate. If using a certificate from an internal Certificate Authority or a self-signed certificate, coordinate with the C3 AI Center of Excellence to pre-stage the required trust chain in the cluster before re-submitting the certificate.

Cluster validation utility reports failures

Cause: One or more infrastructure requirements are not met.

Solution: Review the specific validation failures reported by the utility. Remediate each issue according to the error messages. Re-run the validation utility after remediation. All checks must pass before C3 AI Operations can proceed.

RDS PostgreSQL is inaccessible from EKS

Cause: Security group rules or subnet configuration are preventing traffic between the EKS subnets and the data subnets where RDS is deployed.

Solution: Verify that the RDS instance is deployed in the data subnets. Verify that security group rules allow traffic from the EKS security group to the RDS security group on port 5432.

Terraform apply fails with permission errors

Cause: You are not assuming the correct IAM role, or the role does not have sufficient permissions.

Solution: Verify that you are assuming the IAM role created by the bootstrap module (typically {CLUSTER_NAME}_c3icrole-01). See https://repost.aws/knowledge-center/iam-assume-role-cli for instructions on assuming a role via the AWS CLI.

Terraform apply fails with "Command not found: Terraform"

Cause: The Terraform binary is not in your system PATH.

Solution: Install Terraform using TFSwitch or download the correct binary from the HashiCorp website. Run tfswitch in your working directory to select the correct version. See the Get Started in AWS -- Install Terraform page on the HashiCorp Terraform website.

Appendix A: Cluster Naming Convention

The cluster name is used throughout the deployment to name AWS resources, IAM roles, security groups, subnets, and Kubernetes objects. Choosing a name that follows the required convention is critical. Changing a cluster name after deployment requires rebuilding infrastructure.

Format

Environment	Format	Example
Dev / QA	stgaws{customerabbreviation}	stgawscust
Production	prdaws{customerabbreviation}	prdawscust

Rules

All of the following rules are strictly enforced:

1. **Must start with a letter** -- cannot start with a number.
2. **Lowercase only** -- no uppercase letters are allowed.
3. **No hyphens** -- the - character is not allowed.
4. **No special characters** -- only lowercase letters (a-z) and numbers (0-9).
5. **No diacritics** -- no accented letters (e.g., e, n, u).
6. **Maximum length** -- 15 characters total.
7. **Required prefix** -- must begin with stg (for dev/QA) or prd (for production).
8. **Required cloud indicator** -- must include aws immediately after the environment prefix.

Valid examples

- stgawscust
- prdawsacme
- stgawstest01
- prdawscorp

Invalid examples

Name	Problem
stg-aws-cust	Contains hyphens
StgAwsCust	Contains uppercase letters
stgazurecust	Wrong cloud identifier (should be aws)
awsstgcust	Wrong prefix order (environment must come first)
stgawscustomerlongname	Exceeds 15 characters
1stgawscust	Starts with a number
stgaws-cust	Contains a hyphen

Appendix B: Post-Deployment Information Handoff

After infrastructure deployment is complete, you must provide C3 AI Operations with the following information. Gathering this information in advance will accelerate the platform installation.

#	Information Item	How to Obtain	Example
1	C3 AI Operations IAM user credentials	Create IAM users per C3 AI CoE request	IAM usernames
2	EKS cluster name	AWS Console -> Elastic Kubernetes Service -> identify cluster	stgawscust-kube-01
3	AWS region	Your deployment region	us-east-1
4	Role ARN for kubeconfig	IAM -> Roles -> find role from bootstrap module	arn:aws:iam::123456789012:role/stgawscust-c3icrole-01
5	PostgreSQL endpoint	RDS Services -> locate {cluster}-pg-shared -> Endpoint	stgawscust-pg-shared.us-east-1.rds.amazonaws.com
6	PostgreSQL admin password	RDS -> Modify -> generate new password (or via AWS CLI)	Secure password string
7	EKS Pod subnets with AZs	VPC -> Subnets -> search {cluster}-sn-ekspod -> note Subnet ID and AZ	Subnet ID + AZ mapping for each
8	EKS security group ID	VPC -> Security Groups -> search {cluster}-sg-eks	sg-0abc123def456
9	S3 bucket name	S3 Console -> identify bucket	123456789012--stgawscust
10	Domain name (FQDN)	Your designated domain	c3project.customer.com
11	SSL certificates	Your certificate files from public CA	Private key, public key, certificate chain
12	Cluster validation results	Output from C3 AI Cluster Validation Utility	Validation report file

Support and Escalation

For questions during installation:

- Contact the C3 AI Center of Excellence (CoE). Obtain contact details from your C3 AI account manager.

For critical installation blockers:

- Contact your assigned C3 AI account manager directly.

For post-installation platform issues:

- Contact C3 AI Customer Support. Obtain portal access details from your C3 AI account manager.
