C3 AI Federal Installation Guide – Amazon Web Services GovCloud

Version 8.11

23 July 2026

AWS GOVCLOUD (US) EDITION

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.

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. The use of any third-party trademark in this document is for identification purposes only and does not imply endorsement by, or affiliation with, the trademark owner.
All rights not expressly granted in the applicable license or subscription agreement, or in this notice, are reserved by C3.ai, Inc. and its affiliates.

Table of Contents

AWS GovCloud Overview

AWS GovCloud (US) is an isolated AWS region designed specifically to host sensitive data and regulated workloads in the cloud. It helps U.S. government agencies, defense contractors, and regulated industries meet applicable compliance requirements including FedRAMP High, ITAR, DoD IL2–IL5, CJIS, HIPAA, and CMMC.

This guide covers the C3 AI supported deployment model for installation of the C3 Agentic AI Platform in an AWS GovCloud (US) customer-hosted environment. For C3 AI infrastructure scripts and HashiCorp Terraform deployment documentation, please reference the C3 AI Infrastructure Scripts / HashiCorp Terraform Deployment – Amazon Web Services GovCloud guide.

AWS GovCloud Compliance Requirements

The C3 Agentic AI Platform deploys into AWS GovCloud (US), the isolated AWS partition for regulated U.S. government workloads. GovCloud is operated exclusively by vetted U.S. persons on U.S. soil, supports ITAR/EAR data handling, and carries its own FedRAMP High and DoD Impact Level 2, 4, and 5 authorizations. This means your installation inherits a substantial, independently assessed control baseline from the underlying cloud infrastructure, giving your organization a strong head start toward an Authority to Operate (ATO).

To make the most of that inheritance, your environment needs to be configured correctly from day one. This section walks you through what to verify and set up before installation begins.

What we bring to the deployment

We deliver pre-hardened, pre-configured images STIGed to the latest available benchmark for your target environment. This eliminates a significant portion of the manual hardening work typically required before an ATO review, and ensures your deployment starts from a known, compliant baseline rather than building from scratch. We also deliver Software Bill of Materials (SBOM) documenting every software component and dependency on the platform, giving your security and compliance teams full visibility into what's running in your environment. Ports, Protocols, and Services Management (PPSM) documentation detailing all network facing services, the ports they use, and the protocols they communicate over making it ready to support your authorization boundary review and firewall configuration. Together, these artifacts reduce the documentation burden on your team and give your assessor and Authorizing Official the evidence they need to move through the ATO process efficiently.

Know what you own

AWS GovCloud follows the Shared Responsibility Model: AWS secures the cloud (infrastructure, hypervisor, and authorized services under its FedRAMP High and IL authorizations); you secure what runs *in* it — platform configuration, data, identities, encryption, and network. Your team will need to implement customer owned controls to the applicable NIST SP 800-53 Rev. 5 baseline and document them

for your assessor and Authorizing Official. The AWS Customer Responsibility Matrix is the authoritative reference for understanding exactly where that line falls.

Confirm before you begin

- Target accounts are in the GovCloud (US) partition (aws-us-gov).
- Your workload's data classification fits the GovCloud boundary (FedRAMP High / IL2–IL5).
- All personnel with access meet the citizenship and clearance requirements for your data.
- You have an approved authorization boundary diagram and data-flow description.

Configure to preserve accreditation

- **Authorized services only.** Provision on AWS services that are in-boundary and authorized at your required FedRAMP/IL level; an unauthorized service breaks the inheritance chain.
- **FIPS-validated encryption.** Encrypt data at rest with AWS KMS customer-managed keys and in transit via TLS to GovCloud FIPS endpoints.
- **Least-privilege identity.** Enforce IAM least-privilege with MFA and role-based access, avoid long-lived credentials, and integrate your identity provider.
- **Network isolation.** Deploy in a private VPC with no default internet exposure, segmented security groups, and logged egress; route admin access through approved bastion or PrivateLink paths.
- **Monitoring from day one.** Enable CloudTrail, AWS Config, GuardDuty, Security Hub, or equivalent services during installation to capture required evidence.

Position for ATO

A correct installation maps your control inheritance cleanly to the AWS FedRAMP High / IL package, with customer-owned controls implemented and evidenced. This lets your Authorizing Official inherit the assessed infrastructure controls and focus on the ATO review on the platform and data layer — shortening your path to authorization. Keep your inheritance mapping, STIG results, and monitoring feeds current afterward to sustain it under NIST SP 800-137.

AWS GovCloud Deployment Options: Guidance for Enterprise Customers

C3 AI offers flexible deployment models to meet the diverse needs of enterprise customers. In AWS GovCloud (US) environments, only the Customer-Hosted, C3 AI-Managed deployment model and SMX-hosted, C3 AI-Managed deployment model is available.

Option 1 - SaaS / PaaS Subscription

C3 AI and SMX have a Strategic Integration Partnership focused on delivering and hosting the C3 Agentic AI Platform in the federal government's most secure environments, pairing the platform with SMX's Elevate® platform and Managed Services. The collaboration is focused on enabling deployments that meet FedRAMP Moderate requirements today, with FedRAMP High and DoD Impact Level (IL) 5 in the near future, giving federal agencies, the DoD, and regulated industries an accredited pathway to run AI in their most sensitive operations.

For customers, the benefit is inheritance and acceleration. Hosting C3 deployments in the SMX Elevate environment lets customers inherit independently assessed infrastructure and platform controls, narrow their Authority to Operate (ATO) review to the application and data layer, and reach authorization faster than a self-managed build allows. SMX's certified personnel support the deployment, configuration, and continuous monitoring needed to sustain these accreditation levels — so customers gain both the AI platform and the accredited, managed environment to run it securely.

Due to the nature of the accreditations for FedRAMP Moderate and IL-5, there are some requirements that the customer must provide or meet.

Requirement	Owner (FedRAMP)	Owner (IL5)
Provide OIDC or SAML IdP	Customer	Customer
Provide TLS certificate and key for SSL termination	Customer	Customer
Configure and provide DNS record	Customer	Customer
Configure and deploy AWS GovCloud infrastructure into SMX AWS GovCloud account	C3 AI / SMX	C3 AI / SMX
Configure and deploy the C3 AI Platform	C3 AI / SMX	C3 AI / SMX

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.

FedRAMP Moderate and IL-5 Accreditations

Within the C3 AI–SMX partnership, fully managed C3 AI deployments are hosted in SMX's accredited Elevate® environment. SMX owns the accreditation and operation of that hosting environment, while the customer remains responsible for a defined set of controls at the identity, data, and agency-authorization layer. The split below reflects a standard shared-responsibility model for a managed FedRAMP Moderate or Impact Level 5 deployment; confirm the exact boundaries against SMX's Customer Responsibility Matrix before authorization.

SMX provides (accreditation, management, and access foundation):

- The FedRAMP Moderate–accredited or Impact Level 5-accredited hosting environment and the underlying control implementation, so the customer inherits an independently assessed infrastructure and platform control set rather than building it.
- Full management of the C3 AI deployment — provisioning, configuration to the appropriate baseline, patching, and hardening of the hosted environment.
- Continuous monitoring required to sustain the authorization: vulnerability scanning, logging and security event monitoring, configuration management, and the recurring continuous-monitoring reporting.
- Certified personnel to operate the environment and supply the evidence artifacts (control implementation details, scan results, and supporting documentation) that the customer's Authorizing Official relies on.
- The secure access infrastructure and enforcement points — encrypted connectivity, network isolation, and access-control mechanisms within the managed boundary.

The customer provides (their responsibility):

- **Identity provider (IdP).** The customer supplies and manages their own identity provider and federates it to the environment for authentication. SMX enforces access at the platform, but the source of identity — user accounts, group membership, MFA policy, and provisioning/deprovisioning — remains the customer's responsibility.
- User authorization and role assignment within the platform, consistent with least-privilege.
- The data ingested into the platform, including its classification, handling, and any data-level controls or agreements.
- Application- and use-case-specific configuration the customer builds on top of the managed platform.
- Data Owner approval for application- and use-case-specific data access and utilization.

Option 2 – Customer-Hosted, C3 AI-Managed

For organizations with non-standard data residency, security, or governance requirements, C3 AI supports deployments within a customer's own AWS GovCloud Account. 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:

1. Infrastructure
    - Customer provides AWS GovCloud account access with appropriate IAM permissions for C3 AI personnel
    - (Optional) Customer deploys / provides VPC, subnets, security groups, etc. configured to customer requirements
    - C3 AI infrastructure script deployment performed by C3 AI personnel
    - Customers retain greater control and thus greater responsibility over their AWS subscription and can limit permissions granted to C3 AI.
2. C3 AI Platform deployment (Helm/Kubernetes) performed by C3 AI personnel

Why choose Customer-Hosted, C3 AI-Managed?

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)	SMX + C3 AI	SMX AWS GovCloud account	IdP + data access	FedRAMP Moderate + IL-5 deployments
Customer-Hosted, C3 AI-Managed	C3 AI	Customer AWS GovCloud account	VPC provisioning, access, infrastructure costs	Advanced use cases or requirements

Comparison of Cloud Resources and Ownership

The C3 Agentic AI Platform integrates with core AWS GovCloud 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.

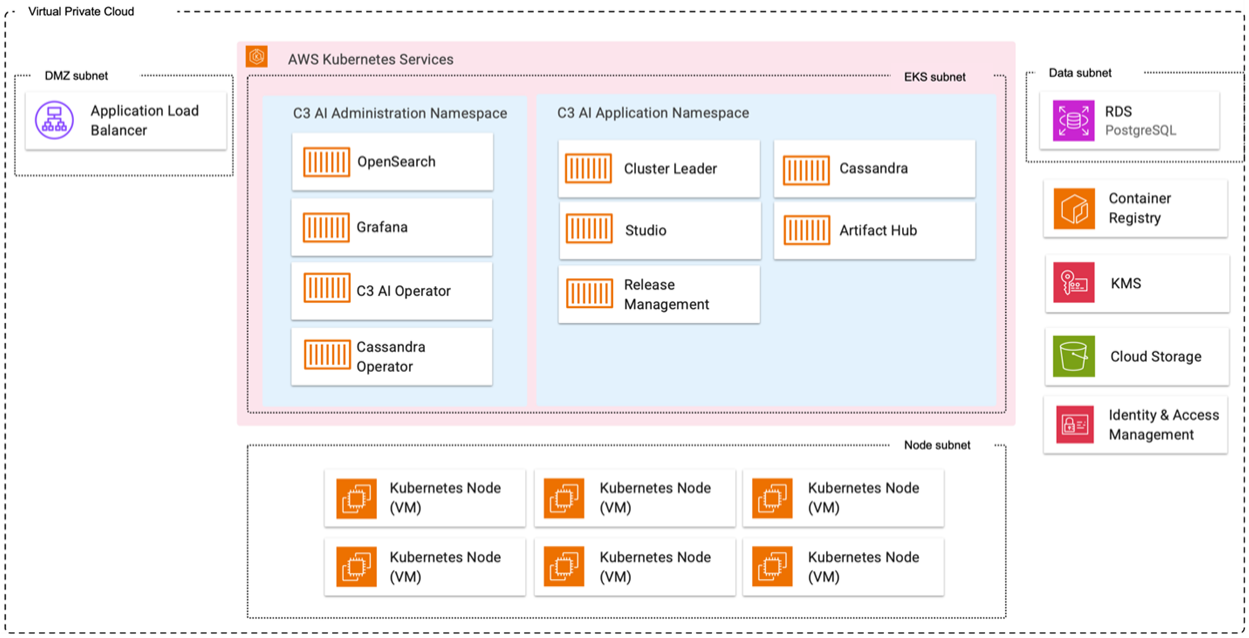

The C3 Agentic AI Platform requires specific AWS GovCloud services and infrastructure for successful deployment, as well as specific access requirements for C3 AI Operations to install, administer, and upgrade the C3 Agentic AI Platform and C3 AI Applications. A comparison of the different deployment options, dependent services, and responsibilities can be found below. The C3 AI Operational Controls Description found on c3.ai/legal are the definitive source of truth for responsibilities for all parties.

Name	Version	Purpose	SaaS/PaaS Subscription	Customer-Hosted, C3 AI-Managed
AWS GovCloud Elastic Kubernetes Service	1.34	Operating environment for deployment, scaling, and management of the C3 Agentic AI Platform.	Managed by SMX	Managed by C3 AI
AWS GovCloud Secrets Manager	-	Scalable, centralized cloud key management.	Managed by SMX	Managed by C3 AI
AWS GovCloud RDS (PostgreSQL)	15	Relational database service for internal operations of the C3 Agentic AI Platform.	Managed by SMX	Managed by C3 AI
AWS GovCloud S3	-	Object storage for application and platform configuration management.	Managed by SMX	Managed by C3 AI
AWS GovCloud IAM	-	Fine-grained access control for cloud service account resources.	Managed by SMX	Managed by C3 AI

AWS GovCloud Virtual Private Cloud	-	Dedicated, isolated network for inter-C3Cluster communication.	Managed by SMX	Provided by customer or managed by C3 AI
AWS GovCloud EC2	-	Compute instances required by Amazon EKS.	Managed by SMX	Managed by C3 AI
AWS Bedrock	-	AWS Bedrock model access for Generative AI applications.	Managed by SMX	Managed by C3 AI
AWS GovCloud KMS	-	Encryption key management for data at rest and in transit.	Managed by SMX	Managed by C3 AI
OIDC or SAML IdP	-	For secure authentication and authorization of end users and administrators	Provided by customer (FedRAMP Moderate) or managed by SMX (IL5)	Provided by customer or managed by C3 AI
TLS certificate and key	-	For secure web application access and SSL termination	Provided by customer	Provided by customer
Domain name / DNS A/AAAA record	-	For secure web application access and domain name resolution	Provided by customer	Provided by customer
OCI-compliant image registry	-	Image registry for C3 AI images. Any OCI-compliant solution is suitable, such as AWS GovCloud ECR, JFrog Artifactory, Sonatype Nexus, and more.	Managed by SMX	Provided by customer or managed by C3 AI

Pre-Installation Preparation

IMPORTANT: Common to both installation options, you must gather the following information. Failure to gather these in advance will cause delays during infrastructure deployment or C3 AI Platform configuration.

You must determine the following before starting:

Name	Description	Example
Cluster Name / ID	Identifier for the install which must follow a specific naming convention. See Appendix: Preferred Naming Convention. Dev/QA clusters: stgaws<customer abbreviation> Production clusters: prdaws<customer abbreviation>	stgawscust prdawscust
AWS Region	AWS region where the deployment will occur	us-gov-west-1 us-gov-east-1
AWS Account ID	The dedicated AWS account for the C3 AI deployment	123456789012
Domain Name / FQDN	The domain name where the platform installation will be accessible. Later in the installation process, C3 AI will provide the IP that the name should resolve to.	c3ai.yourcompany.com
TLS Certificate and Key for SSL Termination	TLS certificate and key matching domain name to be used for SSL termination. Does not need to be a wildcard and only needs to be for the FQDN determined previously.	tls.crt tls.key
Terraform PostgreSQL password	A strong, unique password in terraform.tfvars or via a secrets manager. Never deploy with the module default.	In accordance with organizational requirements

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

Name	Description	Example
IP address allowlist (if applicable)	If remote management from C3 AI VPN IPs is allowed, then reach out for C3 AI VPN IP addresses.	1.2.3.4/32 1.2.3.5/32
(Customer-Hosted) registry.c3.ai / jfrog.c3.ai Credentials	Username and password for access to C3 AI images, Helm charts, and Terraform models. Usage examples contained later.	

(Customer-Hosted) Image BOM, Helm Charts, Terraform Modules, and other installation resources	If the install is to be performed in a restricted environment or otherwise one where the customer is required to facilitate artifact transfers, C3 AI can provide installation artifacts. This includes: • List of C3 AI images to be used for installation • Helm charts (pull instructions or as tar.gz) • Terraform modules	registry.c3.ai/c3:8.11.1_103

Customer-Hosted, C3 AI-Managed Install Guide

Access

C3 AI personnel must be able to perform the infrastructure deployment and likewise require several IAM permissions for this to occur. There are two major options listed below and later explained in more detail.

Name	Pros	Cons
Administrator Permissions (preferred)	• Easiest & quickest to configure • Permissions scope can be limited via setting up dedicated AWS GovCloud account or member account	• Hard to limit permissions
C3 AI Infrastructure Scripts Contributor Role	• Limited in permissions scope; better for shared cloud accounts	• Requires customer engineer to perform part of infrastructure deployment • Requires customer engineer support for all C3 AI infrastructure upgrades

Option 1 – Dedicated AWS GovCloud Account and Administrator Permissions

For this option, the customer's responsibility is to:

1. (Optional) Create a new AWS GovCloud account or subaccount
2. Create IAM Users for C3 AI personnel
3. Assign AdministratorAccess AWS Policy to C3 AI IAM Users

Access Requirement	Description
AWS GovCloud Account or Member Account	Account must be an AWS GovCloud account in a suitable region. Account may be a member account of a pre-existing AWS GovCloud account.
AWS GovCloud IAM Users for C3 AI infrastructure administration personnel	C3 AI personnel must be granted managed policies. See next row. Verify all C3 AI operations personnel with IAM access are U.S. persons. Access may be managed via PAM solution. EC2 instance role is also suitable for this granted

	that C3 AI personnel have a method for accessing such an EC2 instance.
AWS GovCloud IAM Policies	Provide C3 AI Operations with access to AWS managed policies IAMReadOnlyAccess, IAMUserChangePassword, and AdministratorAccess.
(Optional) Secure access to the AWS GovCloud account	Provide C3 AI operations personnel with secure, remote access to a bastion host if required for administrative access. VPNs, government-furnished equipment/customer laptops, virtual desktops, and PAM solutions are all acceptable.
(Optional) Bastion host accessible by C3 AI infrastructure administration personnel	A bastion host with the following software is required: • RedHat Enterprise Linux 8 or greater, or any other modern Linux distribution • AWS CLI • Kubectl v1.34 • Helm v3 • HashiCorp Terraform v1.13 or greater • Python v3.12 • jq v1 • yq v4 • (Optional) Docker or Skopeo for image management Bastion host is required in configurations that utilized entirely private or limited networking or otherwise are unreachable from C3 AI personnel computers.

This is the most straightforward method of granting C3 AI personnel access to a customer's AWS GovCloud account. Customer may restrict access in any of the following ways, including:

- IAM user access is provided through a customer-managed PAM solution
- Administrator access is time-bound or controlled, i.e. after initial deployment, administrator access is replaced with an all-resources read-only
    - The higher the *persistent* access that C3 AI personnel have, the more self-sufficient C3 AI will be able to be in resolving and debugging infrastructure issues, performing upgrades, etc.
- Access is only allowed through customer VPN, EC2 instance role on bastion host, etc.

C3 AI should then have the requisite permissions and access to deploy the C3 AI Platform.

Option 2 – C3 AI Infrastructure Scripts Contributor Role

C3 AI's infrastructure Terraform scripts can also optionally create an infrastructure deployment AWS IAM Role and permit only specific trusted identifier ARNs to assume the infrastructure deployment IAM Role.

The customer's responsibilities for this option are to:

1. Work with C3 AI on downloading C3 AI's Terraform *bootstrap* module and *bootstrap* Terraform module configuration. Additional details are provided in the Appendix section.

2. On C3 AI's behalf, review the deployment plan and deploy the C3 AI *bootstrap* module.

Access Requirement	Description
AWS GovCloud Account or Member Account	Account must be an AWS GovCloud account in a suitable region. Account may be a member account of a pre-existing AWS GovCloud account.
AWS GovCloud IAM Users for C3 AI infrastructure administration personnel	C3 AI personnel must be granted deployed role. See next row. Verify all C3 AI operations personnel with IAM access are U.S. persons. Access may be managed via customer PAM solution given that C3 AI personnel are able to be granted accounts. EC2 instance role is also suitable for this granted that C3 AI personnel have a method for accessing such an EC2 instance.
AWS GovCloud IAM Roles	C3 AI infrastructure administration personnel must be granted the following IAM Role, created via C3 AI's bootstrap Terraform module: • <clustername>-c3cirole-01
(Optional) Secure access to the AWS GovCloud account	Provide C3 AI operations personnel with secure, remote access to a bastion host if required for administrative access. VPNs, government-furnished equipment/customer laptops, virtual desktops, and PAM solutions are all acceptable.
(Optional) Bastion host accessible by C3 AI infrastructure administration personnel	A bastion host with the following software is required: • RedHat Enterprise Linux 8 or greater, or any other modern Linux distribution • AWS CLI • Kubectl v1.34 • Helm v3 • HashiCorp Terraform v1.13 or greater • Python v3.12 • jq v1 • yq v4 • (Optional) Docker or Skopeo for image management Bastion host is required in configurations that utilized entirely private or limited networking or otherwise are unreachable from C3 AI personnel computers.

C3 AI should then have the requisite permissions and access to deploy the rest of the C3 AI Platform.

Infrastructure

C3 AI infrastructure is deployed via Infrastructure-as-Code scripts in the form of Terraform scripts. C3 AI Platform Infrastructure is largely split into two modules:

1. *bootstrap* (previously mentioned) deploys some IAM roles and policies required for C3 AI platform deployments.
2. *c3cluster* contains and deploys the C3 AI Platform computation, storage, and networking components, such as an EKS cluster, S3 buckets, KMS, additional IAM roles/policies for EKS service accounts, optionally a VPC and subnets, and more.

The modules will be configured and deployed by C3 AI personnel. Examples of how to download and use these modules are attached to the appendix.

Both modules support a limited number of customizations, such as existing VPC and subnet reuse. More details on networking customizations available are provided below.

NOTE: The Terraform module enables deletion protection by default for EKS clusters, RDS instances, S3 buckets, and KMS keys (enable_delete_protection = true). This prevents accidental destruction of production infrastructure. Set enable_delete_protection = false only for test/development environments.

Networking

Typically, networking configuration for C3 AI Platform infrastructure deployments fall into one of two categories, which are described below.

Name	Description	Examples
C3 AI Managed	No or basic customer networking requirements. Your organization has some basic requirements, but otherwise C3 AI's default networking deployment is suitable.	• EKS and/or S3 access only allowed via private networking • Ingress only allowed from specific IP CIDR blocks
Customer-Provided	Advanced customer networking requirements / customer provides VPC. Your organization has advanced networking requirements and desires that C3 AI deploy using a VPC, Subnets, and Security Groups that the customer provides.	• Hub-and-spoke cloud networking model • Highly custom ingress and egress paths (egress through firewall, advanced routing, etc.)

In either case, there are overarching networking requirements related to subnet configuration and CIDR block allocations.

VPC Requirements

Name	Description	Requirement	Example
VPC Region	AWS GovCould region where the deployment will occur.	Must be same as overall deployment region.	us-gov-west-1
VPC Sizing	CIDR blocks allocated to the VPC.	C3 AI Agentic AI Platform requires two (2) non-overlapping CIDR blocks, one private and routable, other two are private non-routable. See next two rows.	-

VPC Sizing: Private Routable	Private IPs routable to a public-facing Internet Gateway or equivalent	Contiguous /22 CIDR block of IPs	10.0.0.0/22
VPC Sizing: Private Non-routable (EKS Pods)	Non-routable private IPs, used by EKS pods	Contiguous /16 CIDR block of IPs	172.0.0.0/16
DNS	DNS resolution in VPC	VPC must have DNS hostnames and DNS resolution enabled	-
Subnets	Subnets created within the VPC	C3 AI requires the following subnets: 3x DMZ (one per AZ) 3x data (one per AZ) 3x EKS (one per AZ) 3x EKS Pod (one per AZ) See next four rows.	-
Subnets: EKS	Subnets for EKS nodes (EC2 instances) carved from primary CIDR block	/24 prefix for each of the subnets across AZs	10.0.0.0/24 10.0.1.0/24 10.0.2.0/24
Subnets: DMZ	Subnets for public load balancing / load balancer ENIs carved from primary CIDR block	/27 prefix for each of the subnets across AZs	10.0.3.0/27 10.0.3.32/27 10.0.3.64/27
Subnets: Data	Subnets for RDS PostgreSQL carved from primary CIDR block	/27 prefix for each of the subnets across AZs	10.0.3.96/27 10.0.3.128/27 10.0.3.160/27
Subnets: EKS pod	Subnets for EKS pods carved from *secondary* CIDR block	One /17 and two /18 prefixes across AZs	172.0.0.0/17 172.0.128.0/18 172.0.192.0/18
Subnet Route Table	The route table for the subnets.	The route table for workspace subnets must have 0.0.0.0/0 traffic that targets the appropriate network device (i.e. Internet Gateway, outbound firewall).	-
Security Groups		C3 AI must have at least one AWS Security Group and no more than five. You may reuse existing security groups rather than make new ones. See next four rows.	-
Security Groups: Endpoint Access	SGs for external endpoint or public resources.	See Appendix: Network Resources.	-
Security Groups: Egress	SGs for outbound access.	Egress is limited to those mentioned in Appendix: Network Resources and AWS endpoints.	-
Security Groups: Ingress	SGs for traffic from outside.	See Appendix: Network Resources.	-
Security Groups: Subnet-level ACLs	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 - ALLOW ALL to C3 AI VPC CIDR for internal traffic - ALLOW TCP 443 outbound to external endpoints mentioned in Appendix: Network Resources	

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 VPC uses the address space 10.0.0.0/22. For the Kubernetes pod and service networks, we use secondary CIDR ranges of 172.0.0.0/16 and 172.16.0.0/18, respectively. Note that the EKS service IP block, which hasn't been mentioned so far, is a block of virtual IPs and does not need to be assigned to the VPC and simply needs to be non-overlapping from the perspective of the EKS cluster.

For full C3 AI Platform functionality, access to several public resources or suitable customer-hosted equivalents are required. This includes resources like an OCI-compliant image registry for C3 AI images (previously mentioned), Python Package Index (PyPI) and Anaconda/Conda Forge for Python dependencies, and npm Registry for NodeJS dependencies. See the Appendix: Network Resources for a full list of public resources, suitable hosted alternatives, and descriptions.

AWS WAF Configuration

C3 provisions an AWS WAF by default on the public-facing Application Load Balancer (ALB). This WAF is enabled out of the box and governed by the following variables: managed_waf_rules, waf_rate_based_rules, waf_enforce_managed_rules, and waf_cloudwatch_logging_enabled.

Customers with an existing firewall or edge protection solution must disable the C3-provisioned WAF to avoid duplicate or conflicting controls. To disable the WAF, set the following in your c3cluster configuration: waf_enforce_managed_rules = false

When the WAF is disabled, responsibility for edge protection (rate limiting, managed rule enforcement, and request filtering) shifts to the customer's own firewall solution. Document this ownership in your shared-responsibility matrix so the accreditation boundary remains clear.

Notes for GovCloud deployments:

- Confirm that AWS WAF is an in-boundary, authorized service at your required FedRAMP/IL level before relying on the default configuration.
- Verify that the specific AWS managed rule groups referenced by managed_waf_rules are available in your target GovCloud region, as availability can differ from commercial regions.
- For fully private or air-gapped enclaves with no public ingress, the WAF is out of scope and should be disabled.

Networking Options

Option 1 – C3 AI-Managed

C3 AI provisions networking resources to the specifications declared above, very similar to the example values provided. Several configurations are also available, such as adjusting the subnet start address.

Option 2 – Customer-Provided

Your organization might have advanced requirements for networking. One common case is hub-and-spoke cloud networking models, where careful consideration must be made with respect to chosen routable CIDR block size, routable IPs, and ingress/egress path.

For these cases, C3 AI's infrastructure scripts support leveraging existing networking resources, such as customer-provided VPCs, subnets, and security groups, allowing a customer to create networking resources that conform to their own organizational requirements. Note that basic C3 AI networking infrastructure requirements must still be met, including outbound resource requirements (see Appendix: Network Resources) and those in the Networking section.

Some examples of deviations from requirements and subsequent trade-offs are listed below.

Name	Description	Trade-off
Smaller CIDR block for routable private IPs (/24)	Rather than the default /22, a smaller /24 block is used for routable private IPs on the VPC.	EKS cluster scale is limited but still suitable for small to medium sized deployments.
Subnets only in one AZ	Rather than deploy each subnet across all three AZs, only deploy subnets in one AZ.	More IPs are available in one AZ, but no failover in the event of AWS AZ outage.
Custom Security Groups	You want to customize security groups.	Acceptable if appropriate C3 AI services can still communicate to one-another (e.g. EKS pods to RDS).
Air-Gapped	EKS cluster is air-gapped, and no compute or application workloads can reach the Internet.	Acceptable as long as the EKS cluster can still reach required EKS endpoints.

C3 AI Platform Install via Helm/Kubernetes

While this guide is primarily focused on C3 AI Agentic Platform infrastructure deployment, some infrastructure options have downstream effects on how the Helm and Kubernetes deployments occur. Some general considerations are listed below. Please reference C3 AI Documentation: Non-Standard Deployment Overview for comprehensive deployment documentation.

Tools and Libraries

As mentioned from earlier sections, the following tools are generally required for a Helm installation of the platform:

- RedHat Enterprise Linux 8 or greater
- AWS CLI
- Kubectl v1.34

- Helm v3
- HashiCorp Terraform v1.13 or greater
- Python v3.12
- jq v1
- yq v4
- (Optional) Docker or Skopeo for image management

Air-Gapped Considerations

When deploying with egress limitations, air-gapped environments, or behind strict firewalls, be sure to go through the checklist below to avoid delays in deployment.

Name	**Description**	**Example**
Tools	Tools necessary for deployment are transferred to bastion host ahead of time	Helm, Kubectl, etc. are available on bastion host. See Tools and Libraries.
Images	Images from registry.c3.ai are made available in a suitable registry accessible from within the VPC ahead of time	ECR, other hosted image registry deployed in VPC
Helm Charts and Kubernetes Manifests	Other C3 AI Helm Charts and Kubernetes manifests necessary for the deployment are made available on the bastion host	Helm charts: such as c3crds, c3aiops, and c3, should be made available in tar.gz format. Manifests: secrets, other deployment configurations.
EKS/S3 Public Access Disabled Requires Bastion Host or Equivalent	If no public access is allowed for the EKS API endpoint, then a bastion host or equivalent will be required for EKS cluster management via CLI.	Amazon EKS Docs: Modifying cluster endpoint access

For private or air-gapped deployments where nodes cannot reach public endpoints, configure VPC interface endpoints using the vpc_endpoint_services variable. Supply the required service short-names (e.g., eks, ecr.api, ecr.dkr, sts, kms, logs, monitoring, secretsmanager, elasticloadbalancing, ec2, ssm).

Note that due to the number of possible configurations, this checklist may not be exhaustive. Please reach out to your C3 AI Center of Excellence if you have any questions.