C3 AI Installation Guide – Microsoft Azure

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)

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 Microsoft Azure. Customers may select their preferred Azure 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 Azure Virtual Network (VNet). 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:**
    - Grant C3 Operations owner access to a dedicated Azure subscription for C3-managed deployments, or alternatively, provision a secure Azure VNet using C3 AI's Terraform module with access granted to C3 Operations as described below. 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 Azure 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 Azure Cloud	Minimal	Most organizations
Customer-Hosted, C3 AI-Managed	C3 AI	Customer Azure VNet	Azure subscription, 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 (for example, customer-managed VNet provisioning), please reach out to your C3 AI representative.

Customer Hosted Install Requirements: Checklist

Summary

For C3 AI to operate in a customer-hosted Azure Cloud account, 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 infrastructure creation role ({cluster_name}-c3icrole-01) with temporary administrator privileges** to a dedicated Azure subscription so that C3 AI can perform tasks to set up your deployment.

   C3 AI requires administrator privileges to set up IAM policies, custom roles, and managed identities that allow C3 AI Operations to perform installation, setup, and deployment tasks. The Terraform bootstrap module creates this role automatically. 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. **Three managed identities for ongoing infrastructure management:**

Managed Identity	Purpose
{cluster_name}-c3-mi-01	C3 AI application workloads
{cluster_name}-c3privileged-mi-01	Privileged C3 AI application workloads
{cluster_name}-kube-mi-01	Kubernetes and kubelet operations

   These managed identities are created by the Terraform bootstrap module with the required Azure role assignments. See Appendix B: Terraform Module Reference for the full list of role assignments.

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 VNet and subnets with at least 1024 IP addresses (/22) available for C3 deployment.** The Terraform module creates the VNet and all required subnets automatically. AKS node pools are deployed across three availability zones by default.

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

6. **Public certificate with the complete chain and private key.** 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.

C3 AI Installation Requirements for Azure

The C3 Agentic AI Platform integrates with core Azure services like Azure VM, Azure Virtual Network (VNet), and IAM, enabling cohesive security and infrastructure management.

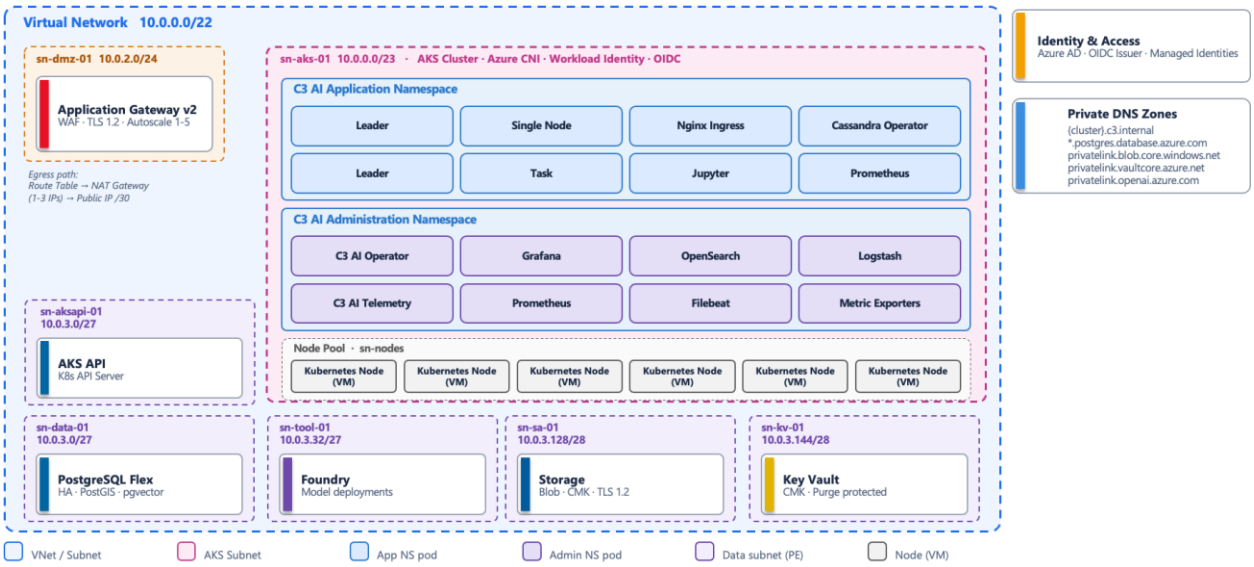

The C3 Agentic AI Platform requires specific Microsoft Azure cloud services and infrastructure for successful deployment, as well as specific access requirements for C3 AI Operations to install, administer, and upgrade the C3 Agentic AI Platform and C3 AI Applications.

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 Microsoft Azure cloud services

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

Microsoft Azure Service	Version	Description
Azure Kubernetes Service (AKS)	1.34	Operating environment responsible for the deployment, scaling, and management of the C3 Agentic AI Platform and C3 AI Applications.
Azure Key Vault	Current version	Securely store and manage sensitive information such as secrets, keys, and certificates.
Azure PostgreSQL (Flexible servers)	15	Relational data required for internal operations of the C3 Agentic AI Platform.

Azure Blob Store	Current version	Reliable and secure object storage used for the management of application and platform configuration and other ancillary tasks.
Azure Identity and Access Management (IAM)	Current version	Fine-grained access control and visibility for centrally managing cloud service account resources. The Terraform bootstrap module creates the required managed identities and custom roles automatically.
AzureRM Provider (Terraform)	4.57.0	HashiCorp Terraform provider for creating and managing Azure resources.
Azure Virtual Network (VNet)	Current version	Dedicated, isolated network for inter-cluster communication.
Azure Virtual Machines	Current version	Compute instances required by Azure Kubernetes Service.
Azure OpenAI Service	Current version	Advanced AI models developed by OpenAI, such as GPT-4, DALL-E, and Codex.

Microsoft Azure 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 Azure service principal 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
- Azure Lighthouse role assignments for designated C3 AI Operations personnel

- An additional **/28 CIDR IP block** for the C3 Software-Defined Perimeter/Management (SDM) peered VNet

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 Requirement	Description
A dedicated Azure subscription	When creating the project, C3 AI requires: (1) Subscription identifier, (2) Cloud region.
Secure internet access to the Azure cloud subscription	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 VNet; 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, Azure Command-Line Interface (CLI) (latest version), kubectl v1.34, Helm 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 Azure Container Registry, JFrog, and Anaconda Enterprise.
X.509 certificate for terminating network encryption	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) from the X.509 certificate 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.

Azure subscription

The Azure subscription must have the following Azure resource providers registered before deployment:

- az provider register --namespace Microsoft.Compute
- az provider register --namespace Microsoft.DBforPostgreSQL
- az provider register --namespace Microsoft.Network
- az provider register --namespace Microsoft.ManagedIdentity
- az provider register --namespace Microsoft.KeyVault
- az provider register --namespace Microsoft.Storage
- az provider register --namespace Microsoft.ContainerService

- az provider register --namespace Microsoft.Quota

Network configuration

To deploy the C3 Agentic AI Platform in your own Azure Virtual Network (VNet), C3 AI will create the VNet following the steps enumerated in the VNet requirements section below.

VNet requirements

The VNet must meet the following requirements to host a C3 AI cluster.

- Azure subscription
- VNet region
- VNet sizing
- VNet IP address ranges
- DNS
- Subnets
- Security groups
- Subnet-level network ACLs

The Azure subscription containing the C3 Agentic AI Platform must have end-to-end encryption enabled using encryption at host. Data stored on the host is encrypted at rest and flows encrypted to the Azure Storage service. Encryption at host must be enabled and can be accomplished using the Azure Portal or CLI. To enable encryption at host using the Azure CLI, run the following:

```
az feature register --name EncryptionAtHost --namespace Microsoft.Compute
```

The Terraform module enforces encryption at host for all AKS node pools (host_encryption_enabled = true by default).

VNet region

The Azure region where deployment will occur. Refer to Azure documentation for a list of available regions.

VNet sizing

Provision a separate VNet per cluster. Sharing a single VNet across multiple clusters in the same Azure subscription is not supported. Size your VNet and subnets to C3 AI specifications.

VNet IP address ranges

C3 AI does not impose strict limits on VNet netmasks. These are our recommendations.

For AKS deployment, C3 AI sets up two non-overlapping network address spaces with distinct purposes:

Primary VNet -- By default, the primary virtual network uses a /22 (we typically use 10.0.0.0/22, but any equivalent /22 range works).

Kubernetes pod and service networks -- In addition to the primary VNet, AKS requires separate private CIDR ranges that should not overlap with the primary VNet. The Terraform defaults are:

Network	Default CIDR	Size
Kubernetes pod network	172.18.0.0/15	131,072 IPs
Kubernetes service network	172.16.0.0/18	16,384 IPs
Docker bridge	172.17.0.0/16	65,536 IPs
DNS service IP	172.16.0.10	--

These secondary CIDR ranges are internal to Kubernetes and are not subnets carved out of the primary /22 VNet.

DNS

The VNet will have DNS hostnames and DNS resolution enabled.

Subnets

The Terraform module creates the following subnets within the VNet for each cluster:

Subnet	Default Name	Default CIDR	Prefix	Purpose
AKS	{cluster_name}-sn-aks-01	Configured via subnet_aks_prefix	/23	Primary AKS node subnet
DMZ	{cluster_name}-sn-dmz-01	Configured via subnet_dmz_prefix	/24	Public DMZ / Application Gateway
Data	{cluster_name}-sn-data-01	Configured via subnet_data_prefix	/27	PostgreSQL (private)
Tool	{cluster_name}-sn-tool-01	Configured via subnet_tool_prefix	/27	Tool services
Bastion	AzureBastionSubnet	Configured via subnet_bastion_prefix	/26	Azure Bastion
Storage Account	{cluster_name}-sn-sa-01	Configured via subnet_sa_prefix	/28	Storage account private endpoints
Key Vault	{cluster_name}-sn-kv-01	Configured via subnet_kv_prefix	/28	Key Vault private endpoints
AKS API	{cluster_name}-sn-aksapi-01	Configured via subnet_aksapi_prefix	/28	AKS API server VNet integration

NOTE: Public IP prefix required for the default Application Gateway. By default, the C3 AI Azure deployment provisions a public-facing Application Gateway in the DMZ subnet, which consumes a Public IP Prefix. In private or peered-VNet clusters where organizational policy does not permit a public IP prefix, the default deployment cannot proceed. In that case, either (1) obtain an exception to allow a public IP prefix for the Application Gateway, or (2) coordinate with the C3 AI Center of Excellence before deployment to use the

nginx-based private load balancer alternative. Raising this with C3 AI FieldOps during the VNet planning phase avoids a deployment-time block.

NOTE: Additional subnets might be required depending on whether it is a C3 AI-managed or customer-managed deployment. Contact the C3 AI CoE for more information.

Subnet route table

The route table ({cluster_name}-rt-01) 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 ({cluster_name}-natgw-01) and internet gateway.
- The NAT gateway is automatically created and associated with the AKS, Data, Tool, Storage Account, and Key Vault subnets.

Security groups

The Terraform module creates the following Network Security Groups (NSGs):

NSG	Default Name	Associated Subnet
DMZ	{cluster_name}-nsg-dmz-01	DMZ subnet
AKS	{cluster_name}-nsg-aks-01	AKS subnet
Data	{cluster_name}-nsg-data-01	Data subnet
Tool	{cluster_name}-nsg-tool-01	Tool subnet
Storage Account	{cluster_name}-nsg-sa-01	Storage Account subnet
Key Vault	{cluster_name}-nsg-kv-01	Key Vault subnet
AKS API	{cluster_name}-nsg-aksapi-01	AKS API subnet

NSG rules are configurable via Terraform variables (dmz_nsg_rules, aks_nsg_rules, data_nsg_rules, tool_nsg_rules, sa_nsg_rules, kv_nsg_rules, aksapi_nsg_rules). Default rules are applied at deployment time. Contact the C3 AI CoE for the required rule configuration for your deployment.

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
- 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 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 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 VNet 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 Azure subscription.

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 Azure account, you must have the following:

- An Azure subscription. Terraform provisions the resource group; you do not need to create one in advance.
- A user with the Owner role on the Azure subscription. This role is required for initial deployment.
- On your local development machine, you must have:
    - **The HashiCorp Terraform CLI** (>= 1.13.0). See Install Terraform on the Terraform website. Select AMD64 or ARM64 to match the client hardware from which you will run the Terraform scripts.
    - **The Azure CLI**, signed in through the az login command with a user that has Owner rights to your subscription to access Microsoft Azure Cloud. See How to install the Azure CLI and Sign in with Azure CLI for more information.
    - **The Kubernetes CLI** (kubectl). See the Kubernetes website for more information about kubectl and related commands for infrastructure creation and deployment.
    - **The Helm CLI**. See Installing Helm on the Helm website for more information.
- 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.

NOTE: As a security best practice, when authenticating with automated tools, systems, scripts, and apps, C3 AI recommends you sign in through the az login command with an Azure Active Directory (Azure AD) service principal. See Sign in with a service principal and Authenticating with Azure Service Principal for more information.

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 Azure Storage Account container with state locking enabled. 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 Microsoft Azure is a multi-step process due to limitations of HashiCorp Terraform and Azure-specific configuration requirements. The installation process is the following:

1. Create the Microsoft Azure Virtual Network (VNet) and required Azure services.
2. Validate the configuration of the VNet and required Azure services and provide C3 AI Operations access to the cluster.

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

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

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 VNet and required Azure services

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

1.1 Run the bootstrap module

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

Cluster Naming Rules:

The cluster name must adhere to the following restrictions:

- Must start with a lowercase letter
- Only lowercase letters and numbers are allowed (no hyphens, special characters, or diacritics)
- Must be at most 15 characters total
- For Dev and QA clusters: <stg><cloud><customerabbreviation> where <cloud> is az. For example, stgazcust
- For Production clusters: <prd><cloud><customerabbreviation> where <cloud> is az. For example, prdazcust

NOTE: The Terraform module enforces length <= 15 via validation. Ensure your cluster name also complies with the letter/number-only constraint.

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

```
module "bootstrap" {
  source                  = "<c3_url>/tf-registry__c3/azure/c3//modules/boots
trap"
  version                 = "VERSION_NUMBER"
  cluster_name            = "CLUSTER_NAME"
  region                  = "REGION"
  setup_role_principal_id = "IDENTITY_OBJECT_ID"
}

provider "azurerm" {
  features {}
  partner_id      = "AZURE_PARTNER_ID"
  tenant_id       = "AZURE_TENANT_ID"
  subscription_id = "AZURE_SUBSCRIPTION_ID"
}

terraform {
  required_version = ">= 1.13.0"
}
```

Replace the following:

Variable	Description
CLUSTER_NAME	The name of the C3 AI cluster. Must follow the cluster naming rules above.
VERSION_NUMBER	The version of the bootstrap module listed on the C3 AI BOM for the release version.
REGION	The region where the infrastructure will be deployed. See the Azure regions mapping list for more information.
IDENTITY_OBJECT_ID	The object identifier used to assume the infrastructure creation role ({cluster_name}-c3icrole-01). To get the Object ID, navigate to Azure Active Directory, search for the user, then select the user and get their Object ID.
AZURE_PARTNER_ID	The Microsoft partner ID. Contact the C3 AI CoE for more information.
AZURE_TENANT_ID	The Azure tenant where the C3 Agentic AI Platform will be deployed.
AZURE_SUBSCRIPTION_ID	The Azure subscription ID.

What the bootstrap module creates:

Resource	Default Name
Resource group	{cluster_name}-rsgp-c3-01
Infrastructure creation role	{cluster_name}-c3icrole-01
C3 application managed identity	{cluster_name}-c3-mi-01
C3 privileged managed identity	{cluster_name}-c3privileged-mi-01
Kubernetes managed identity	{cluster_name}-kube-mi-01
FlexPG custom role	{cluster_name} FlexPG
Application Gateway custom role	{cluster_name} Application Gateway

DiskEncryptionSet custom role	{cluster_name} DiskEncryptionSet
Purge Cognitive Services custom role	{cluster_name} Purge Cognitive Services

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 Azure -- Install CLI 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.

Be sure to login as the service principal that you specified as setup_role_principal_id in the bootstrap module.

Contact your account manager for C3 Control IPs (the list of IP addresses required by C3 AI). These IP addresses enable access for C3 AI Services to manage and maintain the C3 AI cluster. Replace the CIDR_TO_WHITELIST parameter below with the list of IP addresses.

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 VNet 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/azure/c3"
  version      = "VERSION_NUMBER"
  c3_region    = "REGION"
  cluster_name = "CLUSTER_NAME"

  ip_allowlist = [
    "CIDR_TO_WHITELIST",
  ]
```

```
  storage_cors_domains = ["http://*.DOMAIN_NAME"]

  pg_create_mode = "Default"  # Only use Default at creation time
}

provider "azurerm" {
  features {}
  partner_id      = "AZURE_PARTNER_ID"
  tenant_id       = "AZURE_TENANT_ID"
  subscription_id = "AZURE_SUBSCRIPTION_ID"
}

terraform {
  required_version = ">= 1.13.0"
}
```

Replace the following:

Variable	Description
CLUSTER_NAME	The name of the C3 AI cluster. Must follow the cluster naming rules above.
VERSION_NUMBER	The version of the c3cluster module listed on the C3 AI BOM for the release version.
REGION	The Azure region where the C3 Agentic AI Platform will be deployed.
CIDR_TO_WHITELIST	The list of required C3 AI IP addresses.
DOMAIN_NAME	The CORS domain for C3 AI Ex Machina file uploads.
AZURE_PARTNER_ID	The Microsoft partner ID. Contact the C3 AI CoE for more information.
AZURE_TENANT_ID	The Azure tenant where the C3 Agentic AI Platform will be deployed.
AZURE_SUBSCRIPTION_ID	The Azure subscription ID.

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 via storage_cors_domains 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 following 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

Step summary: Validate the configuration of the VNet and required Azure services and provide C3 AI Operations access to the cluster.

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

Contact the C3 AI 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 C3 AI Operations to deploy the C3 Agentic AI Platform.

If the C3 AI Cluster Validation Utility indicates the VNet 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.

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.

Item	Description
C3 AI Operations credentials	Credentials for C3 AI Operations team members.
Subscription ID	Subscription ID used in the execution of the Terraform scripts.
Tenant ID	Azure Tenant ID used in the execution of the Terraform scripts.
Resource Group Name	From Azure Portal, go to "Resource groups" and search your cluster name. By default, it should be {cluster_name}-rsgp-c3-01.
AKS cluster name	The name of the AKS cluster where the C3 Agentic AI Platform will be installed. By default, this will be {cluster_name}-kube-01; confirm this by going to "Kubernetes

	services" in the Azure Portal. The cluster name must adhere to the cluster naming rules described in section 1.1.
Region	The Azure region associated with the AKS cluster, from "Kubernetes services" -> the AKS cluster -> Overview -> Location. See the Azure regions mapping list for more information.
Azure SQL Postgres endpoint	From Azure Portal, go to "Azure Database for PostgreSQL servers" -> {cluster_name}-pg-shared -> Overview -> Server name.
Azure SQL Postgres credentials	Credentials required for the C3 Agentic AI Platform to connect to PostgreSQL. These can be set by pressing "Reset password" in the Azure Portal at "Azure Database for PostgreSQL servers" -> {cluster_name}-pg-shared.
C3 Managed Identity Client IDs	From Azure Portal, go to "Managed identities", filter to the resource group and click Apply. Select the managed identities {cluster_name}-c3-mi-01 and {cluster_name}-c3privileged-mi-01 and click into each. Share the "Client ID" of each managed identity. Map the {cluster_name}-c3privileged-mi-01 managed identity to the C3 AI Kubernetes service account c3-privileged. See Map Cloud Provider Identity to Kubernetes Service Account.
Storage Account Keys	From Azure Portal, go to "Storage accounts" -> filter to resource group {cluster_name}-rsgp-c3-01, and click the storage account. Select "Access Keys" and share key1.
Domain name	A fully qualified domain name for C3 Cluster ingress configuration (for example, c3project.customer.com).
Public and private key	The private and public key (and the certificate chain) from the X.509 certificate. This will be required for ingress configuration.

It is strongly recommended that the sharing of Postgres credentials and certificates occur using Azure Key Vault.

To grant C3 AI Operations AKS cluster administration permissions in the subscription, create a new role assignment for C3 AI Operations users, granting them the Azure Kubernetes Service Cluster Admin role:

```
az login
az role assignment create \
  --assignee <OBJECT_ID_OF_THE_USER> \
  --role "Azure Kubernetes Service Cluster Admin Role" \
  --scope /subscriptions/<SUBSCRIPTION_ID>/resourceGroups/<C3_CLUSTER_RESOURCE_GROUP>
```

Replace the following:

Variable	Description
<OBJECT_ID_OF_THE_USER>	The user, group, or service principal. Supported format: object ID, user sign-in name, or service principal name.
<SUBSCRIPTION_ID>	The subscription identifier.
<C3_CLUSTER_RESOURCE_GROUP>	The resource group associated with the cluster (default: {cluster_name}-rsgp-c3-01).

Appendix A: Default Resource Naming Convention

All resources follow the pattern {cluster_name}-{component}-{sequence}. Default names can be overridden via Terraform input variables.

Resource	Default Name Pattern
Resource Group	{cluster_name}-rsgp-c3-01
AKS Cluster	{cluster_name}-kube-01
AKS DNS Prefix	{cluster_name}-kube-01-dns
Virtual Network	{cluster_name}-vnet-01
AKS Subnet	{cluster_name}-sn-aks-01
DMZ Subnet	{cluster_name}-sn-dmz-01
Data Subnet	{cluster_name}-sn-data-01
Tool Subnet	{cluster_name}-sn-tool-01
Bastion Subnet	AzureBastionSubnet (Azure-required name)
Storage Account Subnet	{cluster_name}-sn-sa-01
Key Vault Subnet	{cluster_name}-sn-kv-01
AKS API Subnet	{cluster_name}-sn-aksapi-01
DMZ NSG	{cluster_name}-nsg-dmz-01
AKS NSG	{cluster_name}-nsg-aks-01
Data NSG	{cluster_name}-nsg-data-01
Tool NSG	{cluster_name}-nsg-tool-01
Storage Account NSG	{cluster_name}-nsg-sa-01
Key Vault NSG	{cluster_name}-nsg-kv-01
AKS API NSG	{cluster_name}-nsg-aksapi-01
Route Table	{cluster_name}-rt-01
NAT Gateway	{cluster_name}-natgw-01
NAT Gateway Public IP	{cluster_name}-natgw-pip-01
Public IP Prefix	{cluster_name}-pipfx-01
WAF Policy	{cluster_name}-waf-pol-01
Application Gateway	{cluster_name}-agw-kube-01
Application Gateway Public IP	{cluster_name}-agw-pip-01
Key Vault	{cluster_name}-keyvault
Disk Encryption Set	{cluster_name}-disk-encrypt-01
PostgreSQL Server	{cluster_name}-pg-shared
Storage Account	{cluster_name (hyphens removed)}store01
Cognitive Services Account	{cluster_name}-openai-01
C3 Managed Identity	{cluster_name}-c3-mi-01
C3 Privileged Managed Identity	{cluster_name}-c3privileged-mi-01
Kubernetes Managed Identity	{cluster_name}-kube-mi-01
Infrastructure Creation Role	{cluster_name}-c3icrole-01
FlexPG Custom Role	{cluster_name} FlexPG
Application Gateway Custom Role	{cluster_name} Application Gateway
DiskEncryptionSet Custom Role	{cluster_name} DiskEncryptionSet
Purge Cognitive Services Custom Role	{cluster_name} Purge Cognitive Services

Appendix B: Terraform Module Reference

The following Terraform modules are used to deploy the C3 Agentic AI Platform infrastructure. See the README.md file in each module directory for detailed input/output documentation.

Terraform Module	Description
bootstrap	Configures the necessary IAM roles, custom roles, managed identities, and policies to allow a Terraform orchestrator to deploy all services required for the C3 Agentic AI Platform on Azure.
aks-cluster	Configures Azure Kubernetes Service (AKS), including VNet configuration, endpoint access, authorized IP addresses, and the version of Kubernetes used by the cluster.
aks-nodepool	Configures the AKS node groups, including default instance size, required subnet, and permissions assigned to each node.
aks-sa	Configures the managed identity to be used by the C3 AI cluster via workload identity federation.
network	Configures the VNet, including public and private subnets, internet gateway, CIDR blocks, DHCP, NAT gateway, and Application Gateway.
firewall	Configures ingress and egress NSG security rules.
iam	Configures the required IAM roles, managed identities, and role assignments.
delegated-iam	Configures the IAM roles and role assignments for delegated access to the C3 AI cluster.
kms	Configures the Azure Key Vault and encryption keys for disk encryption and storage encryption.
postgres	Creates an Azure PostgreSQL Flexible Server database and assigns the database to the data subnet.
storage-account	Configures the Azure storage account(s) to be used with the C3 AI cluster.
resource-group	Manages the Azure resource group lifecycle.
ai-platform	Configures Azure Cognitive Services (AI Services) for AI model deployments including Azure OpenAI.
log-analytics	Configures the Azure Log Analytics workspace for monitoring PostgreSQL query insights and storage blob logging.
azure-policy	Configures Azure Policy assignments for resource delete protection and governance enforcement.

NOTE: The root module (referred to as c3cluster in deployment steps) coordinates the execution of all sub-modules listed above. It is the top-level entry point, not a separate sub-module.