

Nutanix Government Cloud Clusters - Nutanix Government Cloud Clusters on AWS

[PDF generated February 25 2026. For all recent updates please see the Nutanix Bible releases notes located at https://nutanixbible.com/release_notes.html. Disclaimer: Downloaded PDFs may not always contain the latest information.]

The Nutanix Government Cloud Clusters (GC2) solution meets the stringent security and compliance requirements of government and highly regulated enterprises deploying hybrid cloud infrastructure on bare-metal Amazon Web Services (AWS). Unlike the traditional Nutanix Cloud Clusters (NC2) architecture, which requires you to add AWS credentials and maintain continuous outbound connectivity to Nutanix software as a service (SaaS) endpoints, GC2 places complete control within the AWS environment. You retain full ownership of your AWS credentials, eliminate external dependencies, and operate your Nutanix clusters entirely within your own virtual private cloud (VPC), addressing the core security and sovereignty concerns that prevent adoption in government cloud environments.

Supported Configurations

The solution is applicable to the configurations below (list may be incomplete, refer to documentation for a fully supported list):

Core Use Case(s):

- On-Demand / burst capacity
- Backup / DR
- Cloud Native
- Geo Expansion / DC consolidation
- App migration

Management interfaces(s):

- Prism Central (PC) - Nutanix Management
- AWS Console - AWS Management

• EC2 Metal Instance Types:

- m5d.metal
- z1d.metal
- i3en.metal
- i4i.metal
- m6d.metal
- i7i.metal-48xl
- i7ie.metal-48xl

• Supported Regions:

- GovCloud US-East
- GovCloud US-West

Upgrades:

- Part of AOS

Compatible Features:

- AOS Features
- AWS Services

Key terms / Constructs

The following key items are used throughout this section and defined in the following:

- Region
 - A geographic landmass or area where multiple Availability Zones (sites) are located. A region can have two or more AZs. These can include regions like US-East-1 or US-West-1.
- Availability Zone (AZ)
 - An AZ consists of one or more discrete datacenters inter-connected by low latency links. Each site has its own redundant power, cooling, network, etc. Comparing these to a traditional colo or datacenter, these would be considered more resilient as a AZ can consist of multiple independent datacenters. These can include sites like US-East-1a or US-West-1a.
- VPC
 - A logically isolated segment of the AWS cloud for tenants. Provides a mechanism to secure and isolate environment from others. Can be exposed to the internet or other private network segments (other VPCs, or VPNs).
- S3
 - Amazon's object service which provides persistent object storage accessed via the S3 API.
- Cloud Formation Template (CFT)
 - A Cloud Formation Template simplifies provisioning, but allowing you to define a "stack" of resources and dependencies. This stack can then be provisioned as a whole instead of each individual resource.

Cluster Architecture

The GC2 architecture shifts orchestration intelligence from the NC2 console directly into Prism Element running in your cluster. Cluster provisioning uses AWS CloudFormation templates that use precreated network configurations, enabling flexible deployment into government cloud regions while maintaining strict network isolation. Day 2 operations, including cluster expansion, node replacement, and upgrades, run entirely within your environment through Prism Element, ensuring that sensitive workloads never expose credentials or operational telemetry beyond your boundaries.

In the standard Nutanix Cloud Clusters (NC2) model, you extend permissions to the NC2 SaaS service to provision clusters on your behalf, and it provisions all the necessary cloud infrastructure. After you deploy the infrastructure, Nutanix deploys agents to the Controller VMs (CVMs) to synchronize with the NC2 console through outbound internet connectivity. This architecture relies on three critical agents working together:

Host agent

The host agent continuously gathers telemetry on hardware health, network status, storage conditions, and system performance, emitting regular heartbeats back to the NC2 console for centralized monitoring.

Clusters agent

The clusters agent is a conduit between the NC2 console and local infrastructure, listening for orchestration intents issued by the NC2 console, verifying their authenticity, and forwarding them to the appropriate local services for implementation.

Infrastructure gateway

The infrastructure gateway translates these high-level intents into actionable AWS API calls to provision Elastic Compute Cloud (EC2) instances, configure networking, manage storage attachments, hibernate nodes, and perform cluster scaling operations.

While this architecture works effectively for most enterprise deployments, it presents challenges for government cloud environments. Extending AWS permissions to the NC2 console to act on behalf of the customer and maintaining persistent outbound connectivity to external endpoints violates fundamental security principles mandated by government compliance frameworks that demand complete data sovereignty and zero-trust network architectures. Highly regulated industries can't permit their infrastructure credentials to leave their controlled environments, nor can they allow continuous telemetry streams to external commercial software as a service (SaaS) platforms. These security constraints led Nutanix to develop Government Cloud Clusters (GC2), where all orchestration intelligence moves from the NC2 console to Prism Element, running entirely within your AWS virtual private cloud (VPC) to eliminate credential sharing and external connectivity requirements while preserving full operational capabilities.

The new infrastructure manager in the GC2 service assumes the role previously held by the NC2 console as the intent generator, while the infrastructure gateway continues as the engine. By colocating both services in the private environment, GC2 preserves all the operational capabilities of NC2 while delivering a secure and autonomous experience that meets government cloud requirements.

The infrastructure manager uses a leader-based design with Zookeeper maintaining leadership across CVMs, ensuring high availability and fault tolerance. Every CVM runs an infrastructure manager instance, but only one node actively processes user inputs and generates intents at any given time. If the leader node fails, the remaining infrastructure manager instances automatically elect a new leader and resume operations without manual intervention, guaranteeing continuous cluster management capabilities even during node failures.

The primary operational relationship in GC2 exists between the infrastructure manager and infrastructure gateway. When you initiate an operation through the command-line interface (CLI), the infrastructure manager validates the input, performs preflight checks such as instance type compatibility, constructs a standardized intent, and invokes infrastructure gateway APIs to perform necessary Prism Element cluster operations. The infrastructure manager also directly interacts with AWS APIs to manage cloud infrastructure tasks like provisioning Elastic Compute Cloud (EC2) instances, eliminating the need to share credentials with external platforms.

The following shows a high-level overview of the GC2 Communication in an AWS VPC :

GC2 - Overview

GC2 Agents - Overview

This clear separation of responsibilities creates a streamlined and resilient pipeline from user request to cluster state change. The infrastructure manager handles orchestration logic and cloud resource management, while the infrastructure gateway focuses exclusively on implementing cluster-level changes in Prism Element. Zookeeper-backed state persistence provides fault tolerance and automatic recovery in the event of failures, maintaining operational continuity without requiring external intervention or connectivity.

AHV Networking with Amazon Web Services

Nutanix integration with the Amazon Web Services (AWS) networking stack means that every VM deployed on Government Cloud Clusters (GC2) on AWS receives a native AWS IP address when using native networking, so that applications have full access to all AWS resources as soon as you migrate or create them on NC2 on AWS. Network performance remains high because the Nutanix network capabilities are directly on top of the AWS overlay, and resource consumption is low because you don't need additional network components.

AHV uses Open vSwitch (OVS) for all VM networking. You can configure VM networking through Prism or the AOS command-line interface (aCLI), and each virtual network interface card (vNIC) connects to a tap interface. Native networking uses the same networking stack as on-premises. The AHV host, VMs, and physical interfaces use ports to connect to the bridges, and both bridges communicate with the AWS overlay network. Each host has the required drivers to use the AWS overlay network.

GC2 - OVS Architecture

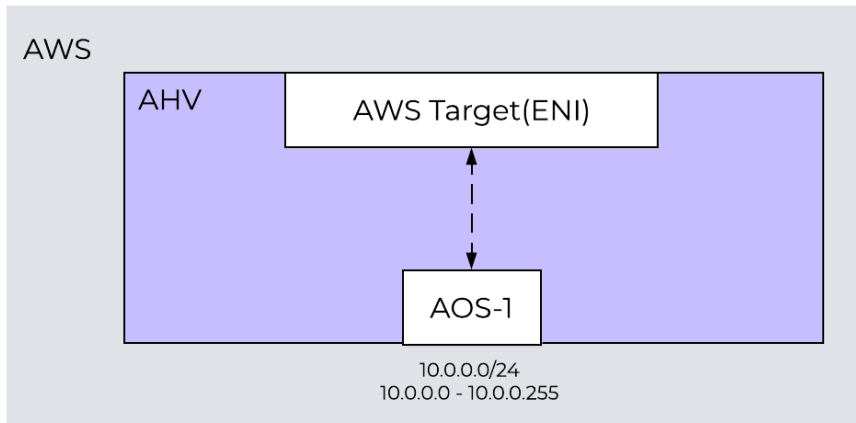
GC2 on AWS - OVS Architecture

Nutanix uses native AWS API calls to deploy AOS on bare-metal Elastic Compute Cloud (EC2) instances and consume network resources. Each bare-metal EC2 instance has full access to its bandwidth through an elastic network interface (ENI), so if you deploy Nutanix to an `i7i.metal` instance, each node has access to up to 100 Gbps. An ENI is a logical networking component in a VPC that represents a virtual network card, and each ENI can have one primary IP address and up to 49 secondary IP addresses. AWS hosts can support up to 15 ENIs.

When deployed with GC2 on AWS, AHV runs the Cloud Network Controller service on each node as a leaderless service and runs an OpenFlow controller. Cloud Network Controller uses an internal service called Cloud Port Manager to create and delete ENIs and assign ENI IP addresses to guest VMs.

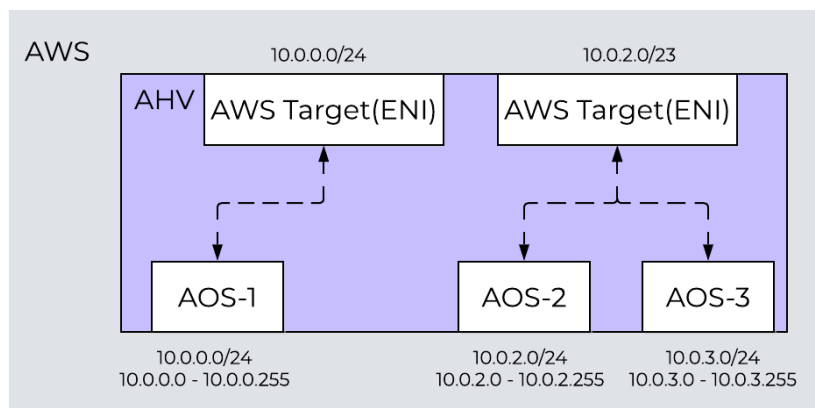
Cloud Port Manager can map large Classless Inter-Domain Routing (CIDR) ranges from AWS and allows AHV to consume all or a subset of the range. If you use an AWS subnet of `10.0.0.0/24` and then create an AHV subnet of `10.0.0.0/24`, Cloud Port Manager uses 1 ENI (cloud port) until active VMs consume all of the secondary IP addresses. When the 49 secondary IP addresses are

used, Cloud Port Manager attaches an additional ENI to the host and repeats the process. Because each new subnet uses a different ENI on the host, this process can lead to ENI exhaustion if you use many AWS subnets for your deployment.



One-to-One Nutanix AHV and AWS Subnet Mapping

To prevent ENI exhaustion, use an AWS subnet of 10.0.2.0/23 as your AWS target. In AHV, create two subnets of 10.0.2.0/24 and 10.0.3.0/24. With this configuration, Cloud Port Manager maps the AHV subnets to one ENI, and you can use many subnets without exhausting the bare-metal nodes' ENIs. Because a GC2 cluster can have multiple AWS subnets if enough ENIs are available, you can dedicate ENIs to any subnet for more throughput.



One-to-One and One-to-Many Nutanix AHV and AWS Subnet Mapping

When you consume multiple AHV subnets in a large AWS CIDR range, the network controller generates an Address Resolution Protocol (ARP) request for the AWS default gateway with the ENI address as the source. After the AWS default gateway responds, the network controller installs ARP proxy flows for all AHV subnets with active VMs on the ENI (cloud port). The ARP requests to a network's default gateway reach the proxy flow and receive the media access control (MAC) address of the AWS gateway in response. This configuration allows traffic to enter and exit the cluster, and it configures OVS flow rules to ensure traffic enters and exits on the correct ENI. For more information on the Nutanix implementation of OVS, see the AHV Administration Guide.

GC2 on AWS creates a single default security group for guest VMs running in the Nutanix cluster. Any ENIs created to support guest VMs are members of this default security group, which allows all guest VMs in a cluster to communicate with each other.

From the AWS VPC dashboard, click on 'subnets' then click on 'Create Subnet' and input the network details:

Creating a Subnet

An AWS Region is a distinct geographic area. Each Region has multiple, isolated locations known as Availability Zones (AZs), which are logical datacenters available for any AWS customer in that Region to use. Each AZ in a Region has redundant and separate power, networking, and connectivity to reduce the likelihood of two AZs failing simultaneously.

Create a subnet in AWS in a VPC, then connect it to AOS in Prism Element. Cloud network, a new service in the CVM, works with AOS configuration and assigns a VLAN ID (or VLAN tag) to the AWS subnet and fetches relevant details about the subnet from AWS. The network service prevents you from using the AHV or CVM subnet for guest VMs by not allowing you to create a network with the same subnet.

You can use each ENI to manage 49 secondary IP addresses. A new ENI is also created for each subnet that you use. The AHV host, VMs, and physical interfaces use ports to connect to the bridges, and both bridges communicate with the AWS overlay network. Because each host already has the drivers that it needs for a successful deployment, you don't need to do any additional work to use the AWS overlay network.

Always keep the following best practices in mind:

- Don't share AWS guest-VM subnets between clusters.
- Have separate subnets for management (AHV and CVM) and guest VMs.
- If you plan to use VPC peering, use nondefault subnets to ensure uniqueness across AWS Regions.
- Divide your VPC network range evenly across all usable AZs in a Region.
- In each AZ, create one subnet for each group of hosts that has unique routing requirements (for example, public versus private routing).
- Size your VPC CIDR and subnets to support significant growth.

Guest AHV IP Address Management

AHV uses IP address management (IPAM) to integrate with native AWS networking. GC2 on AWS uses the native AHV IPAM to inform the AWS DHCP server of all IP address assignments using API calls. NC2 relies on AWS to send gratuitous Address Resolution Protocol (ARP) packets for any additions to an ENI's secondary IP addresses. We rely on these packets to ensure that each hypervisor host is notified when an IP address moves or new IP addresses become reachable. For guest VMs, you can't share a given AWS subnet between two GC2 on AWS deployments. You can, however, use the same management subnet (AHV and CVMs) for multiple clusters.

The cloud network controller, a service in the AHV host, helps with ENI creation. The cloud network controller runs an OpenFlow controller, which manages the OVS in the AHV hosts and handles mapping, unmapping, and migrating guest-VM secondary IP addresses between ENIs or hosts. A subcomponent of the cloud network controller called cloud port manager provides the interface and manages AWS ENIs.

IPAM avoids address overlap by sending AWS API calls to inform AWS which addresses are being used.

The AOS leader assigns an IP address from the address pool when it creates a managed vNIC, and it releases the address back to the pool when the vNIC or VM is deleted.