



Cisco Cloud Network Controller Release Notes, Release 26.0(1)

Introduction

If you have a private cloud, you might run part of your workload on a public cloud. However, migrating workload to the public cloud requires working with a different cloud provider interface and learning different ways to set up connectivity and define security policies. Meeting these challenges can result in increased operational cost and loss of consistency. Cisco Cloud Network Controller can be used to solve these problems by extending a Cisco Multi-Site fabric to Amazon Web Services (AWS), Microsoft Azure, or Google public clouds. You can also mix AWS, Azure, and Google Cloud in your deployment.

This document describes the features, issues, and limitations for the Cisco Cloud Network Controller software. For the features, issues, and limitations for the Cisco APIC, see the appropriate [Cisco Application Policy Infrastructure Controller Release Notes](#). For the features, issues, and limitations for the Cisco Multi-Site Orchestrator/Nexus Dashboard Orchestrator, see the appropriate [Cisco Nexus Dashboard Orchestrator Release Notes](#).

For more information about this product, see "Related Content."

Note: The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Date	Description
March 31, 2023	Release 26.0(1c) became available.

New Software Features

Product Impact	Feature	Description
Performance and Scalability	Scale enhancements	Several scale enhancements are now available as part of Release 26.0(1). For more information, see the Cisco Cloud Network Controller Verified Scalability Guide, Release 26.0(1) .
	Enhancements to the Cisco Cloud Network Controller GUI	The screenshots and diagrams in this document have been updated to reflect the recent UI changes. While the look and feel of the new UI is different, the arrangement of the screens and configuration options remains the same.

Supported Upgrade Paths

Cisco Cloud Network Controller supports policy-based upgrades for the following upgrade paths:

- Release 25.0(4) to 26.0(1)
- Release 25.0(5) to 26.0(1)

- Release 25.1(1) to 26.0(1)

Changes in Behavior

There are no changes in behavior in this release.

Open Issues

Click the bug ID to access the Bug Search tool and see additional information about the bug. The "Exists In" column of the table specifies the 26.0(1) releases in which the bug exists. A bug might also exist in releases other than the 26.0(1) releases.

Bug ID	Description	Exists in
CSCWe26882	After a Cisco Cloud Network Controller upgrade, there is a banner on Cisco Cloud Network Controller Dashboard which says that the system is still in progress. The banner usually disappears when the Cisco Cloud Network Controller shows that all the Vnet peering is OK. However, the banner could disappear before the Cisco Cloud Network Controller shows all the VNet Peering are OK. There is no functional impact but can cause confusion to the users.	26.0(1c) and later
CSCWe68147	After the L4-L7 configuration is posted, the connector crashes on adding more than 25 firewall VMs on Azure in one request.	26.0(1c) and later
CSCWe73030	User is unable to log into the Visore tool after the session times out. The browser continues to show the login page.	26.0(1c) and later
CSCWe75135	Cisco Cloud Network Controller dashboard shows Cisco Catalyst 8000V routers in error/down state. This issue is seen in rare scenarios when new Cisco Catalyst 8000V routers are deployed.	26.0(1c) and later
CSCWe75281	In some cases when IAM roles are not configured before restoring a huge configuration on AWS Cisco Cloud Network Controller, the system may not deploy the resources in the cloud. Cisco Cloud Network Controller will show faults indicating permission issues.	26.0(1c) and later
CSCWe80768	Traffic that is destined to another tenant in a same or different site and routed via Cisco Catalyst 8000V may be dropped on the source side Cisco Catalyst 8000V.	26.0(1c) and later
CSCWd03385	The "Cloud Access Privilege" for certain brownfield Cloud Context Profiles in the tabular view under "Application Management" may sometime show as "Not Applicable"	25.1(1e) and later
CSCWd10789	When we scale up CCRs or manage the CCRs in a region within 1 min after CCR scale down or unmanaging CCRs in a region, the new CCRs will not be programmed with licenses. A fault on the UI will indicate that the CCR is not programmed with licenses and the hcplatformlicense operational state is down.	25.1(1e) and later
CSCWd10816	Flow filter options may still be available in the statistics chart dropdown after the filter is deleted.	25.1(1e) and later
CSCWc71587	In some cases, the BGP session between the on-premises router and the Cisco Catalyst 8000V router in the cloud is down. Although the crypto sessions and tunnels are up, there could be some traffic loss when this happens.	25.0(5k) and later
CSCWc83944	You might not be able to SSH to the Cisco Cloud Network Controller on Azure after launching it for the first time during first 5-10 minutes of system bringup.	25.0(5k) and later

Bug ID	Description	Exists in
CSCwa97199	No functional issue, stale licenses are configured on the Cisco Catalyst 8000V. On programming a T2 license on the Cisco Catalyst 8000V, a stale license entry is seen on both the Cisco Catalyst 8000V and the Cisco Smart account.	25.0(4k) and later
CSCwb01378	The following fault appears on the Cisco Cloud APIC indicating that the license is not configured on the CCR: Oper State of HcplatformLicense is down with administrative-down.	25.0(4k) and later
CSCwc01909	When staying on the VRFs table, the VRF internal/external statuses sometimes do not update for a long time.	25.0(4k) and later
CSCwc28787	When BFD is enabled between sites, the BFD sessions go down and come back up quickly with no user intervention/trigger.	25.0(4k) and later
CSCwc39392	In a few specific cases, on importing a brownfield VNet with a Routing & Security access policy, an NSG might not get created and/or VNets might not peer.	25.0(4k) and later
CSCvu72354	Adding an EPG endpoint selector fails with an error message saying the selector is already attached.	25.0(1c) and later
CSCvx16601	When the "AllowAll" flag is enabled on a service device such as a native load balancer or on the logical interface of a third party device, it is possible that to see some specific rules apart from a rule that allows all traffic from any source to any destination.	25.0(1c) and later
CSCvy89617	Cloud routers may not get created if external network objects are not configured. External network configuration is required for configuring cloud routers.	25.0(1c) and later
CSCvy97972	Cisco Cloud APIC in this release limits the number of regions where we can deploy the hubnetwork in order to establish external connectivity. When you attempt to deploy/configure hubnetwork in more than four regions, the configuration will be rejected with the following error: Invalid Configuration CT_INTNETWORK_REGION_MAXIMUM: At present, there can be at most 4 cloudRegionName in cloudtemplateIntNetwork uni/tn-infra/infranetwork-default/intnetwork-default; current count = <total-hubnetwork-regions-attempted>	25.0(1c) and later

Resolved Issues

Click the bug ID to access the Bug Search tool and see additional information about the bug. The "Fixed In" column of the table specifies whether the bug was resolved in the base release or a patch release.

Bug ID	Description	Fixed in
CSCvz17160	Customers are restricted to shorter key value pairs than they need to be.	26.0(1c)
CSCvz38067	Incorrect DNS server is configured on Cisco Cloud APIC with Google Cloud. Though this is not directly used when deploying Cisco Cloud APIC with Google Cloud, an incorrect IP address is configured.	26.0(1c)
CSCvz52773	When performing a Cisco Cloud APIC upgrade (but not also performing a CSR upgrade), before the upgrade is finished and when the Cisco Cloud APIC is reconciling the CSR configurations, if you delete certain configurations and add the same configurations back (for example, if you delete a VRF and add the VRF back), a traffic drop may happen. Eventually it should recover.	26.0(1c)

Bug ID	Description	Fixed in
CSCwa03277	<p>When the brownfield VPC is imported into Cisco Cloud APIC, you need to take care of the creation and management of the route table and route table entries, security group rules, and transit gateway VPC attachment.</p> <p>After creating the transit gateway VPC attachment with the infra transit gateway for the brownfield VPC in the AWS console, the corresponding cloudCtxPeerOper for the brownfield cloud context profile will move from Failed state to Configured state.</p> <p>After that, if the created transit gateway attachment for the brownfield VPC is deleted in the AWS console, cloudCtxPeerOper is not moving back to Failed state.</p>	26.0(1c)
CSCwa43845	<p>This issue occurs if there is a misconfiguration done where the local subnet was provided under routes leaked from external VRF to internal VRF.</p> <p>This is an Azure Cloud APIC only issue, since AWS does not allow programming routes that overlap with VPC's CIDR.</p>	26.0(1c)
CSCwa44822	<p>When the tunnels are created from source interface Gig 2 or 4, we put allow all for security group. In future the allow all rule needs to be removed and provide explicit IP needs to be allowed.</p>	26.0(1c)
CSCwa48929	<p>If Cloud APIC is rebooted, in a rare case an expected rule may not be programmed on the cloud due to a timing mismatch in the reconcile and programming workflow.</p>	26.0(1c)
CSCwa92033	<p>Cisco Catalyst 8000V License is not getting deregistered from Cisco Smart License Server as the Cisco Catalyst 8000V lose connection to the internet before getting deleted.</p>	26.0(1c)
CSCwa99935	<p>Disabling "Hub peering" triggers a CSR deployment, which restricts the CSR count and increases the bandwidth together in a single action.</p>	26.0(1c)
CSCwc66761	<p>When we move from private peering to public internet, stale rules are present on GCP. These are specific IP rules.</p> <p>When we move from private peering to public internet, NDO does not delete the cloudtemplateNextHopIp MOs (which represents the onprem tep pool). In the case of AWS and Azure, the Cloud Network Controller detects if private peering is static (no-IPsec) or dynamic (IPsec enabled) and sets the route reachability on BgpPeerP as Static or Dynamic. Today we program static routes in Cisco Cloud Routers and create firewall rules in AWS and Azure only if the route reachability is Static. We delete static routes in Cisco Cloud Routers and delete firewall rules if the route reachability changes to Dynamic.</p> <p>This behavior was causing issues in GCP since we always need to program firewall rules for GCP, for private peering. Since cloudtemplateNextHopIp is never deleted, we cannot depend on the existence or deletion of this MO to program or cleanup firewall rules on GCP. We need to introduce a new type for GCP private peering so that we can still program firewall rules but not program static routes in Cisco Cloud Routers if IPsec is enabled. Currently in GCP, the firewall rules depend on the existence or deletion of cloudtemplateNextHopIp MOs. As they are never deleted by NDO, rules are never deleted on the Cloud Network Controller.</p>	26.0(1c)

Bug ID	Description	Fixed in
CSCwc76230	<p>On AWS Cloud Network Controller with transit gateway Connect enabled, under any of the following conditions -</p> <ol style="list-style-type: none"> 1. If the Infra AWS account is not cleaned before deploying the Cloud Network Controller and if there are any stale transit gateway Connect peers on the stale transit gateway Connect that is not cleaned 2. If the new transit gateway Connect peers are assigned the same inner IP as that of the stale transit gateway Connect peers <p>Then the new transit gateway Connect peers will fail until the stale transit gateway Connect peers are deleted. When this happens, the Cloud Network Controller should raise a fault, but currently the fault is not raised by the Cloud Network Controller.</p>	26.0(1c)
CSCwc77209	On the Azure Cloud Network Controller user interface, the BGP and OSPF sessions are seen as "down" after upgrading to the 25.0(5k) release version. On a fresh deployment, this issue will not be seen.	26.0(1c)
CSCwd17270	Existing Traffic Blackhole when Catalyst 8000Vs are added to a new region.	26.0(1c)
CSCwd17434	In Scale setup (with ~300 Spoke VNETs), deploying the CCR on non-home region took more than 4 hours. The delay is in Azure for creating ~300 VNet peerings and subnets. In Azure activity log of the resource group, we can see there are subnet creation retry-able failures, those failures are due to the previous subnet creation not finished, those failures are ok and retries will eventually succeed.	26.0(1c)

Known Issues

Click the bug ID to access the Bug Search tool and see additional information about the bug. The "Exists In" column of the table specifies the 26.0(1) releases in which the bug exists. A bug might also exist in releases other than the 26.0(1) releases.

Bug ID	Description	Exists in
CSCvo06626	When a cloudExtEpg matches on a 0/0 network and has a bi-directional contract with two cloud EPGs, such as cloudEpg1 and CloudEpg2, this can result in inadvertent communication between endpoints in cloudEpg1 and cloudEpg2 without a contract between the two EPGs themselves.	25.0(1c) and later
CSCvo30542	TACACS monitoring of the destination group is not supported through the GUI.	25.0(1c) and later
CSCvo55112	Logs are lost upon stopping the Cloud APIC instance.	25.0(1c) and later
CSCvo95998	There is traffic loss after a Cloud APIC upgrade. Traffic will eventually converge, but this could take a few minutes.	25.0(1c) and later
CSCvq11780	Creating VPN connections fail with the "invalidCidr" error in AWS or the "More than one connection having the same BGP setting is not allowed" error in Azure.	25.0(1c) and later
CSCvq76039	When a fault is raised in the Cloud APIC, the fault message will be truncated and will not include the entire cloud message description.	25.0(1c) and later
CSCvr01341	REST API access to the Cloud APIC becomes delayed after deleting a tenant with scaled EPGs and endpoints. The client needs to retry after receiving the error.	25.0(1c) and later

Bug ID	Description	Exists in
CSCvu05329	The Ctx Oper managed object is not deleted after the attachment is deleted.	25.0(1c) and later
CSCvu64277	Stats seen on Cisco Cloud APIC are sometimes not in sync with Azure stats.	25.0(1c) and later
CSCvu78074	Route nextHop is not set to the redirect service node specified in the service graph.	25.0(1c) and later
CSCvu81355	Traffic gets dropped after downgrading to the 5.0(1) release. Cloud Services Router has incompatible configurations due to an issue with reading configurations using SSH.	25.0(1c) and later
CSCvu88006	On the Dashboard, fewer VNet peerings are shown than expected.	25.0(1c) and later
CSCvw32664	When the CSR bandwidth needs to be increased, the user needs to undeploy all the CSRs in all the regions and redeploy with the desired bandwidth, which can cause traffic loss.	25.0(1c) and later
CSCvw81647	When an invalid Cloud Services Router license token is configured after initially configuring a valid token, the Cloud Services Router fails the license registration and keeps using the old valid token. This failure can only be found from the CSR event log.	25.0(1c) and later
CSCvw05821	Redirection and UDR does not take effect when traffic coming through an express route and destined to a service end point is redirected to a native load balancer or firewall.	25.0(1c) and later
CSCvw07392	Inter-site VxLAN traffic drops for a given VRF table when it is deleted and re-added. Packet capture on the CSR shows "Incomplete Adjacency" as follows: Punt 1 Count Code Cause 1 10 Incomplete adjacency <<<<<<<< Drop 1 Count Code Cause 1 94 Ipv4NoAdj	25.0(1c) and later
CSCvw07781	There is complete traffic loss for 180 seconds.	25.0(1c) and later
CSCvw24376	Inter region traffic is black-holed after the delete trigger for contracts/filter. It was observed that the TGW entry pointing to the remote region TGW is missing for the destination routes. On further debugging it was found that post delete trigger as part of re-add flow, when a describe call is sent to AWS got a reply with the state of this entry as "active" because of which a new create request is not being sent.	25.0(1c) and later
CSCvw39814	Infra VPC subnet route table entry for 0.0.0.0/0 route with TGW attachment as nh, is left as a stale entry upon being undeployed. There is no functional impact. Upon being redeployed, this entry is updated with the correct TGW attachment ID as nh.	25.0(1c) and later
CSCvw40737	SSH to a virtual machine's public IP address fails, despite the NSG allowing the traffic inbound. SSH to the private IP address of the virtual machine from within the VNet works.	25.0(1c) and later
CSCvw40818	After upgrading Cloud APIC, the Cloud Services Routers will be upgraded in two batches. The even set of CSRs are triggered for upgrade first. AFTER their upgrade is complete and all of the even CSRs are datapathReady, only then the odd set of CSRs will be triggered for upgrade. When even one of the upgrade of the even CSRs fail and they don't become datapathReady, the odd set of CSRs will not be triggered for upgrade. This is the behavior followed to avoid any traffic loss.	25.0(1c) and later

Bug ID	Description	Exists in
CSCvw48190	When Cloud APIC is restart, the VPN connection from a tenant's VNets will get deleted and re-created, one by one. This can be seen in the Azure activity logs. It should not impact traffic, as all connections are not deleted at the same time.	25.0(1c) and later
CSCvw49898	When the downgrading from the 5.2(1) release to the 5.0(2) release, traffic loss is expected until all of the CSRs are downgraded back to the 17.1 release. The traffic loss occurs because when the CSRs are getting downgraded to the 17.1 release, the CSR NIC1s will be in the backendPools and traffic from the spokes will still be forwarded to the native load balancer. The traffic gets blackholed until the CSRs get fully programmed with all the configurations in the 17.1 release.	25.0(1c) and later
CSCvw50918	Upon downgrading Cloud APIC, VPN connections between Cloud APIC and the cloud (AWS/Azure VPN gateway) will be deleted and re-created, causing traffic loss. Traffic loss is based on how quickly the VPN connections are deleted and re-created in AWS due to AWS throttling.	25.0(1c) and later
CSCvw51544	A user who is assigned a large number of security domains may not be able to create other Cisco ACI policies.	25.0(1c) and later
CSCvw55088	A user who is assigned a large number of security domains may not be able to create other Cisco ACI policies.	25.0(1c) and later
CSCvx91010	When TGW Connect is disabled, traffic loss is observed for about 8 minutes.	25.0(1c) and later
CSCvx98260	When delete followed by add operations of tenant or other resources (such as VPCs and contracts) are done within a short span of time, it is possible that the resource deployment in Google Cloud may get out of sync with the configuration on Cisco Cloud APIC. The likelihood of this happening is directly proportionate to the scale of configuration and how quickly the operations are done. We may see resources either not created or not deleted on Google Cloud to match the user configuration on Cisco Cloud APIC.	25.0(1c) and later
CSCvy06610	The eventmgr crashes when handling a fault triggered by a new cloud account.	25.0(1c) and later
CSCvy10936	Downgrading Cisco Cloud APIC from release 5.2(1) to 5.1(2) may cause CSRs to not be downgraded. The CSR release for 5.2(1) is 17.3.2, and the CSR version for release 5.1(2) is 17.3.1. After the Cisco Cloud APIC downgrade, the CSR version should be downgraded to 17.3.1, but it will not happen due to this bug.	25.0(1c) and later
CSCvy12722	Loss of traffic between a cloud and Cisco ACI On-Premises deployment.	25.0(1c) and later
CSCvy13369	After upgrading AWS, infra vPC peering does not get deleted.	25.0(1c) and later
CSCvy19286	There is traffic loss after downgrading from 5.2(1) to 5.1(2).	25.0(1c) and later
CSCvy28890	There is a loss in SSH connectivity to the Cisco Cloud APIC across reboots. But, after a few minutes, the connection should come back and users will be able to SSH in to the Cisco Cloud APIC again.	25.0(1c) and later
CSCvy28896	There is an increase in the connector's memory utilization. All of the CSR workflows rerunning might happen even after the setup is in the steady state.	25.0(1c) and later

Bug ID	Description	Exists in
CSCvy30314	After upgrading the Cisco Cloud APIC, on the TGW route tables, the default route (0.0.0.0/0) does not point to infra VPC attachment or is missing. In this case, traffic intended to get forwarded to the CSR will be dropped or forwarded to an invalid next-hop.	25.0(1c) and later
CSCvy33435	There is intersite traffic loss when TGW Connect is enabled.	25.0(1c) and later
CSCvy34180	Cloud Intersite traffic is dropped due to the CSR in the cloud site not advertising the EVPN routes.	25.0(1c) and later
CSCvy45517	The Cisco Cloud APIC GUI shows the total allowed count for CtxProfile, VRF (fvCtx), EPGs, and contracts. These numbers have been validated only for Azure-based deployments. For AWS deployments, the numbers supported are much lower.	25.0(1c) and later
CSCvy77233	<p>Routes for subnets that are not yet configured in Google Cloud may become visible on an external device. When you configure routes to be advertised to an external device, but don't actually configure subnets in the cloud that you intend to advertise the routes for, those routes are still advertised.</p> <p>Remote router may see routes that are advertised even when the subnets are not yet configured.</p> <p>The traffic will get dropped because the subnets are not actually configured.</p>	25.0(1c) and later
CSCvz11574	The cloud VRF egress route table is missing the route for 0.0.0.0/0 via the Internet Gateway (IGW), which leads to issues with ssh for VMs in the cloud VRF.	25.0(1c) and later
CSCvz20282	An upgrade to or downgrade from the Cloud APIC 5.2(1g) release to any release while using "Ignore Compatibility Check: no" will fail. The following fault is raised: "The upgrade has an upgrade status of Failed Due to Incompatible Desired Version."	25.0(1c) and later
CSCvz49747	<p>When delete followed by add operations of tenant or other resources (such as VPCs and contracts) are done within a short span of time, it is possible that the resource deployment in Google Cloud may get out of sync with the configuration on Cisco Cloud APIC. The likelihood of this happening is directly proportionate to the scale of configuration and how quickly the operations are done.</p> <p>We may see resources either not created or not deleted on Google Cloud to match the user configuration on Cisco Cloud APIC.</p>	25.0(1c) and later
CSCwa08564	UI dashboard shows the wrong status for inter-region connectivity.	25.0(2e) and later
CSCwa28888	This issue is hit in some cases where Cloud APIC is unable to deploy infra configuration, such as creating cloud routers in overlay-1 VPC. This is sometimes seen in new deployments, but not in the case of upgrade scenario. Cloud routers and other configurations do not get deployed in Google Cloud.	25.0(2e) and later
CSCwa36940	Transit gateway external connectivity is not getting deployed in regions where cloud context profiles are deployed.	25.0(2e) and later
CSCwa40705	When an IKEv1 tunnel is configured to a destination while another tunnel already exists to the same destination but with a different source interface, this tunnel will remain with the protocol shown as down.	25.0(2e) and later
CSCwa45047	This is not a functional issue. There will be no fault shown in the Cloud APIC UI if the border gateway protocol sessions of the transit gateway external connectivity are down.	25.0(2e) and later

Bug ID	Description	Exists in
CSCwa49263	<p>In VPC route table, the route table entry for a destination CIDR pointing to transit gateway is sometimes missing when a quick delete and add of tenant or contract is done or when we move from transit gateway connect to legacy transit gateway solution. This happens only with legacy transit gateway solution in either of the cases.</p> <p>This is a timing issue with the legacy transit gateway solution, where we create two transit gateways per hub network in a region. This can happen either if we move to legacy transit gateway solution or if legacy transit gateways are coming up for the first time.</p> <p>These conditions result in deleting the route table entry for a given destination CIDR and adding back the same entry at the same time. Due to an issue with the AWS API which returns a deleted route table entry as a non deleted entry, Cloud APIC deletes the wrong entry.</p>	25.0(2e) and later
CSCwa92698	<p>If the Cloud APIC infra CIDR has a collision with the reserved CIDR 172.17.0.0/16, connectivity to the Cat8kv VMs from Cloud APIC VM might fail. If the connectivity fails, the configuration push to Cat8kv will fail and Cat8kv will remain unreachable from Cloud APIC. A fault will be raised in the Cloud APIC.</p> <p>Currently 172.17.0.0/16 is reserved by Cloud APIC and it cannot be used as Infra CIDR. 172.17.0.0/16 is used by the docker network running on Cloud APIC.</p>	25.0(3k) and later
CSCwc11244	<p>In this case the HcloudCtxOper of infra VNet, which is used to deploy NLB/ALB, is down. Since one of the network interface associated with the Vnet is in a failed state, Cisco Cloud APIC is unable to add a new Cidr to the VNet.</p>	25.0(4k) and later
CSCwd20102	<p>Statistics page for Azure stats might contain missing data in the GUI when Azure cloud reports data incorrectly.</p>	25.1(1e) and later

Compatibility Information

This section lists the compatibility information for the Cisco Cloud Network Controller software. For Cisco Cloud Network Controller services compatibility information, see the [Nexus Dashboard and Services Compatibility Matrix](#). In addition to the information in this section, see the appropriate [Cisco Application Policy Infrastructure Controller Release Notes](#) and [Cisco Nexus Dashboard Orchestrator Release Notes](#) for compatibility information for those products.

- Refer to the [Nexus Dashboard and Services Compatibility Matrix](#) to see which version of Cisco Nexus Dashboard Orchestrator is compatible with Cloud Network Controller release 26.0(1c).
- Cloud Network Controller supports the following AWS regions:
 - Asia Pacific (Hong Kong)
 - Asia Pacific (Mumbai)
 - Asia Pacific (Osaka-Local)
 - Asia Pacific (Seoul)
 - Asia Pacific (Singapore)
 - Asia Pacific (Sydney)
 - Asia Pacific (Tokyo)
 - AWS GovCloud (US-Gov-West)

-
- Canada (Central)
 - EU (Frankfurt)
 - EU (Ireland)
 - EU (London)
 - EU (Milan)
 - EU (Stockholm)
 - South America (São Paulo)
 - US East (N. Virginia)
 - US East (Ohio)
 - US West (N. California)
 - US West (Oregon)
 - Cloud Network Controller supports the following Azure regions:
 - Australiacentral
 - Australiacentral2
 - Australiaeast
 - Australiasoutheast
 - Brazilsouth
 - Canadacentral
 - Canadaeast
 - Centralindia
 - Centralus
 - Eastasia
 - Eastus
 - Eastus2
 - Francecentral
 - Germanywestcentral
 - Japaneast
 - Japanwest
 - Koreacentral
 - Koreasouth
 - Northcentralus
 - Northeurope

- Norwayeast
 - Southafricanorth
 - Southcentralus
 - Southeastasia
 - Southindia
 - Switzerlandnorth
 - Uaenorth
 - Uksouth
 - Ukwest
 - Westcentralus
 - Westeurope
 - Westindia
 - Westus
 - Westus2
- Cloud Network Controller supports the following Azure Government cloud regions:
 - US DoD Central
 - US DoD East
 - US Gov Arizona
 - US Gov Texas
 - US Gov Virginia

Note: The US Gov Iowa region is not supported in Cisco Cloud Network Controller because Azure has deprecated support for this region.

- Cloud Network Controller supports all Google Cloud regions.

Related Content

See the [Cisco Cloud Network Controller](#) page for the documentation.

See the [Cisco Application Policy Infrastructure Controller \(APIC\)](#) page for the verified scability, Cisco Application Policy Infrastructure Controller (APIC), and Cisco Nexus Dashboard Orchestrator (NDO) documentation.

The documentation includes installation, upgrade, configuration, programming, and troubleshooting guides, technical references, release notes, and knowledge base (KB) articles, as well as other documentation. KB articles provide information about a specific use case or a specific topic.

By using the "Choose a topic" and "Choose a document type" fields of the APIC documentation website, you can narrow down the displayed documentation list to make it easier to find the desired document.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, send your comments to apic-docfeedback@cisco.com. We appreciate your feedback.

Legal Information

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL:

<http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2023-2024 Cisco Systems, Inc. All rights reserved.