White paper
Cisco public

Cisco Cloud ACI on Microsoft Azure

Contents

Cisco Cloud Application Centric Infrastructure (Cisco Cloud ACI) overview

In today's world, enterprises are undergoing increasing pressure to innovate rapidly, to keep up with competition and to increase IT agility to meet customer demands. To achieve these goals, businesses are choosing different infrastructure environments for deploying different types of applications. Some applications may be best suited to be hosted on-premises, whereas other applications may be best suited to be hosted in a public cloud, and yet others may benefit from hybrid deployments. In fact, hybrid cloud is becoming the new normal for many businesses.

However, in a hybrid cloud environment it is becoming more and more challenging to maintain a homogeneous enterprise operational model, comply with corporate security policies, and gain visibility across hybrid environments. Cisco® Cloud Application Centric Infrastructure (Cisco Cloud ACI®) is a comprehensive solution that provides simplified operations, consistent policy management and visibility across multiple on-premises data centers and public clouds or hybrid cloud environments. Cisco Cloud ACI allows customers running Cisco ACI™ in their on-premises data centers to extend their Cisco ACI policies to public clouds.

In an on-premises Cisco ACI datacenter, Cisco Application Policy Infrastructure Controller (APIC) is the single point of policy configuration and management for all the Cisco ACI switches deployed in the data center. When there is a need to seamlessly interconnect multiple Cisco ACI-powered data centers and selectively extend Cisco ACI constructs and policies across sites, Cisco Nexus Dashboard Orchestrator (NDO) (formerly known as Cisco Multi-Site Orchestrator [MSO]) enters the scene. NDO is a software solution that represents a single point of policy orchestration and visibility across multiple geographically dispersed ACI sites.

With the new Cisco Cloud ACI capabilities delivered in Cisco ACI Release 4.1 with AWS and Release 4.2 with Microsoft Azure, NDO can manage policies across multiple on-premises Cisco ACI data centers as well as public clouds. The policies configured from NDO can be pushed to different on-premises Cisco ACI sites and cloud sites. Cisco APIC controllers running on premises receive this policy from NDO, then render and enforce it locally. When extending Cisco ACI to the public cloud, a similar model applies. But there is a twist. Public cloud vendors do not speak Cisco ACI natively. Things such as Endpoint Groups (EPGs) or contracts are not familiar concepts there. NDO policies therefore need to be translated into cloud-native policy constructs. For example, contracts between Cisco ACI EPGs need to be translated into Application Security Groups and Network Security Groups on Microsoft Azure first, then applied to Microsoft Azure Virtual Machines. This policy translation and programming of the cloud environment is performed using a new component of the Cisco Cloud ACI solution called the Cisco Cloud Application Policy Infrastructure Controller (Cisco Cloud APIC or Cloud APIC).

The Cisco Cloud APIC runs natively on supported public clouds[1] to provide automated connectivity, policy translation and enhanced visibility of workloads in the public cloud. The Cisco Cloud APIC translates all the policies received from NDO and programs them into cloud-native constructs such as VNets (Virtual Network), application security groups, network security groups, outbound rules, inbound rules, etc.

This new solution brings a suite of capabilities to extend your on-premises data center into true hybrid cloud architectures, helping drive policy and operational consistency regardless of where your applications reside. It provides a single point of policy orchestration across hybrid environments, operational consistency, and visibility across different types of clouds such as AWS and Azure.

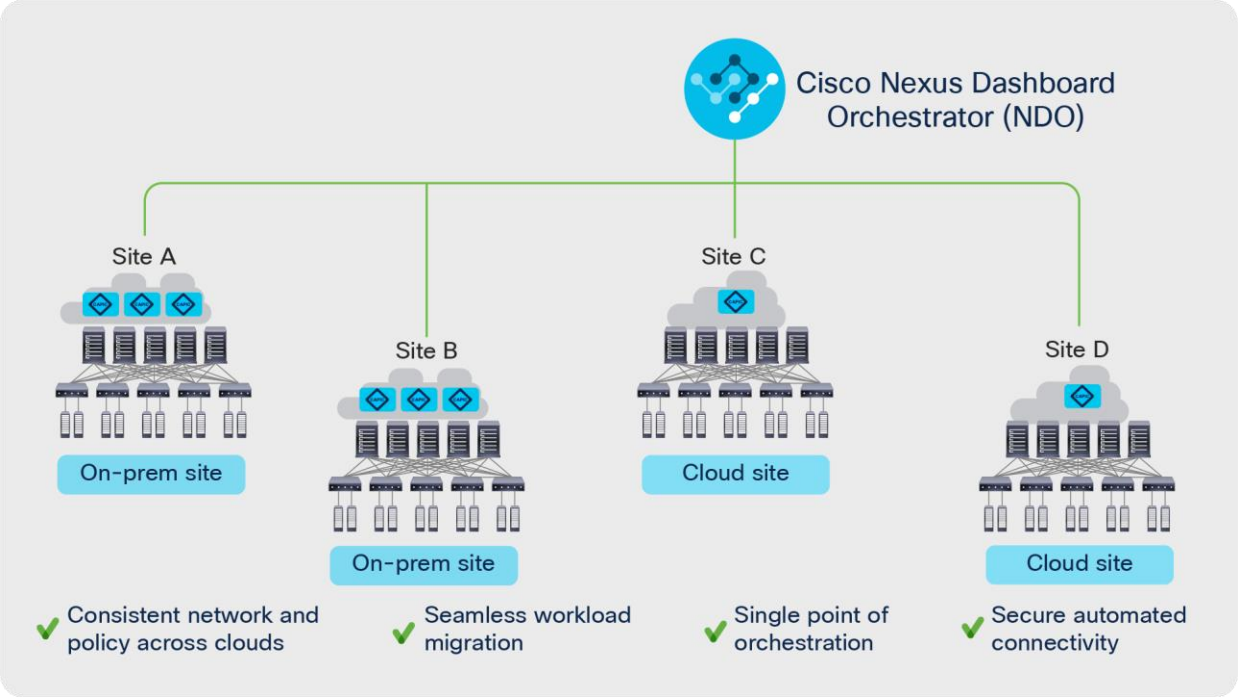

Figure 1.
High-level architecture of Cisco Cloud ACI

Figure 1 shows the overall high-level architecture of Cisco Cloud ACI with Cisco Nexus Dashboard Orchestrator acting as a central policy controller, managing policies across multiple on-premises Cisco ACI data centers as well as hybrid environments, with each cloud site being abstracted by its own Cloud APICs. The rest of this white paper discusses the architecture, benefits, use cases, and deployment of Cisco Cloud ACI on Microsoft Azure.

[1] See the data sheet and release notes for public cloud environment support information.

Challenges in hybrid cloud environments

As the adoption of hybrid cloud strategies grows, the industry is demanding consistent policy, security, and visibility everywhere with a simplified and consistent operating model. At the same time, the total cost of the solution must be kept under control to benefit from the hybrid cloud advantages.

The main challenges in building and operating a hybrid cloud environment are:

1. Automating the creation of secure interconnects between on-premises and public clouds
2. Dealing with the diverse and disjoint capabilities across on-premises private cloud and public cloud
3. Multiple panes of glass to manage, monitor, and operate hybrid cloud instances
4. Inconsistent security segmentation capabilities between on-premises and public clouds
5. Facing the learning curve associated with operating public cloud environment
6. Inability to leverage a consistent L4–L7 services integration in hybrid cloud deployments

Cisco ACI has delivered on the Software-Defined Networking (SDN) promise of intent-based network configuration and automation and further simplified operations by delivering control and visibility based on network policies that closely mimic your applications. The next phase of Cisco ACI must now address extending this policy-driven automation to hybrid environments. The Cisco Cloud ACI solution offers a coherent hybrid cloud strategy delivering on the key pillars of automation, security, and simplicity.

High-level architecture of Cisco Cloud ACI on Microsoft Azure

As briefly explained above and further depicted in Figure 2, an instance of NDO orchestrates multiple independent sites using a consistent policy model and provides a single pane of glass for centralized management and visibility. The sites can be on-premises Cisco ACI fabric sites with their own site-local APIC clusters, or cloud sites in Microsoft Azure with a Cloud APIC to manage the cloud site. Just as with a normal Cisco ACI Multi-Site architecture, all the sites are interconnected via a “plain” IP network. No need for IP multicast or Dynamic Host Configuration Protocol (DHCP) relay here. Just take care of providing IP connectivity, and NDO will be responsible for setting up the intersite overlay connectivity.

For more details on the Cisco ACI Multi-Site solution, refer to the following white paper:

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.pdf

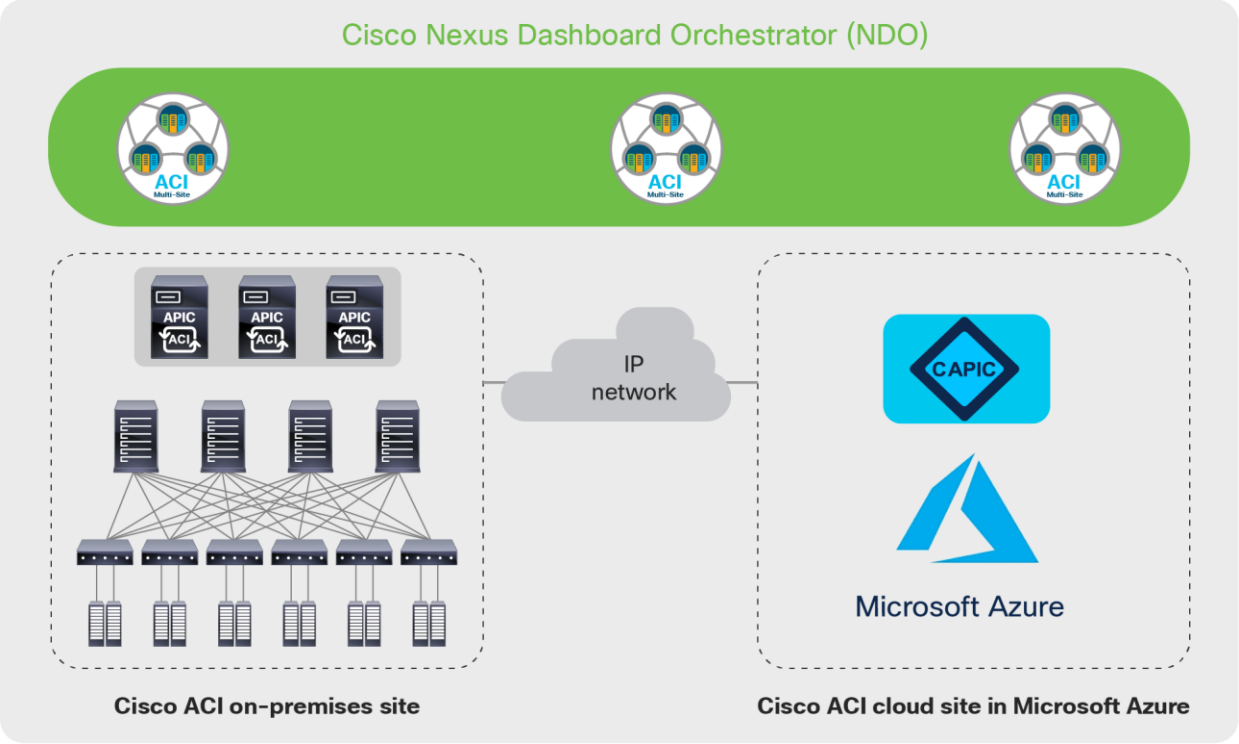

Figure 2.
Cisco Cloud ACI on Microsoft Azure Architecture

The key building blocks of the Cisco Cloud ACI architecture include the following:

- Cisco Nexus Dashboard Orchestrator (NDO)
- On-premises Cisco ACI site running Cisco ACI Release 4.1 or later (ACI release 4.2 or later is required for Azure cloud site support). The on-premises site needs to be equipped with at least one second-generation spine model (EX, FX, C, or GX) for intersite connectivity.
  - **Note:** Starting with ACI release 4.2, an on-premises ACI site is optional in the Cisco Cloud ACI architecture. The architecture also supports designs that have only cloud sites.
- Cisco Cloud ACI sites on public clouds that are supported by Cisco Cloud ACI solution. For the focus of this white paper, the public cloud in discussion is Microsoft Azure.
- Intersite connectivity. This includes the connectivity between on-premises and cloud sites, as well as the connectivity between cloud sites.
- The network policy mapping between the Cisco ACI policy model and the corresponding cloud-native policy model. For the focus of this white paper, an Azure-native policy model is used.

Cisco Nexus Dashboard Orchestrator (NDO)

In a Cisco ACI Multi-Site architecture, the Cisco Nexus Dashboard Orchestrator (NDO) is the single pane of glass for management of all the interconnected sites. It is a centralized place to define all the inter-site policies that can then be published to the individual Cisco ACI sites where the site-local APICs render them on the physical switches that build those fabrics.

With the Cisco Cloud ACI, NDO's orchestration functions expand to the cloud sites. It is responsible for site registration of both on-premises Cisco ACI data center sites and the cloud sites. It automates the creation of overlay connectivity between all the sites (on-premises and cloud). Continuing to be the central orchestrator of intersite policies, NDO now can, not only publish policies to on-premises Cisco ACI data center sites, but also push the same policies to cloud sites in Microsoft Azure. It is also capable of instrumenting the policy deployment among different sites by selectively distributing the policies to only the relevant sites. For instance, NDO can deploy the web front tier of an application into the cloud site in Microsoft Azure while keeping its compute and database tiers in the on-premises site. Through the NDO interface, network administrators can also regulate the communication flow between the on-premises site and Microsoft Azure as required by applications.

Cisco Cloud APIC on Microsoft Azure

The Cisco Cloud APIC is an important new solution component introduced in the architecture of Cisco Cloud ACI. It plays the equivalent of APIC for a cloud site. Like APIC for on-premises Cisco ACI sites, a Cloud APIC manages network policies for the cloud site that it is running on, by using the Cisco ACI network policy model to describe the policy intent. Cloud APIC is a software-only solution that is deployed using cloud-native instruments, for example, Azure Resource Manager (ARM) templates on Microsoft Azure. Network and security policies could be locally defined on the Cloud APIC for the cloud site, or globally defined on NDO and then distributed to the Cloud APIC. While the on-premises APIC renders the intended policies onto the Cisco ACI switches of the site, the Cloud APIC renders the policies onto the Microsoft Azure cloud network infrastructure. It accomplishes this task by translating the Cisco ACI network policies to the Microsoft Azure-native network policies and uses the Microsoft Azure-native policy API to automate the provisioning of the needed Microsoft Azure-native cloud resources, such as Virtual Network (VNet), cloud routers (Cisco CSR 1000V Series IOS® XE SD-WAN Routers and Microsoft Azure Virtual Network Gateway), Application Security Group, Network Security Group, etc. In a nutshell, the key functionalities of Cloud APIC include the following:

1. Provides a north-bound REST interface to configure cloud deployments
2. Accepts Cisco ACI Policy Model and other cloud-specific policies directly or from NDO
3. Performs endpoint discovery in the cloud site
4. Performs Cisco ACI Cloud Policy translation
5. Configures the cloud router's control plane
6. Configures the data-path between Cisco ACI Fabric and the cloud Site

Cisco Cloud APIC is a micro services-based software deployment of the APIC controller. Cisco Cloud APIC on Microsoft Azure is deployed and runs as an Microsoft Azure Virtual Machine using persistent storage volumes in Azure Managed Disks. The Virtual Machine for Cisco Cloud APIC are available at Microsoft Azure marketplace, and use Bring Your Own License (BYOL) model for licensing.

As ACI APIC for an on-premises ACI fabric, ACI Cloud APIC contains only policies and is not in the data forwarding path. Any downtime of the Cloud APIC will not impact network forwarding functionality or performance in the cloud site. The Microsoft Azure Virtual Machine of the Cloud APIC takes advantage of Azure Storage redundancy, high availability, and durability. Upon a failure in the Microsoft Azure Virtual Machine, it can always relaunch or restore to the previous state by rebuilding configuration and states from persistent storage and provide seamless Cloud APIC functionalities. Therefore, for simplicity and cost effectiveness, Cloud APIC is deployed as a single Microsoft Azure Virtual Machine in the initial release of Cisco Cloud ACI on Microsoft Azure. In the future, clustering of multiple virtual instances will be introduced for Cloud APIC to achieve higher scalability and instance level redundancy.

Key benefits of using Cisco Cloud ACI

The Cisco Cloud ACI solution provides the following key benefits:

Automation of on-premises to cloud interconnect

Cisco Cloud ACI on Microsoft Azure automates the configuration of end-to-end connectivity between an on-premises Cisco ACI fabric and Microsoft Azure. This connectivity can take place using IPsec VPN. Cisco Cloud APIC deploys a pair of Cisco CSR 1000V Series routers in Microsoft Azure and programs them to form an IPsec tunnel to the IPsec terminator that is installed on premises. Once the IPsec tunnel is up, NDO configures the BGP EVPN control plane as well as VXLAN tunnel between the on-premises second-generation Cisco ACI spines and the Cisco CSR 1000V Series router deployed on Microsoft Azure. This end-to-end automation makes hybrid cloud connectivity seamless, reducing the configuration time, risk of errors, and accelerating the pace of deployment and rate of change. Later in the document, more technical details are provided regarding this automated intersite connectivity.

Universal policy model

Both Cisco ACI and Microsoft Azure use group-based network and security policy models. In a nutshell, the logical network constructs of the Cisco ACI network policy model consist of tenants, Bridge Domains (BDs), bridge-domain subnets, Endpoint Groups (EPGs), and contracts. Microsoft Azure uses slightly different constructs: Resource Groups, Virtual Network (VNet), Application Security Groups (ASG), and Network Security Groups (NSGs), plus outbound rules and inbound rules. As shown in Figure 3, Cisco ACI classifies endpoints into EPGs and uses contracts to enforce communication policies between these EPGs. Microsoft Azure uses Application Security Groups (ASGs), Network Security Groups (NSGs), outbound rules, and inbound rules for classification and policy enforcement. Figure 4 gives one example of Microsoft Azure security policy enforcement using Network Security Groups.

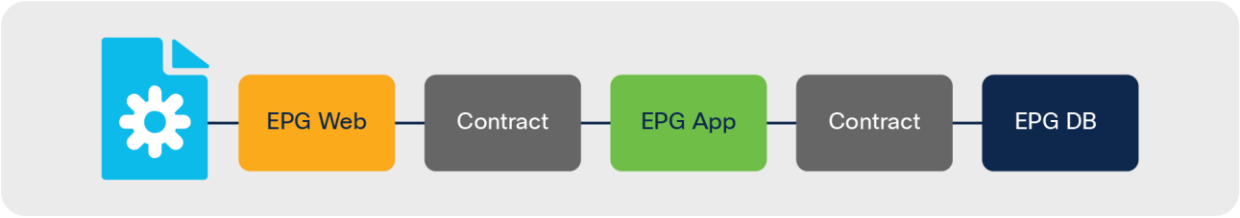

Figure 3.
Cisco ACI EPG-based network model

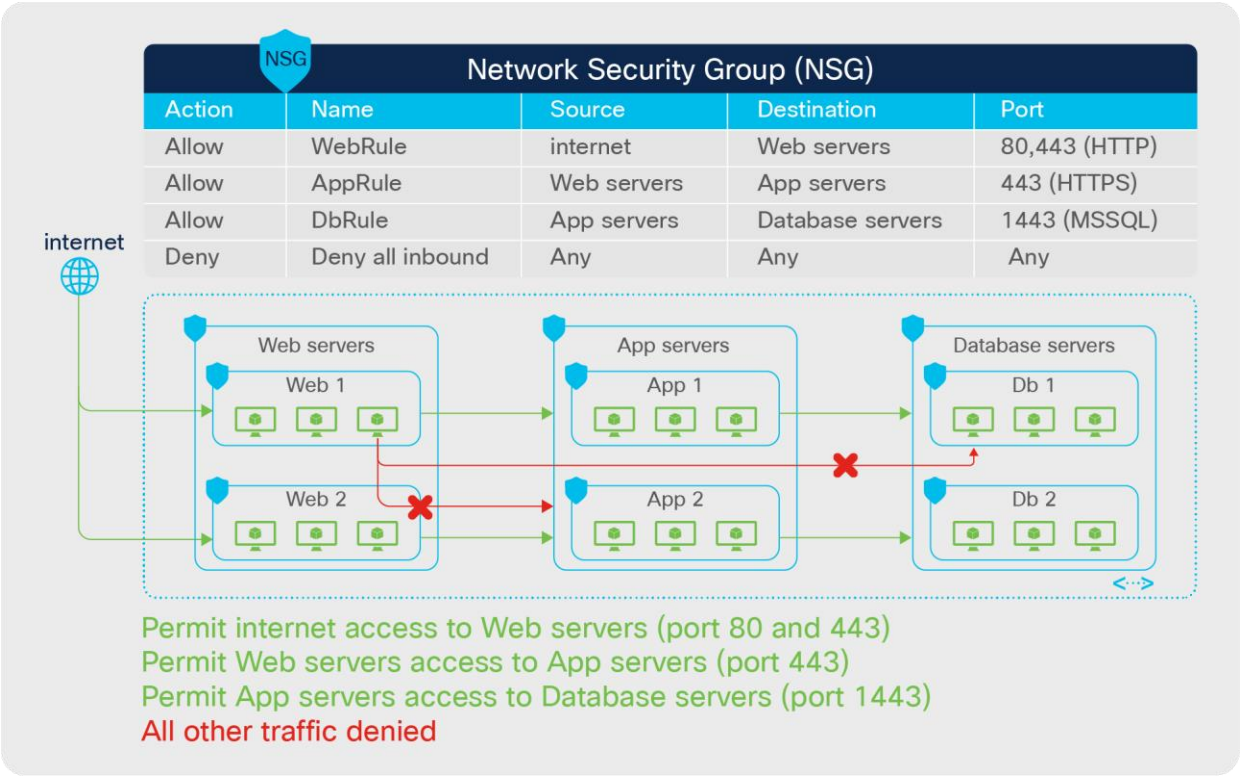

Figure 4.
Microsoft Azure Security Group-based network model

Granular and accurate mapping between these two network policy models is crucial to ensure the correct deployment of network policies across Cisco ACI and Microsoft Azure. Figures 5 and 6 shows how Cloud APIC handles this policy mapping.

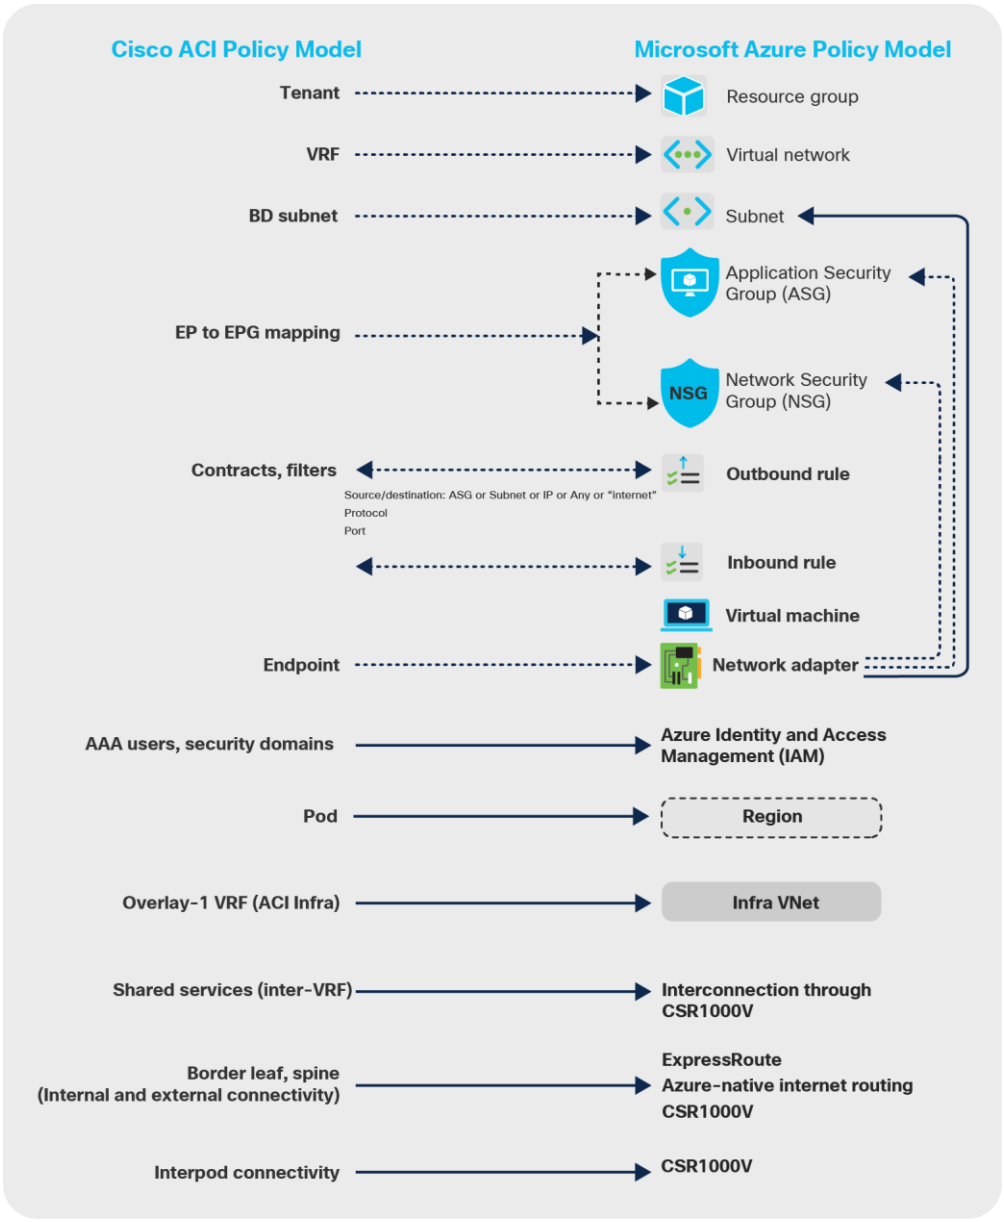

Figure 5.
Cisco ACI Policy Model to Microsoft Azure Policy Model mapping (NSG per subnet)

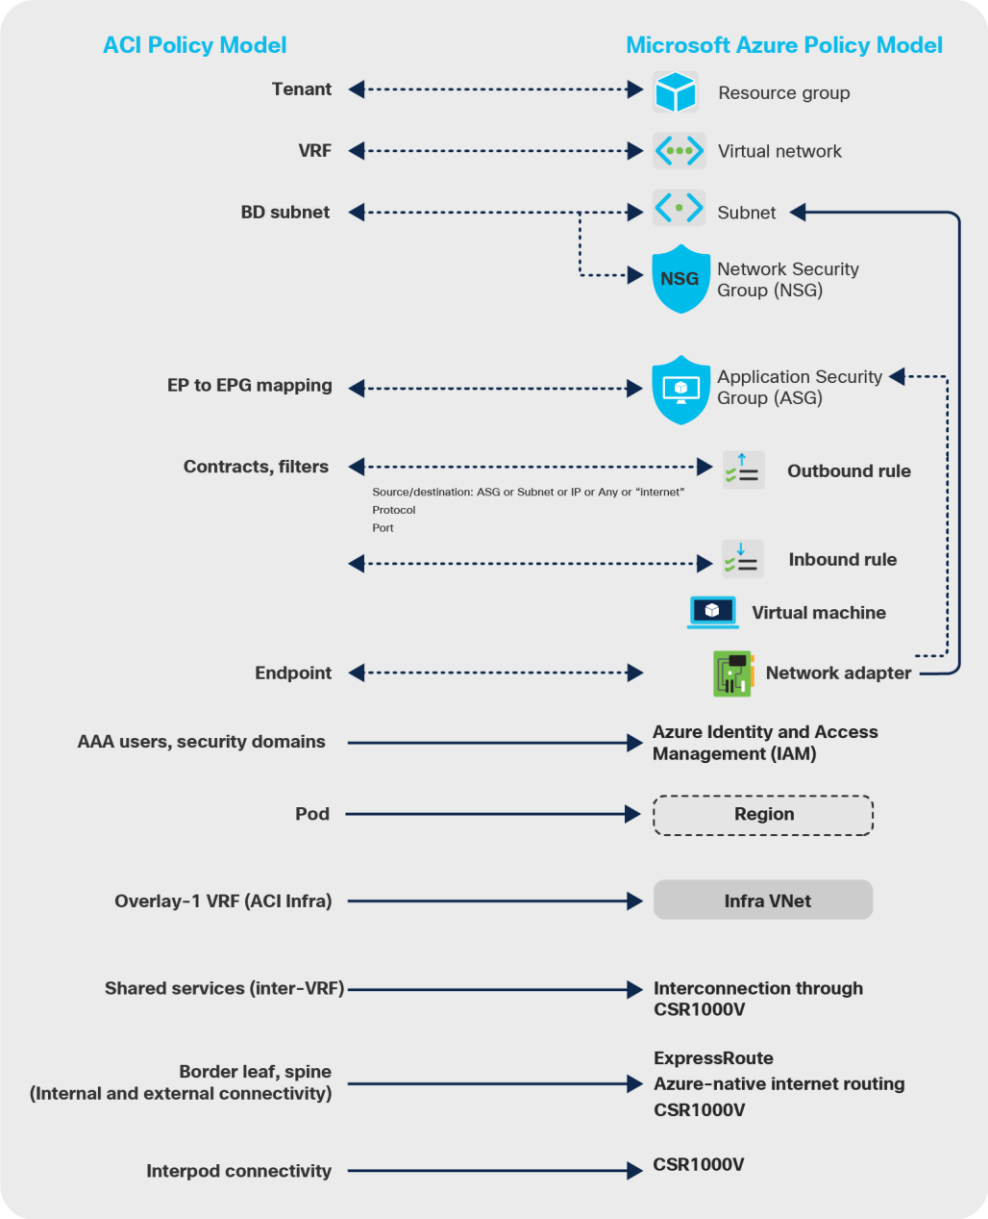

Figure 6.
Cisco ACI Policy Model to Microsoft Azure Policy Model mapping (NSG per EPG)

Unified network management and operations

Cisco Nexus Dashboard Orchestrator (NDO) provides end-to-end visibility and health of all the endpoints managed by it across on-premises and public cloud environments, giving a single place to monitor the health, performance, and operational status of hybrid cloud environments. NDO being the single point of policy configuration and orchestration, this highly reduces the operational complexity of operating across hybrid environments.

Consumption of cloud-native services

Cisco Cloud ACI uses cloud-native constructs for enforcing security and workload segmentation policies. This facilitates the consumption of cloud-native services by workloads managed by Cisco Cloud ACI. Administrators can define policies from NDO regulating which workloads can access which cloud-native services. Inside the Microsoft Azure environment, these policies then get programmed as security group rules that either allow or deny the workloads to access specific services. To the APIC administrator, whether an application is deployed in part on premises and in part in the cloud does not matter. Familiar EPGs and contracts govern that application's communications. One of the key benefits of this solution is to make Cloud-bursting easy.

Digging further: looking inside an ACI Cloud site

Hub-and-spoke topology with an infra VNet

When running Cisco Cloud ACI, a hub-and-spoke Microsoft infra VNet topology is deployed in the Microsoft Azure-native network infrastructure, based on the translation of Cisco ACI policies into Microsoft Azure-native policies. The hub is an infra VNet while the spokes are user VNets where the application endpoints are deployed. Figure 7 shows this concept using IPsec with Azure Virtual Network Gateway (VNG), and Figure 8 shows this concept using Azure Virtual Network peering (VNet peering).

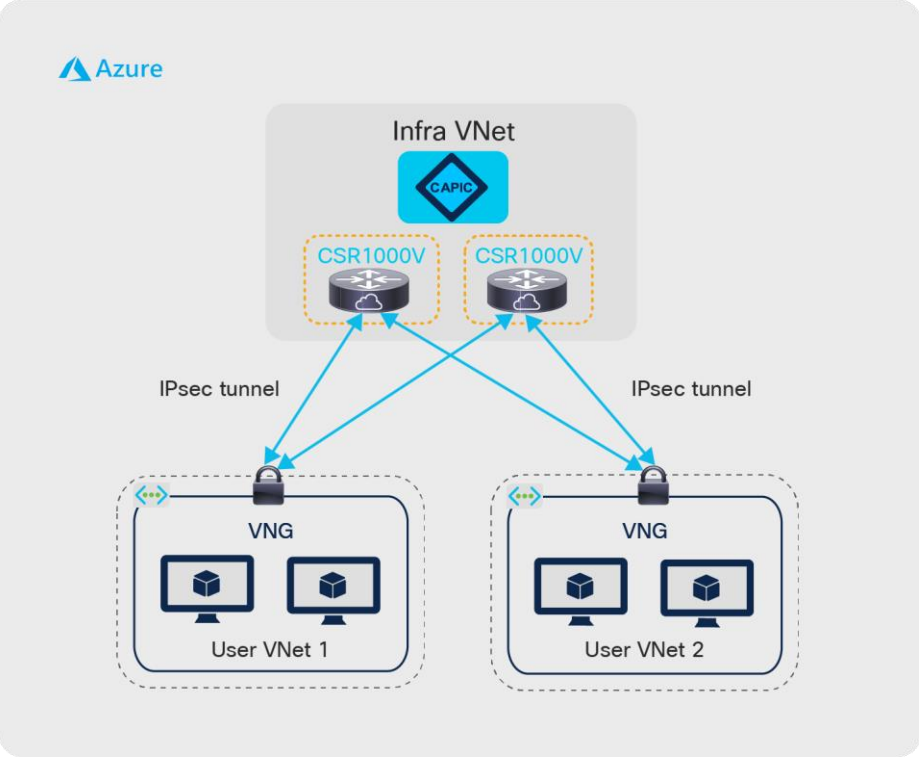

Figure 7.
Inside the cloud using IPsec tunnel with VNG

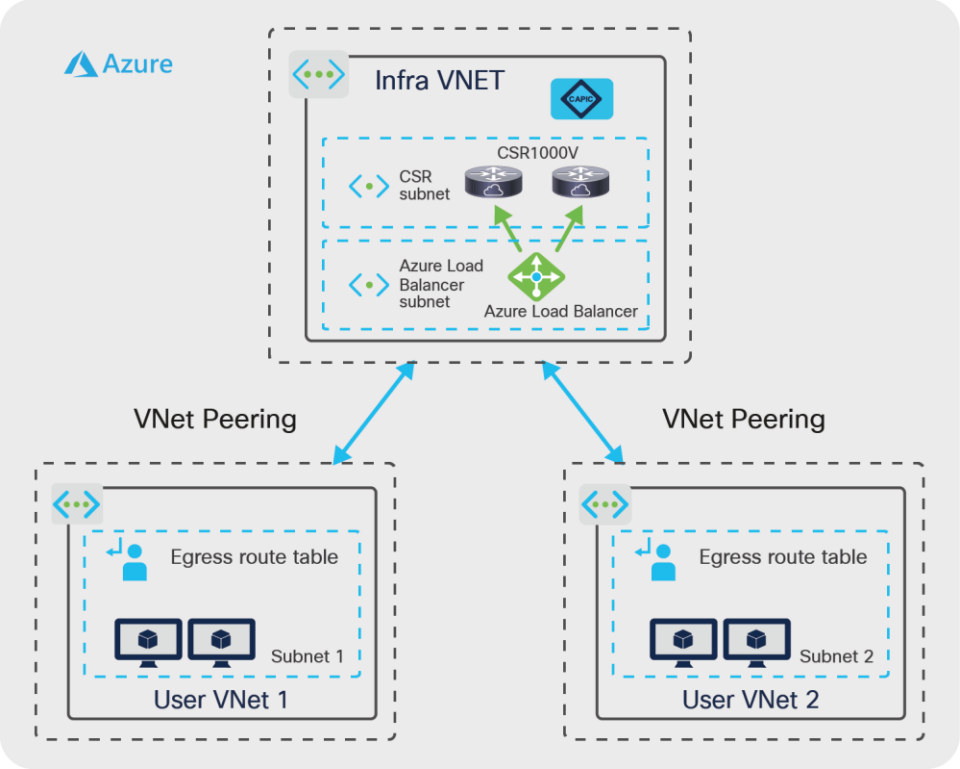

Figure 8.
Inside the cloud using VNet peering

The infra VNet carries the logical role of the on-premises Cisco ACI infra VRF. It is where Cloud APIC is deployed and running. It is automatically created during the deployment of Cloud APIC. Cloud APIC then deploys a pair of Cisco CSR 1000V Series routers in this infra VNet as cloud routers responsible for providing the virtual underlay connectivity, including the internal connectivity within the cloud site and the intersite connectivity to the on-premises Cisco ACI sites or other cloud ACI sites.

A user VNet is equivalent to a tenant VRF in the Cisco ACI network policy model. Cloud APIC creates a user VNet when an ACI tenant VRF needs to be deployed or extended to the cloud site. Within the user VNet, Cloud APIC provisions an Azure Virtual Network Gateway (VNG) or VNet peering to connect to the Cisco CSR 1000V Series routers in the infra VNet.

In case of the example described in Figure 7 using IPsec tunnels with VNGs, the Cloud APIC automatically provisions IPsec tunnels between the infra VNet Cisco CSR 1000V Series routers and the user VNet Azure VNGs to provide proper network connectivity within the cloud site. The infra VNet functions as a transit VNet to support route exchanges between user VNets through the IPsec tunnels. Endpoint communication between user VNets, and one between a user VNet and the on-premises site goes through VNGs in user VNet and Cisco CSR 1000Vs in the infra VNet.

In case of the example described in Figure 8 using VNet peering, the Cloud APIC automatically provisions the Azure Load Balancer in the infra VNet in addition to Cisco CSR 1000V Series routers, and provisions UDRs (User Defined Routes) in user VNets to provide network connectivity within the cloud site. Endpoint communication between user VNets, and one between a user VNet and the on-premises site go through the Cisco CSR 1000Vs via the Azure Load Balancer in the infra VNet based on the UDRs in the user vNet.

An ACI Cloud site across multiple Microsoft Azure regions

An ACI cloud site in Microsoft Azure can span across multiple Microsoft Azure regions. While the entire cloud site is managed by the same Cloud APIC, each region can have its own infra VNet with a pair of Cisco CSR 1000V Series routers for networking (as shown in Figure 9 and 11) or share the infra VNet and the Cisco CSR 1000Vs in other regions (as shown in Figure 10 and 12). User VNets can be deployed into any of the regions with an Azure Virtual Network Gateway (VNG) or Azure Virtual Network Peering (VNet peering).

In case of using VNG, all of the user VNets are connected to all the infra VNets through IPsec tunnels. IPsec tunnels are established between the Virtual Network Gateway in each user VNet to all Cisco CSR 1000V Series routers in the Infra VNet. This provides a full intra-region and inter-region connectivity with virtually a spine-leaf CLOS architecture between infra VNets and user VNets. The provisioning of the VNets, the cloud routers, and the IPsec tunnels is fully automated by the Cloud APIC. Meanwhile, NDO automates the provisioning of overlay network connectivity between the on-premises and cloud sites.

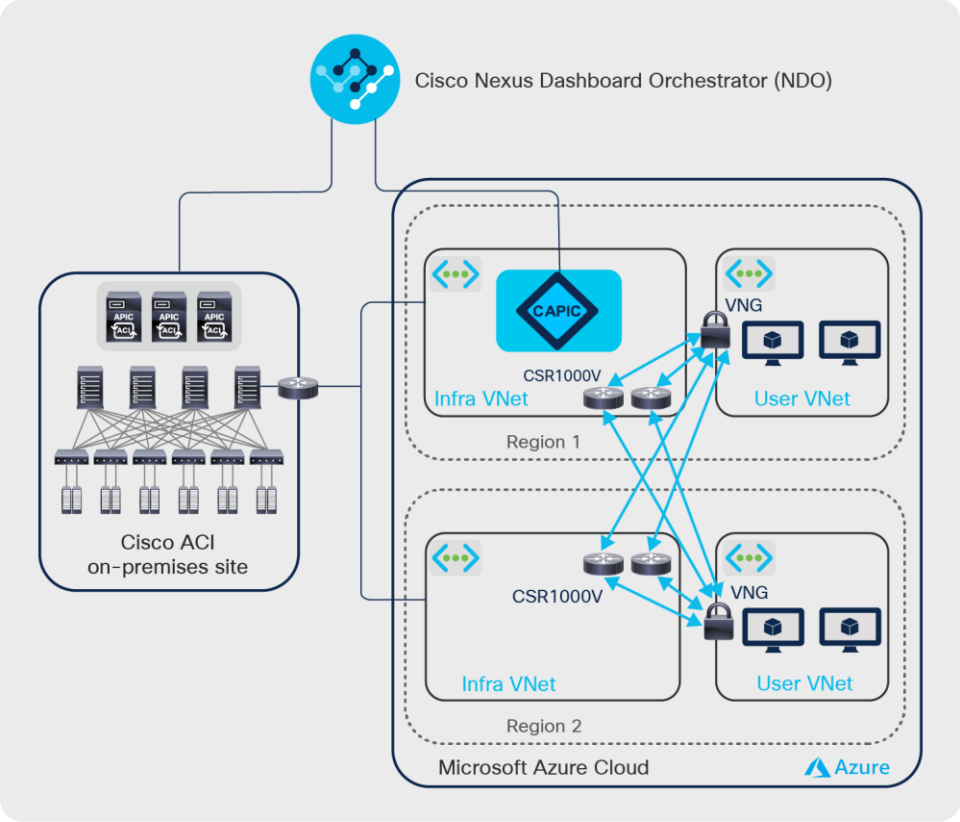

Figure 9.
Cisco Cloud ACI Microsoft Azure multi-region site with dedicated Infra VNet using IPsec tunnels with VNGs

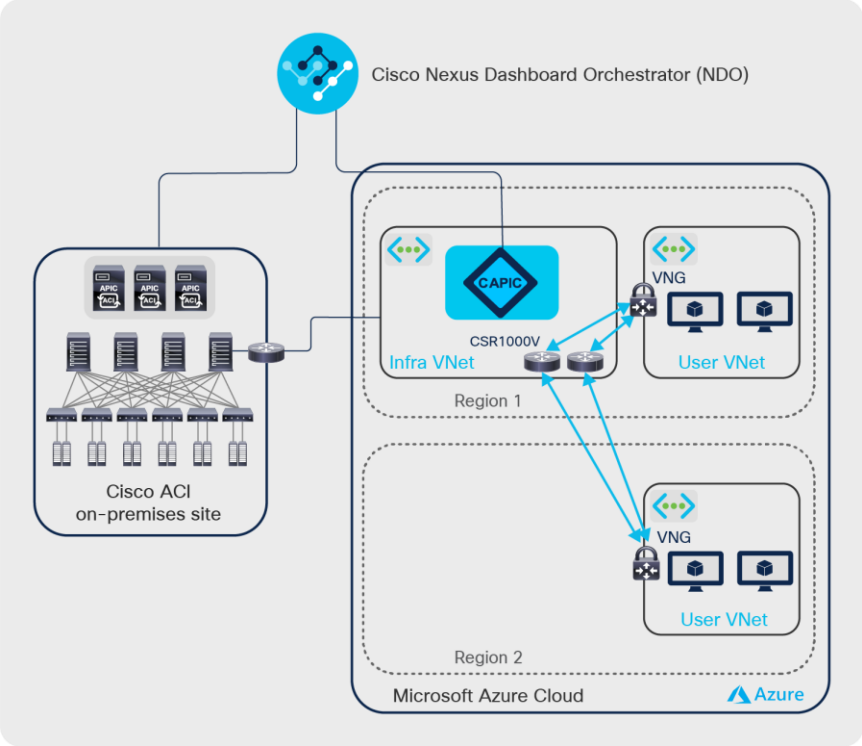

Figure 10.
Cisco Cloud ACI Microsoft Azure multi-region site with shared infra VNet using IPsec tunnels with VNGs

In case of using VNet peering, the Azure Load Balancer is automatically deployed in the infra VNet and configured by Cloud APIC in addition to the Cisco CSR 1000Vs Series routers. Cloud APIC also provisions UDRs (User Defined Routes) in user VNets to provide network connectivity within the region and across the regions. Inter-region connectivity goes through the global VNet peering with infra VNet in the other region.

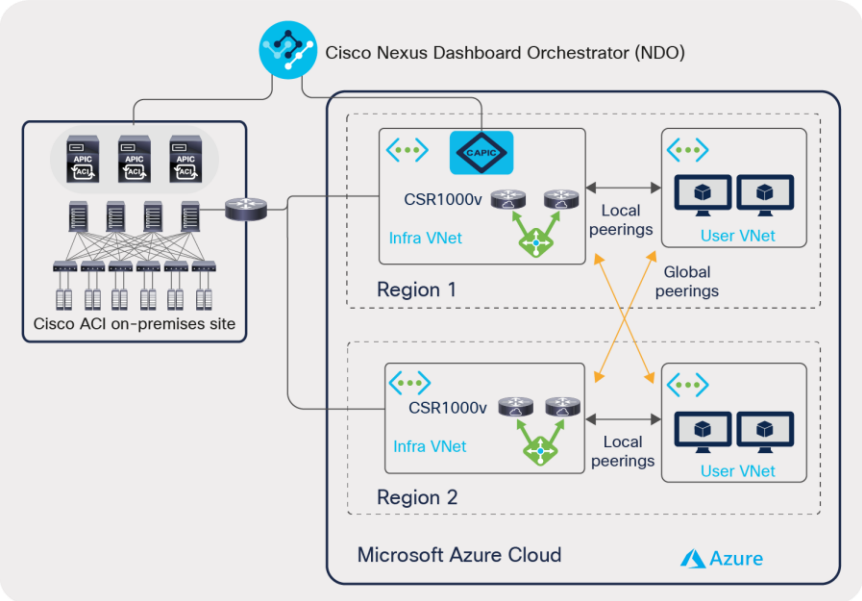

Figure 11.
Cisco Cloud ACI Microsoft Azure multi-region site with dedicated infra VNet using VNet peering

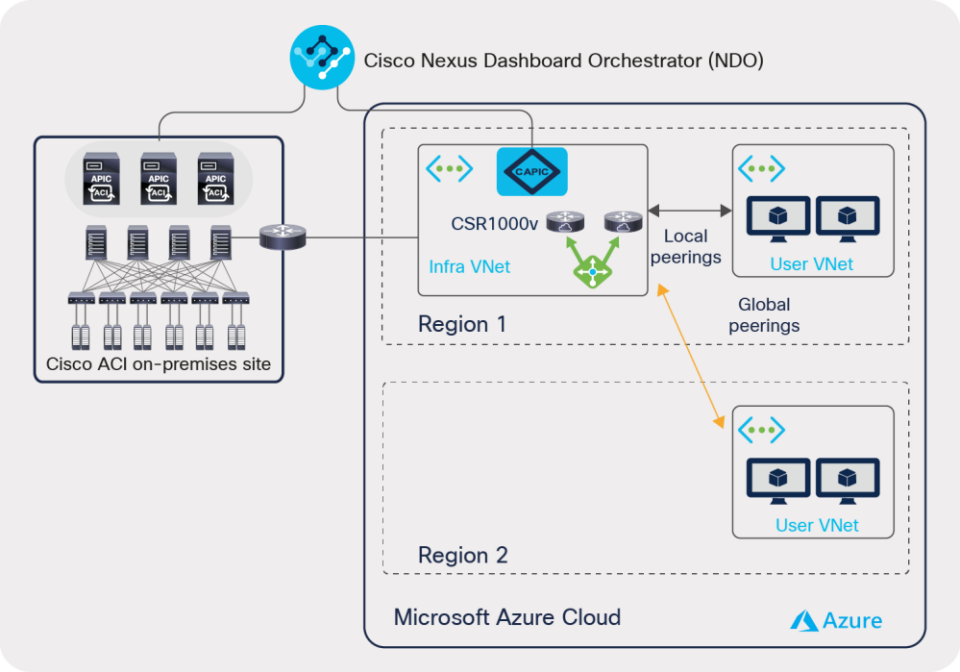

Figure 12.
Cisco Cloud ACI Microsoft Azure multi-region site with shared infra VNet using IPsec tunnels with VNet peering

Traffic flows inside the cloud site

Traffic between two endpoints in the same VNet is routed locally in the VNet. It does not need to go through the infra VNet or hairpin back on premises. Traffic between two endpoints that are in different VNets needs to be routed through the Cisco CSR 1000V Series router in the infra VNet. Finally, traffic between an endpoint in a cloud site and an endpoint on premises needs to be routed through the Cisco CSR 1000V Series router in the infra VNet.

In case of VNet peering, the traffic from user VNet to the other VNet is redirected to the Azure Load Balancer in the infra VNet first based on the UDR programmed in the user VNet and then the Azure Load Balancer load balances the traffic to one of the Cisco CSR 1000V Series routers in the infra VNet. Then traffic is routed through the router.

Intersite connectivity

The underlay network between on-premises and cloud sites

The on-premises Cisco ACI sites and the ACI cloud site in Microsoft Azure are connected through an IP network that can be IPsec VPN over the internet, or through Azure ExpressRoute[2]. In case of IPsec VPN, the Cloud Cisco CSR 1000V Series routers in the Infra VNet need to be programmed to establish IPsec peering with the IPsec device located on premises (at the physical Cisco ACI site). This underlay network provides IP reachability for the overlay control plane and data plane between the two sites. This is represented in Figure 13:

[2] As of Cisco ACI Release 5.1, though this solution does not take care of an automatic creation of Azure ExpressRoute, Azure ExpressRoute can be used in conjunction with the solution.

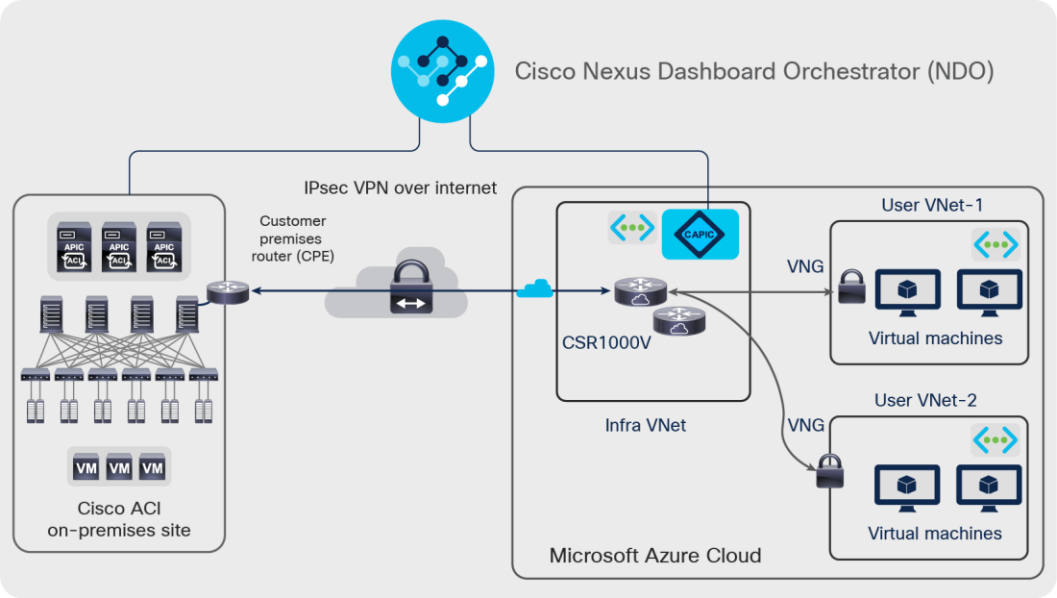

Figure 13.
The underlay network between on-premises and cloud sites

Note: Though the topology described in this section utilizes IPsec tunnels with VNGs in user VPCs, VNet peering is also supported in conjunction with intersite connectivity using IPsec VPN over the Internet.

The overlay network between on-premises and cloud sites

The overlay network between the on-premises and cloud sites runs BGP EVPN as its control plane and uses VXLAN encapsulation and tunneling as its data plane. The use of VXLAN is to identify the right routing domain for VRF stretch across on-premises Cisco ACI fabric and the clouds.

BGP EVPN sessions are established between the on-premises Cisco ACI spine switches and the Cisco CSR 1000V Series cloud routers in the Infra VNet of the cloud site. Tenant host routes and prefix routes are exchanged between the two sites as BGP EVPN route type-2 (host) and type-5 (prefix). The provisioning of this overlay network connectivity is automated by NDO. Figure 14 zooms in on this logical/physical architecture. On-premises spines connect to an intersite network (called ISN or IPN for inter-pod network). That IPN layer then connects to an on-premises IPsec router that initiates IPsec tunnels to the Cisco CSR 1000V Series routers in the Microsoft Azure infra VNet. MP-BGP EVPN sessions are established between the ACI spine switches and the Cisco CSR 1000V Series routers in the Microsoft Azure infra VNet over the IPN network and the IPsec tunnels.

You might need to adjust Maximum Transmission Unit (MTU) size on ACI Control Plane MTU policy for BGP EVPN control-plane and on your endpoints for data plane to avoid fragmentation because of IPsec tunnels and VXLAN encapsulation overhead. Otherwise, fragmentation by devices in the network could degrade overall performance. For example, if MTU of the involved endpoints is adjusted to 1300 bytes, this would account for the additional 50 bytes from VXLAN and around 100 bytes for IPsec overhead to go over the internet where the common value of MTU is 1500 bytes. If adjusting the MTU size on your endpoints is not allowed or not preferable, you need to configure the TCP Maximum Segment Size (MSS) Adjustment on CSR 1000Vs from cAPIC. This configuration option is available starting from Cisco Cloud APIC Release 4.2(4q), 4.2(5n), and 5.0(2i).

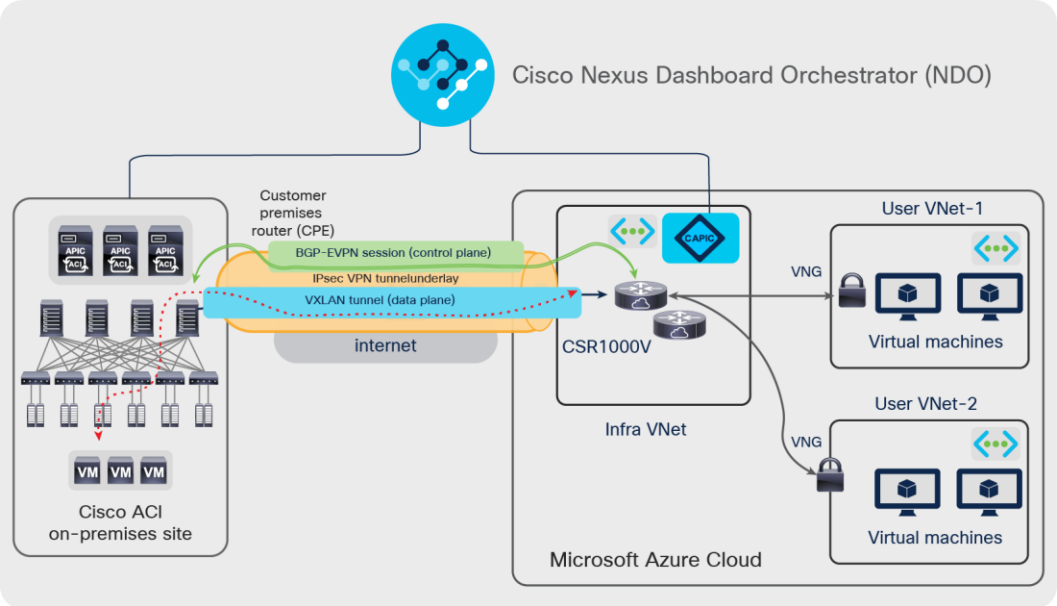

Figure 14.
The overlay network between on-premises and cloud sites

Use-case scenarios

Cisco Cloud ACI enables customers to achieve the following scenarios:

High Availability of applications across on-premises and cloud sites

Cisco Cloud ACI enables customers to deploy an application in High Availability by stretching it across on-premises and cloud sites. This makes it possible to have a multitier application deployed across a hybrid-cloud environment all in the same Virtual Routing and Forwarding domain (VRF).[3]

Customers having application tiers deployed on an on-premises Cisco ACI site can now add new application tiers in a cloud site interacting with the on-premises tiers using consistent policies.

Applications can fail over between the on-premises Cisco ACI site and the ACI cloud site during a disaster recovery, or the application can be deployed in an active/active mode, where both on-premises application tiers and cloud tiers are active. A Global Load Balancer can be configured to distribute traffic between the two sites.

For example, the web and application tiers of an application (two EPGs) in the same VRF running in the on-premises Cisco ACI data center. By stretching the same VRF to the cloud site, we can deploy web and application tiers in the cloud site and can be configured as an active/active between the on-premises and cloud sites. You can also deploy the on-premises web and application tiers as active, and the cloud tiers can act as standby, and can fail over to a cloud tier in case of a disaster recovery.

[3] Note: Extending a broadcast domain between an on-premises site and the cloud is not possible. Cloud vendors typically do not run broadcast or multicast and never face unknown unicast situations.

All of this can be achieved by using NDO as a single point of orchestration; you can configure contracts between these tiers spread across the hybrid cloud environment. Simply publishing this policy from NDO to both sites programs the end-to-end constructs required to implement the workload segmentation policy. This is shown in Figure 15.

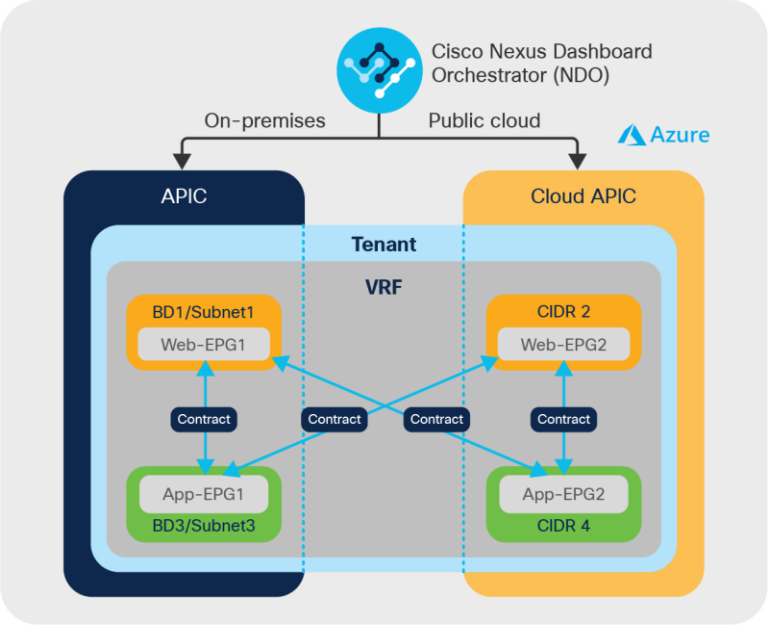

Figure 15.
Stretched routing domain with intersite contracts

Cloud bursting: stretch an application tier (EPG) to cloud with consistent segmentation

Cisco Cloud ACI enables customers to stretch an application tier across the on-premises and cloud sites in Microsoft Azure, which means that an ACI EPG can be stretched across the on-premises and Microsoft Azure sites. This enables customers to burst a tier to Microsoft Azure Cloud during times of peak load and access other tiers of the application on premises via secure segmentation policies. Even multiple tiers can be burst to the cloud site in Microsoft Azure and still maintain the same level of policy and consistent segmentation irrespective of where the workloads are deployed.

From NDO, you can either create a new EPG that can then be stretched or import an existing EPG from the on-premises site and stretch it to Microsoft Azure. This is achieved just as you would with regular Cisco ACI Multi-Site, using templates and schemas. Once that is done, configure the site-local properties that define how the EPG should classify its member endpoints.

When associating an EPG with an on-premises ACI site, you can either associate the EPG to a Virtual Machine Manager (VMM) domain, or to static ports or a VLAN/port combination to classify endpoints on premises. When the same EPG is associated with a cloud site in Microsoft Azure through NDO, EPG membership classification criteria can then be based on region, tag, or IP address/subnets.

Stretching an EPG does not mean extending a broadcast domain from an on-premises site to the cloud, though; it simply means that you can create an EPG with members on premises and in the cloud, using different subnets. Once two or more endpoints are in the same EPG, communication flows freely inside that EPG.

Example: Let's say you have an application with web and application tiers deployed in an on-premises ACI site. During a peak load time, you burst either web tier or both web and application tiers to the cloud site in Microsoft Azure. You can do that seamlessly with just a couple of button clicks from NDO and stretch the tiers to Microsoft Azure with the same level of security and segmentation as their on-premises workloads. Now, contracts between stretched Web EPG and on-premises EPGs or Cloud EPGs can be configured as you normally would with an on-premises ACI. Cloud-bursting doesn't get any easier. This is shown in Figures 16.

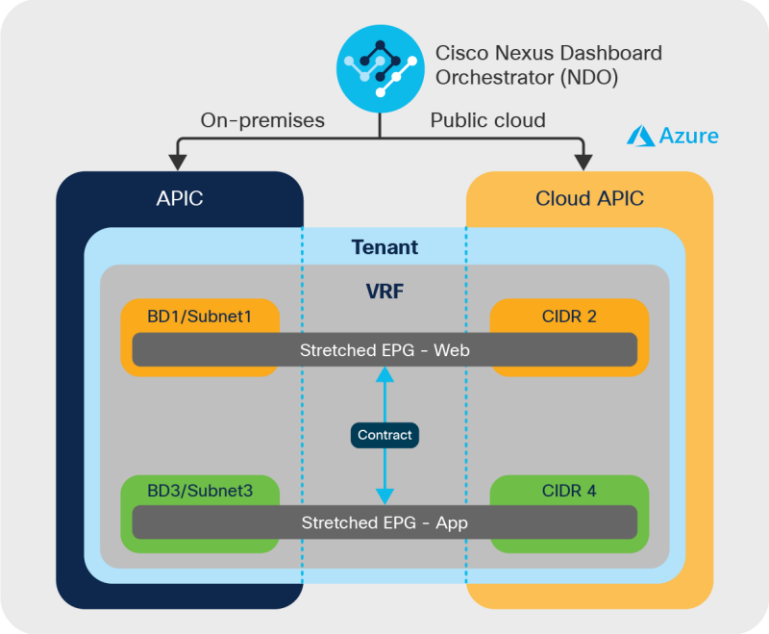

Figure 16.
Stretched EPGs across sites

Shared services across hybrid cloud

Shared services such as DNS, Active Directory (AD), or other authentication services deployed in a tenant in an on-premises ACI site can be securely consumed by endpoints in other tenants spread across other Cisco ACI and Microsoft Azure sites. This enables services from a single on-premises provider tenant to be consumed by multiple consumer tenants spread across the on-premises ACI site and cloud sites in Microsoft Azure.*

This makes it easier to deploy new applications in the cloud and consume shared services from the brownfield network on premises, without having to redeploy them for applications hosted in the cloud.

* Prior to Cisco ACI Release 5.1, inter-tenant shared service is not supported; for example, a contract between a provider EPG in the cloud and a consumer EPG on premises is not supported if the EPGs are in different tenants. Starting from Cisco ACI Release 5.1, it's supported if the tenants are stretched across an on-premises site and a cloud site. Stretching the consumer and provider VRFs across on-prem and cloud sites are not mandatory, the VRFs can be local.

Example: Let's say there are DNS servers deployed in Tenant-1, an on-premises ACI site. Workloads that are part of the Web-EPG deployed on the cloud site in Microsoft Azure in Tenant-1 can access these DNS-shared services from the on-premises ACI site via an inter-VRF contract between the DNS-EPG and the Web-EPG. This is shown in Figure 17.

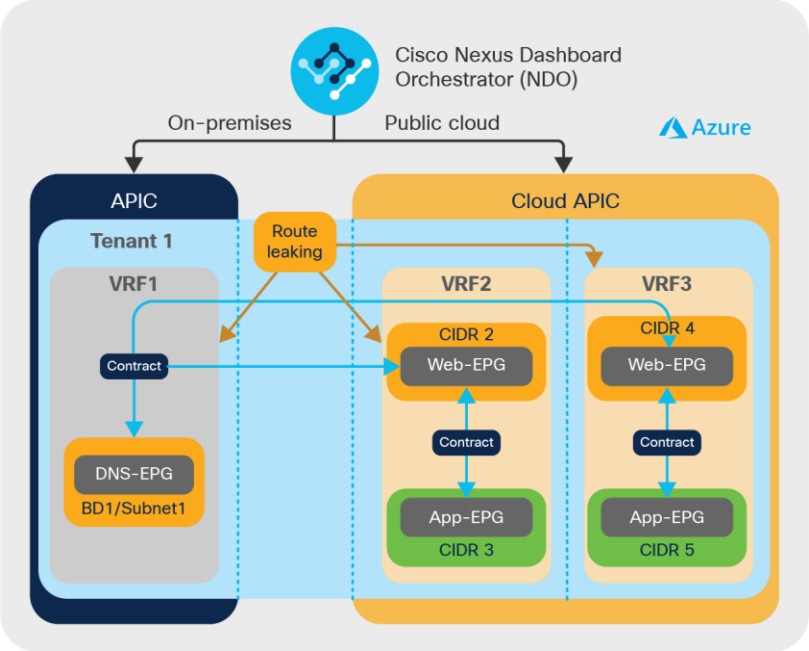

Figure 17.
Cross-site shared services

Note: As of this writing, inter-tenant shared service is not supported.

External connectivity to the internet via the cloud or on-premises

External connectivity to the internet for workloads in Microsoft Azure Cloud can be configured in two ways:

1. Cloud L3Out: A cloud-local Internet connection (also called L3Out in Cisco ACI terminology) can be defined for workloads deployed on Microsoft Azure. It is achieved by configuring a cloud-External EPG in NDO for the cloud site in Microsoft Azure. Once an EPG has a contract with the cloud-External EPG, security rule will be created in the Network Security Group on Microsoft Azure site through NDO. So that workloads in the EPG can reach external network. Appropriate routes will be programmed into the VNet by Azure internally.
2. On-premises L3Out: Some customer environments require all the traffic from a VNet in Microsoft Azure to transit to an on-premises site and be inspected through an on-premises firewall/IDS/IPS before the traffic exits to, or enters from, the internet. This can also be achieved by defining an on-premises L3Out as the internet exit for traffic and associating the cloud endpoints to that EPG via a contract.

Customers have full control over external network connectivity options for workloads deployed in the cloud and can choose to redirect the traffic for inspection by various services deployed on premises by using Cisco ACI service graphs. All this can be done through a single policy, and end-to-end connectivity can be automated and orchestrated by Cisco Nexus Dashboard Orchestrator, greatly simplifying the operational workflows.

Example: When an administrator configures a cloud-local L3Out in the Microsoft Azure environment (as shown in Figure 18), each Microsoft Azure VNet will have an external connectivity, and the Microsoft Azure Virtual Machines in the VNet can directly communicate with the internet based on the policies defined.

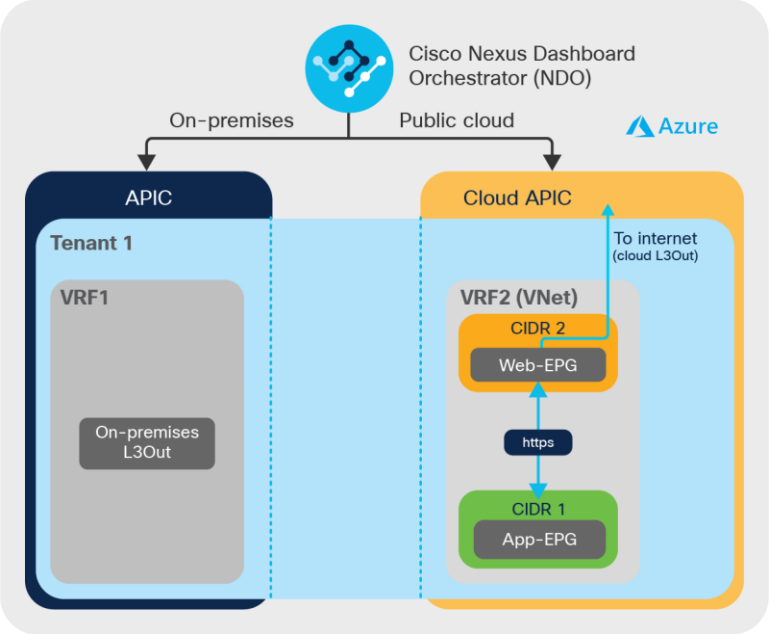

Figure 18.
An example of cloud L3Out

If the administrator defines an on-premises Cisco ACI L3Out (as shown in Figure 19) and forces cloud instances to use that L3Out, then all traffic from Microsoft Azure Virtual Machines reaches the Cisco CSR 1000V Series router via the VPN tunnels, and will be sent on premises over the VXLAN tunnel running over the IPsec tunnel. Traffic can then exit via the on-premises Cisco ACI L3Out instead of using the internet access from VNet directly.

Once traffic reaches the on-premises Cisco ACI site, the administrator may choose to subject this traffic to various inspections using the service chain options in Cisco ACI and then let the traffic exit to the Internet.

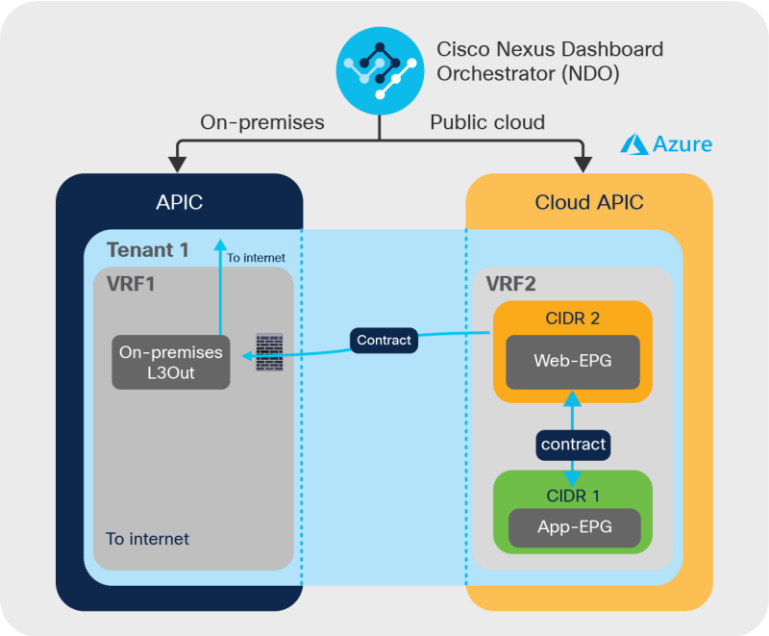

Figure 19.
An example of on-premises L3-Out for cloud endpoints

Cloud-native and on-premises services

Cisco Cloud ACI makes it easy to securely consume services that are either cloud-native or third-party virtual appliance in Microsoft Azure, or on premises.

1. Cloud-native services in Microsoft Azure: Using service-graph concepts similar to those for on-premises ACI, you can seamlessly integrate cloud-native services such as Azure Application Gateway (that is called Application Load Balancer: ALB in cAPIC) and Azure Load Balancer (that is called Network Load Balancer: NLB in cAPIC)[4] with your application stack. This can be fully configured from NDO, and the Cloud APIC will automate end-to-end deployment and configuration of ALB and NLB in the VNet and associate it with endpoints in a given endpoint group.
2. Third-party virtual appliances[5] in Microsoft Azure: Using service-graph concepts similar to those for on-premises ACI, you can seamlessly integrate third-party virtual appliances such as Cisco Adaptive Security Appliance (ASAv), Palo Alto VM-Series Next-Generation Firewall, F5 BIG-IP Virtual Edition, etc., with your application stack. Cloud APIC will automate end-to-end service chaining those third-party firewall and load balancers in the VNet and associate it with endpoints in a given endpoint group.
3. On-premises services: For application tiers that are split across on-premises and Microsoft Azure sites, you can seamlessly insert on-premises services (such as load balancer or firewall). This can be achieved by configuring ACI service graphs from NDO.

Cloud-native service example: Let's say we need a load balancer in front of a group of virtual machines in a cloud site, and the Virtual IP (VIP) is accessed from an external network or within the cloud. As shown in Figure 20, this can be configured as part of a contract using cAPIC service graphs from NDO. cAPIC automatically configures NSGs for the service device interfaces. If the load balancer is cloud-native, such as ALB or NLB, then virtual machine IPs are added, accordingly, as part of the listeners of the VIP on the load balancer.

[4] Azure Load Balancer (NLB in cAPIC) is supported starting from Cisco ACI Release 5.0(2).

[5] Third-party firewall is supported starting from Cisco ACI Release 5.0(2), and third-party load balancer is supported starting from release 5.1(2).

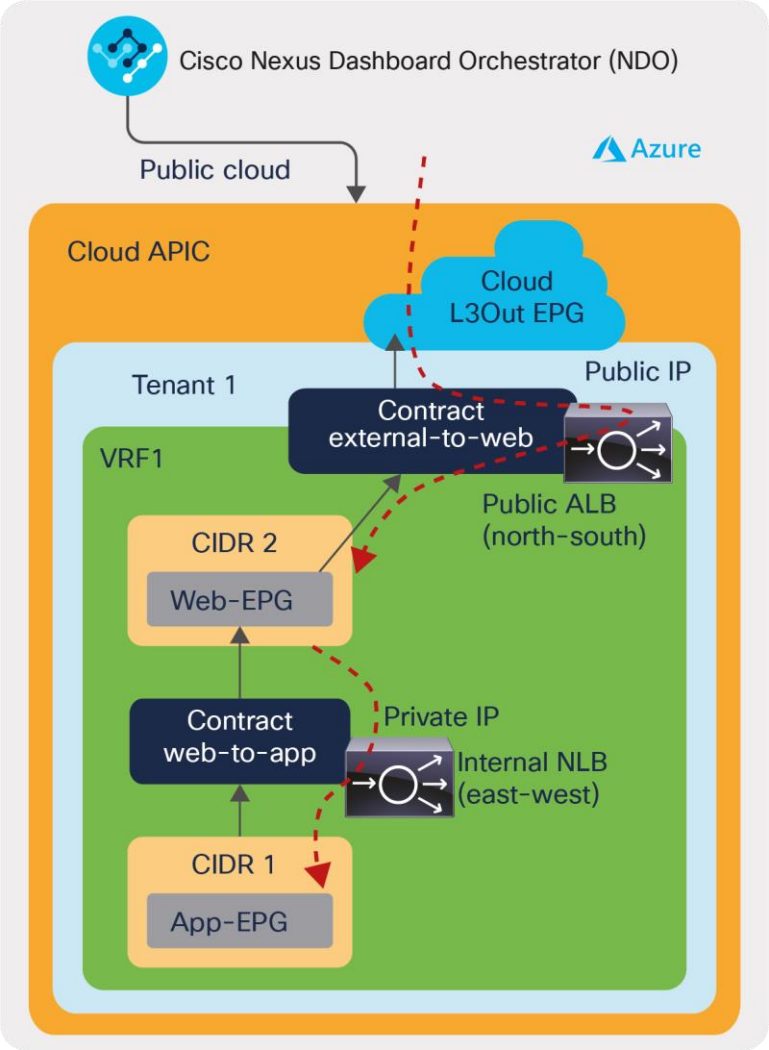

Figure 20.
An example of Azure Application Gateway (ALB), Azure Load Balancer (NLB), or a third-party load balancer for cloud endpoints

Third-party virtual appliances example: Let's say we need to insert firewall for Source Network Address Translation (SNAT) for the traffic from a group of virtual machines in a cloud to an external network. As shown in Figure 21, this can be configured as part of a contract using cAPIC service graphs with redirect from NDO. cAPIC automatically configures NSGs for the third-party firewall interfaces and UDRs to redirect traffic to the firewall if the traffic is from the virtual machines in the cloud EPG to the external network based on the contract. The firewall performs SNAT, and traffic will go to the external network. Because of SNAT on the firewall for the outgoing traffic, the return traffic coming back to the cloud will arrive on the firewall first and then go back to the virtual machine.

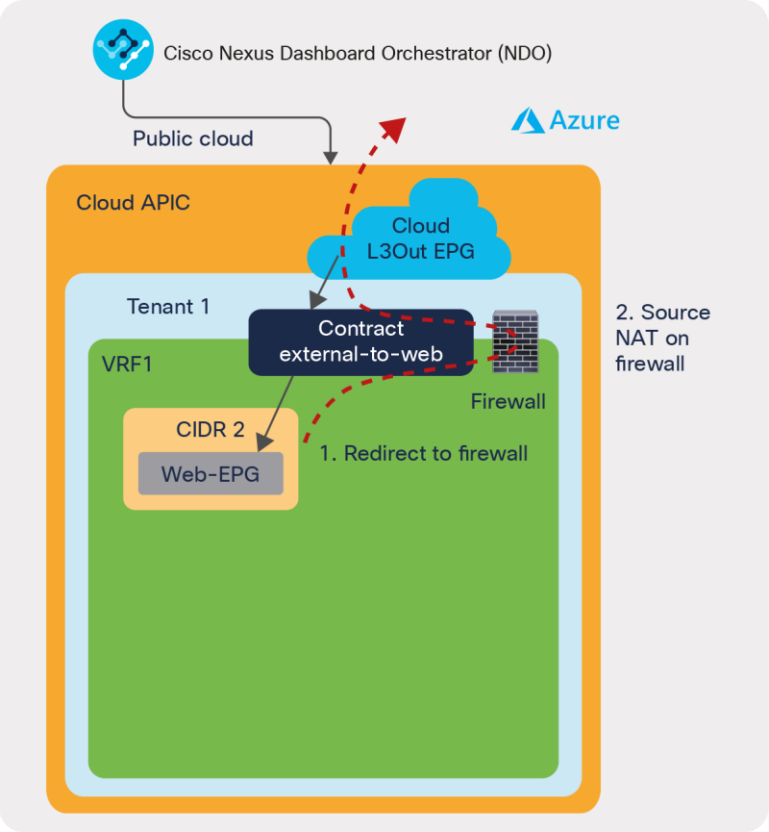

Figure 21.
An example of third-party firewall for cloud endpoints

Multi-node services insertion[6] example: Let's say we need to insert redundant firewalls for Source Network Address Translation (SNAT) for the traffic from a group of virtual machines in a cloud to an external network. As shown in Figure 22, this can be configured as part of a contract using cAPIC multi-node service graphs with redirect from NDO. In this example, the multi-node service graph has NLB and a third-party firewall. NLB is used for traffic load balancing to firewalls. cAPIC automatically configures NSGs for the service device interfaces and UDRs to redirect traffic to the NLB if the traffic is from the virtual machines in the cloud EPG to the external network based on the contract. After the NLB load balances the traffic to one of the firewalls, the firewall performs SNAT, and traffic will go to the external network similar to the previous example. Because of SNAT on the firewall for the outgoing traffic, the return traffic coming back to the cloud will arrive on the firewall first and then go back to the virtual machine without going through the NLB.

[6] Multi-node service graph is supported in ACI Release 5.0(2).

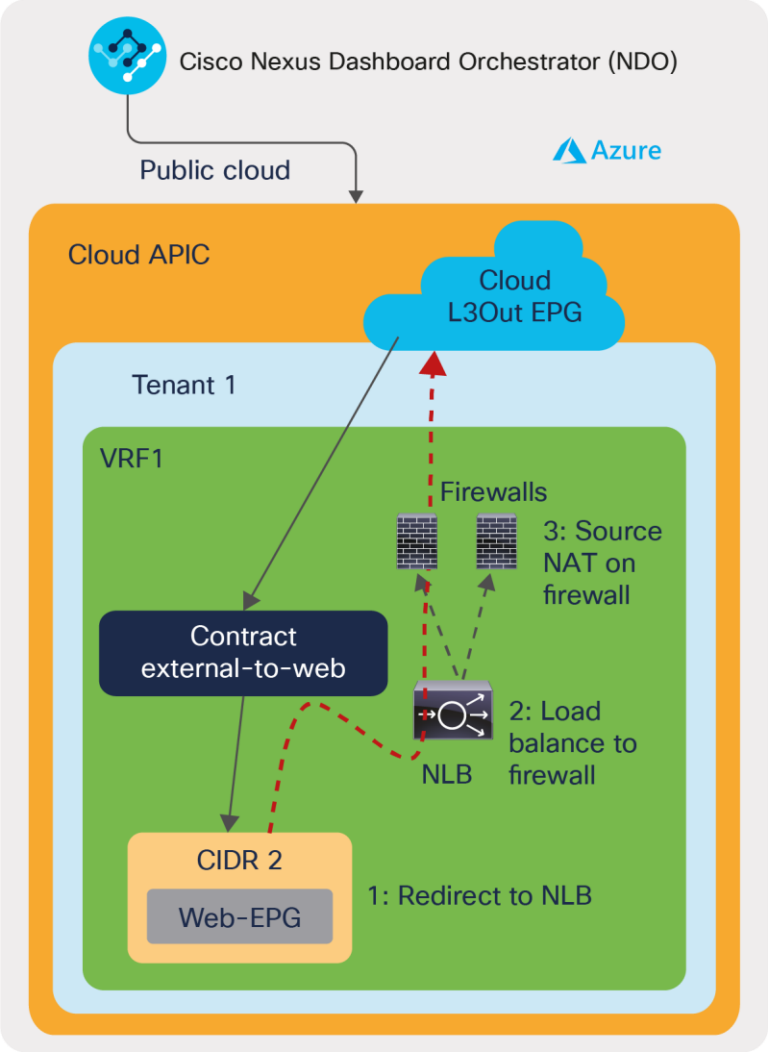

Figure 22.
An example of multi-node service insertion for cloud endpoints

On-premises services example: Let's say we need an east-west firewall between the database tier deployed on premises and a web tier that is stretched across the hybrid cloud environment. As shown in Figure 23, this can be configured with much ease using ACI Service graphs from NDO, and the traffic between the database and web tiers always passes through this firewall irrespective of where the endpoints of the web tier are located across the hybrid cloud.

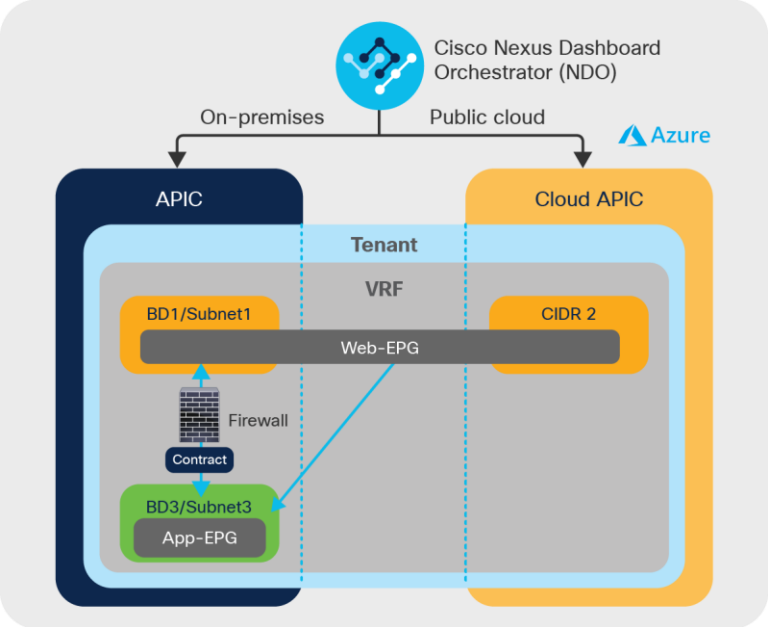

Figure 23.
On-premises service chaining for a stretched application tier

Cloud sites only deployment

Starting with ACI release 4.2, Cisco Cloud ACI supports multi cloud deployments with cloud ACI sites only, i.e., without an on-premises ACI site. The cloud sites can be on any supported public clouds, including Azure. Each cloud site can consist multiple regions on the corresponding public cloud. Figure 24 shows a cloud-sites-only deployment with multiple cloud ACI sites within Azure, each having two regions. Figure 25 shows a cloud-sites-only deployment with cloud ACI sites on different supported public clouds.

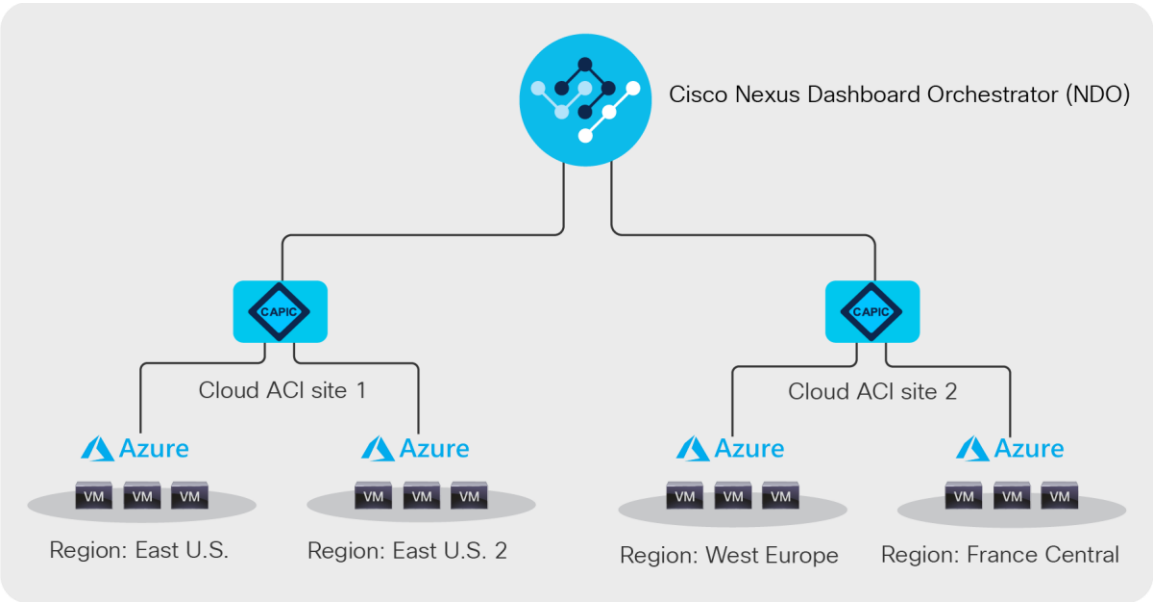

Figure 24.
Multi-cloud ACI sites within Azure

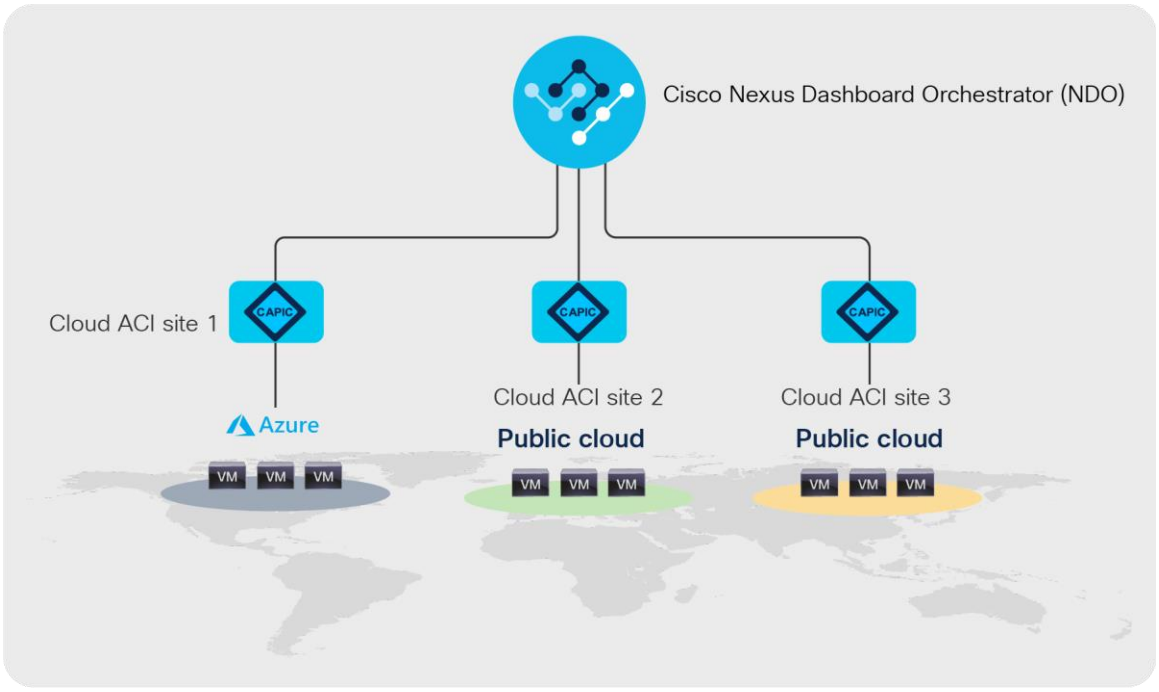

Figure 25.
Multi-cloud ACI sites on different public clouds

In a cloud-sites-only deployment, the network policies are still centrally defined on Cisco Nexus Dashboard Orchestrator (NDO) and distributed to the desired cloud ACI sites. Even without an on-premises ACI site, Cisco Cloud ACI offers the same benefits mentioned earlier in this document. Especially for multi-cloud environments where each cloud has different cloud-native policy constructs, NDO contributes to operational simplicity by representing a single pane of glass for the provisioning of the configuration in different cloud sites. The site-local Cloud APIC translates the network policies to the cloud-native policy model and deploys the cloud native network infrastructure for that site. NDO automatically programs the Cisco CSR 1000v Series routers of the cloud sites to establish the intersite connectivity using IPsec tunnels. It also configures BGP EVPN as the control plane between the CSR 1000v routers to exchange routes over the IPsec tunnel. VXLAN tunnels are then established between CSR 1000v routers for intersite data transfer. Unless on-premises specific feature is included, the use cases mentioned in the previous section are applicable to a cloud-sites-only deployment too.

How to deploy the solution

Deploying Cisco Cloud APIC and the Infra VNet

The Cisco Cloud APIC network virtual appliance is available at Microsoft Azure marketplace and uses the BYOL model for licensing at the time of the writing of this white paper. Readers are encouraged to refer to the ordering guide for any other options.

The first step to deploy Cisco Cloud ACI is to configure the Infra VNet and to launch Cisco Cloud APIC in Microsoft Azure. Those steps can be performed by using a Cisco-provided Microsoft Azure ARM template, which allows you to create automation workflows called stacks that will execute a series of steps on your behalf.

Azure Resource Manager (ARM) uses a JSON template to describe the automation workflows. Cisco provides an ARM template to execute the necessary steps to configure Cisco Cloud ACI integration. The ARM template is available in Microsoft Azure Marketplace, as shown in Figure 26, and can be executed from there. Alternatively, it can be downloaded before deployment through Microsoft Azure Marketplace as show in Figure 27.

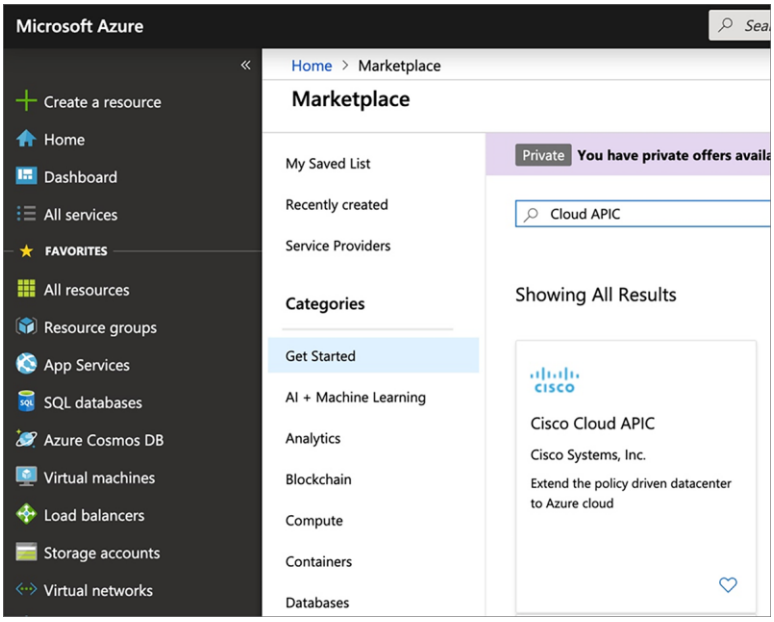

Figure 26.
Cisco Cloud APIC on Microsoft Azure Marketplace

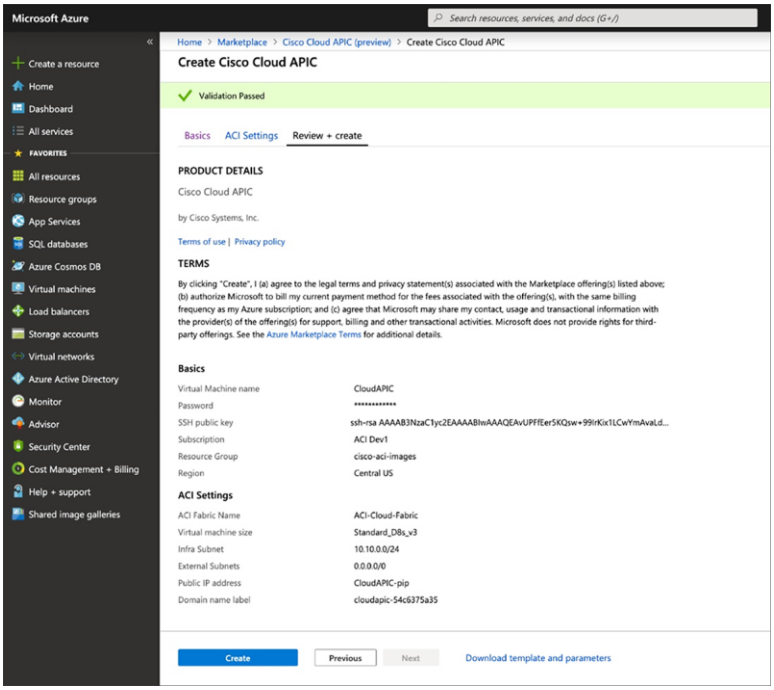

Figure 27.
Download template

Once Cisco Cloud APIC is deployed successfully, your Cisco Cloud APIC becomes accessible via its Web UI and API. As shown in Figure 28, you can find the public IP address of Cisco Cloud APIC. Connect to the Cisco Cloud APIC UI to complete the installation through the getting-started wizard.

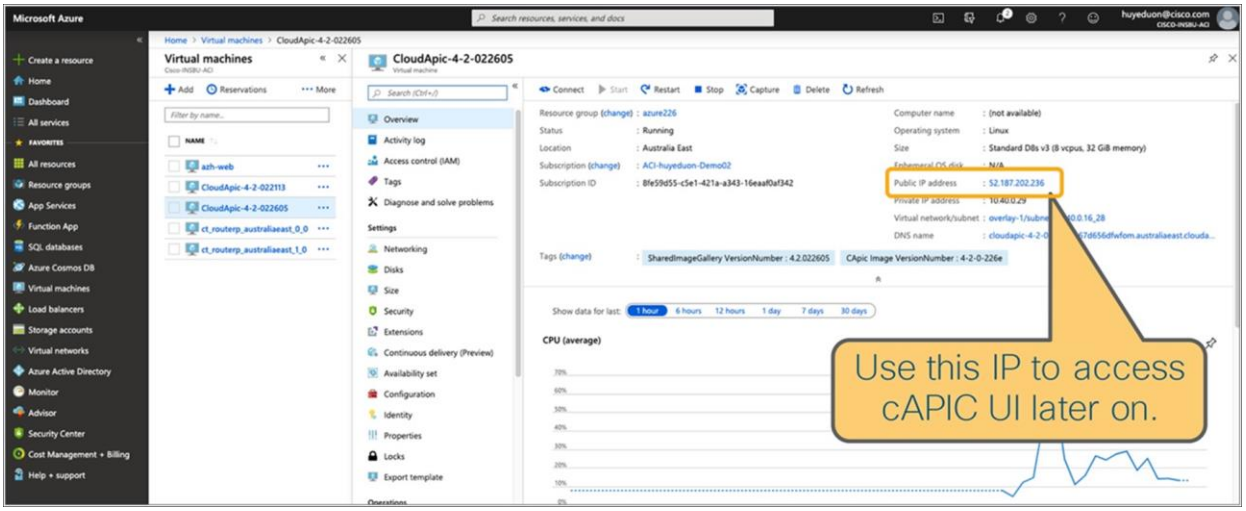

Figure 28.
Cisco Cloud APIC Virtual Machine with Public IP address

Cisco Cloud APIC's First Time Setup wizard

The first time you connect to Cisco Cloud APIC UI, the First Time Setup wizard (shown in Figure 29) automatically kicks off. This wizard helps you configure some of the Cisco Cloud APIC required settings, such as DNS, the TEP pool, the regions to be managed, and IPsec connectivity options. At the end of the First Time Setup wizard, Cisco Cloud APIC configures the Microsoft Azure infrastructure needed to become fully operational, such as the pair of Cisco CSR 1000V Series routers. The provisioning of the Microsoft Azure infrastructure is fully automated and carried out by Cisco Cloud APIC. After this step, you will be able to start deploying your Cisco ACI policy on Microsoft Azure.

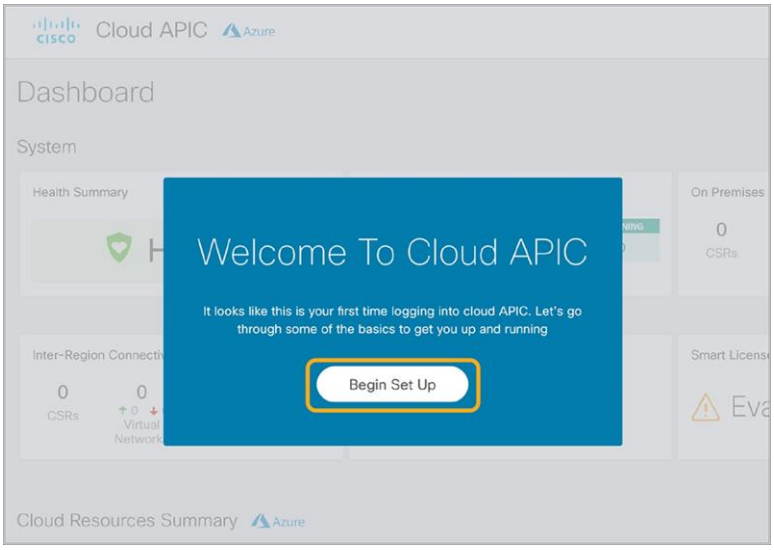

Figure 29.
First time setup Wizard of Cisco Cloud APIC

Registering a Cisco ACI cloud site in NDO

Each Cisco Cloud APIC represents a Cisco ACI site. To extend policy across sites, Cisco ACI uses the Cisco Nexus Dashboard Orchestrator (NDO). As shown in Figure 30, when you register a Cisco Cloud APIC in NDO, it will appear as a new site and will allow you to deploy existing or new schemas to Microsoft Azure. NDO ensures that you specify the required site-specific options, such as subnets and EPG membership classification criteria, which are different for each site.

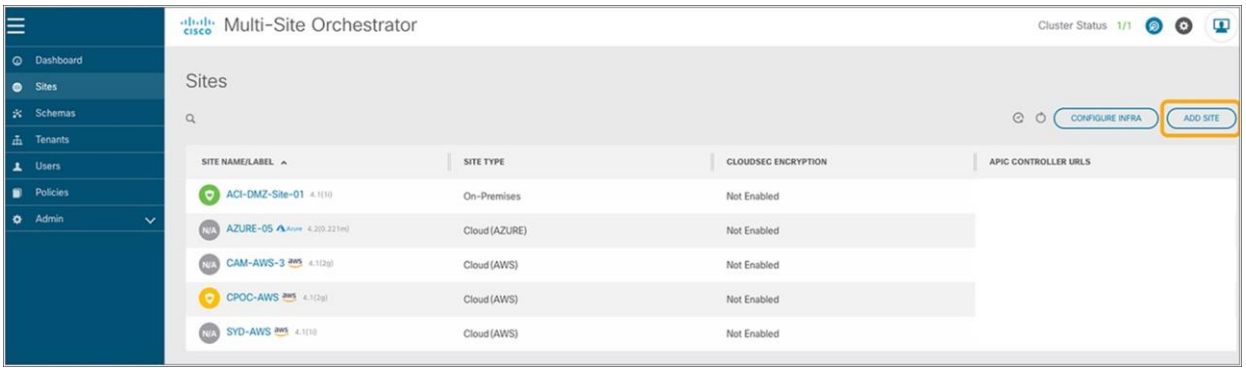

Figure 30.
Register a Cisco ACI Cloud site in NDO

Cisco Cloud APIC lets you create networks on Microsoft Azure using the Cisco ACI object model representation. In the backend, Cisco Cloud APIC translates Cisco ACI objects into Microsoft Azure-native constructs. This means that Cisco Cloud ACI adheres to Microsoft Azure networking specifications. As those differ slightly from what you might be used to with Cisco ACI, they are detailed below.

As shown in Figure 31, in Microsoft Azure, a subnet is bound to a VNet, which itself is bound to a region. A VNet and a subnet span multiple availability zones.

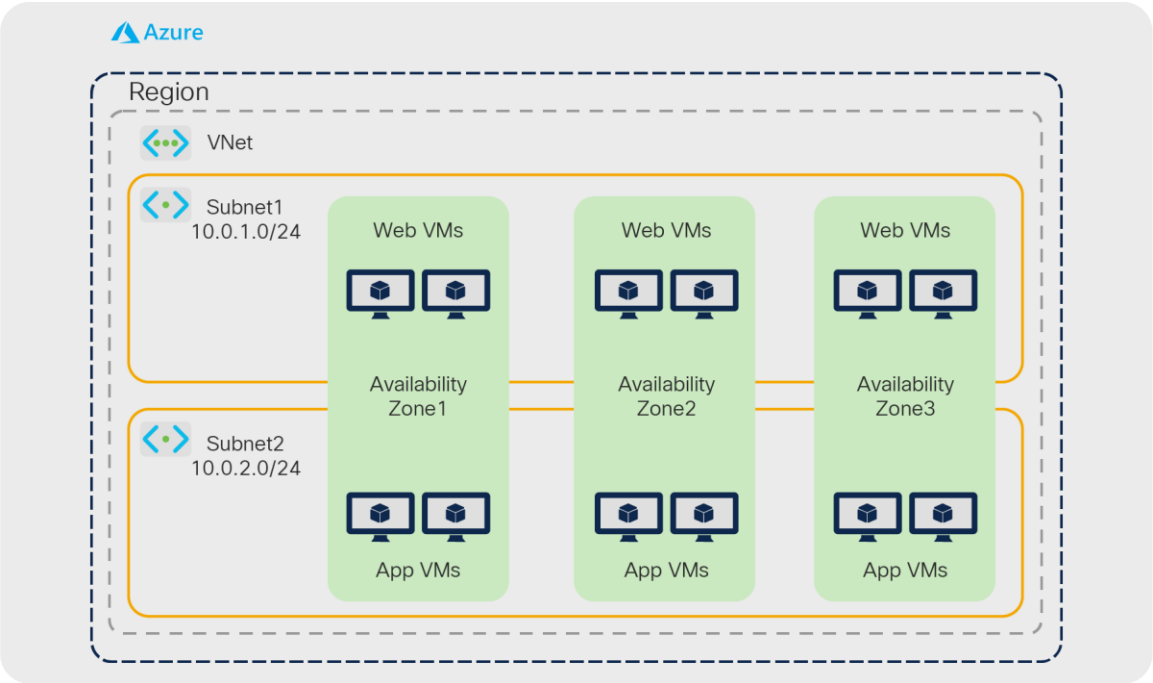

Figure 31.
Microsoft Azure-native network construct

Note: If you have existing vNets and subnets that need to connect to vNets and subnets created through Cisco Cloud APIC, please see "Importing Existing Brownfield Azure Cloud VNets Into Cisco Cloud APIC".

This means that, between Availability Zones, VNets, or regions, the traffic is routed. There is no concept of extending L2 from one VNet to another VNet or from the on-premises site to a VNet in Microsoft Azure. To respect this design philosophy, Cisco Cloud ACI extends on-premises networks using L3 only.

Cisco Cloud APIC also provides a view of the Microsoft Azure-native constructs used to represent the Cisco ACI policy. This allows network administrators to slowly familiarize themselves with Microsoft Azure networking constructs. Figure 32 below demonstrates the native cloud resources view on the Cloud APIC UI. As an example, it shows the display of the provisioned Microsoft Azure VNets in a cloud site.

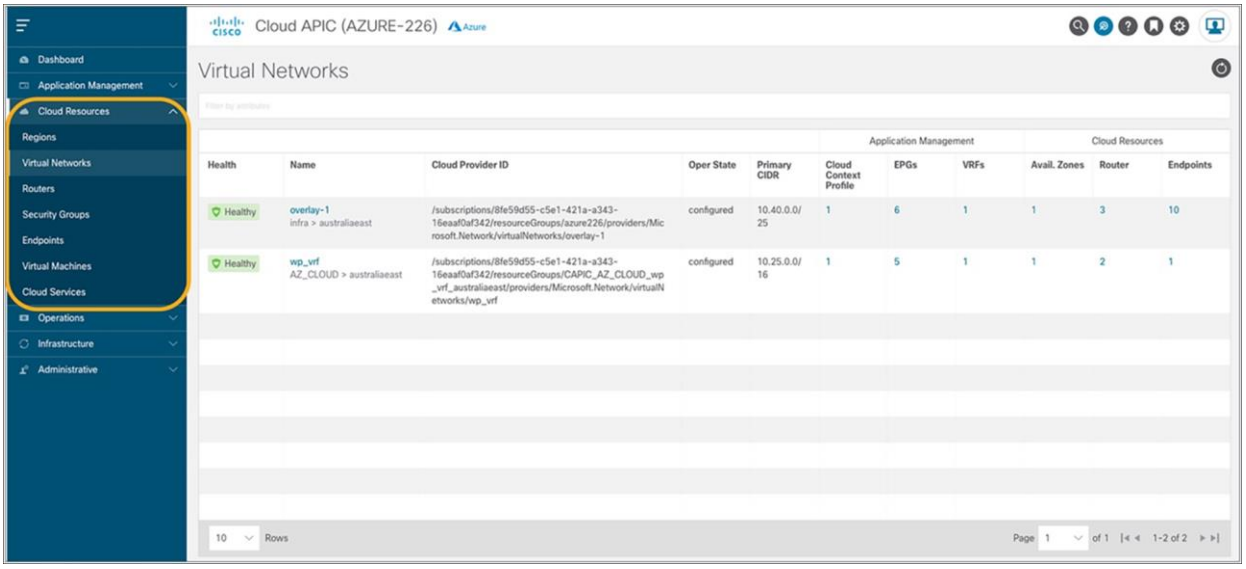

Figure 32.
Native-cloud-resources view on the Cloud APIC User Interface

Deploying a multi-tier application in a hybrid scenario

We use a classic three-tier application as an example in this session. The application consists of a database (DB), App, and Web tier. To deploy it across an on-premises data center and the Microsoft Azure cloud using Cisco Cloud ACI integration, you will need to configure a schema on NDO that represents this policy. As shown in Figure 33, it should contain at least one VRF, one Application Profile, and EPGs as well as contracts between the tiers. For example, the DB tier can be deployed on premises and the Web tier in Microsoft Azure. Or you can use any permutation of this set as you see fit, including (as explained previously) the option to stretch one or more EPGs between the on-premises data center and the cloud.

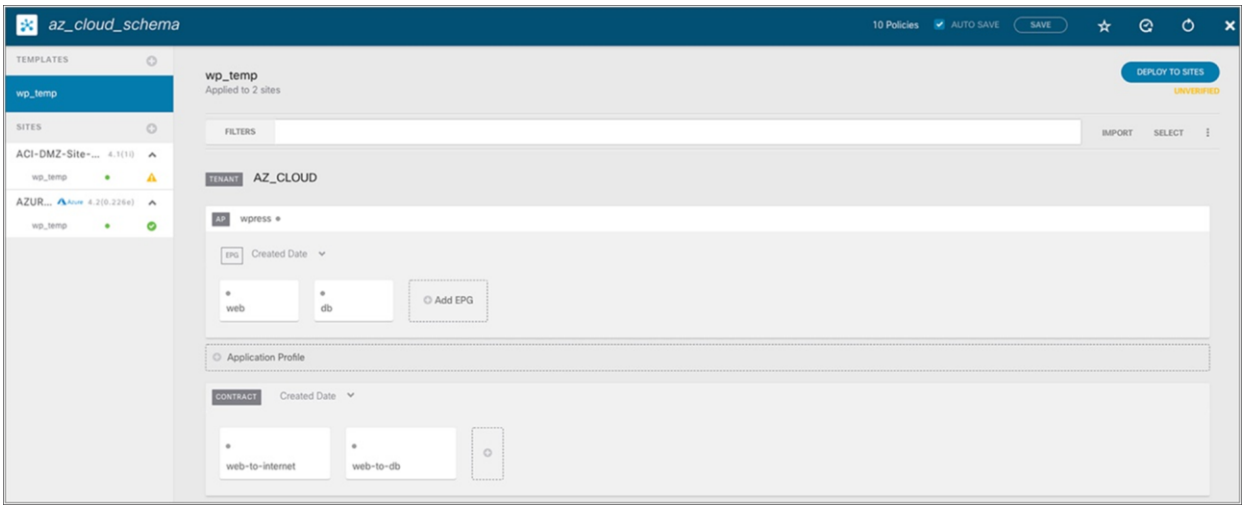

Figure 33.
Application Scheme on NDO

The schema can then be associated with the on-premises site and the Cisco Cloud ACI site. Once the association is made, you then define the subnets to be used for the VRF on Microsoft Azure. Cisco Cloud APIC model associates BD subnets to VRF because, in Microsoft Azure, VRFs are mapped to VNets and subnets are in a VNet. Though a VNet and a subnet span multiple availability zones, there is no concept of extending L2 from one from the on-premises site to a VNet in Microsoft Azure. This means that you need to define a separate subnet for the Web EPG in each site. You will also define the membership criteria for cloud instances to join the Web EPG. As shown in Figure 34, once you are satisfied with the NDO schema and you have completed the required site-specific configuration steps, you can deploy the configuration to both Cisco ACI sites using the NDO one-click deployment button.

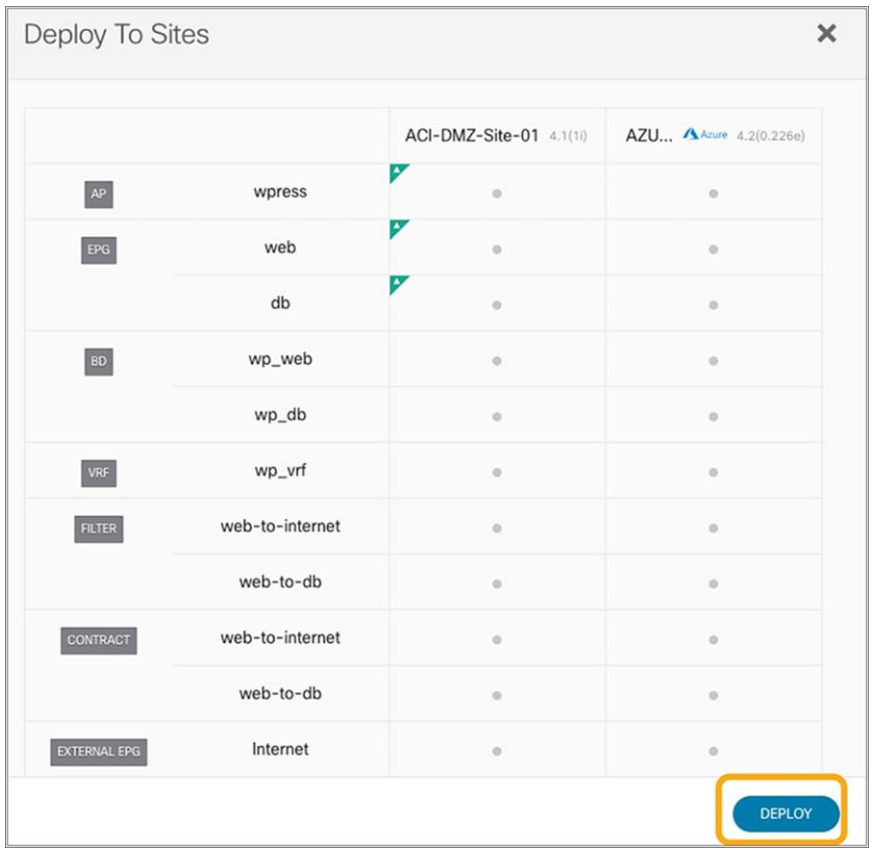

Figure 34.
Deploy application to on-premises and cloud sites in Microsoft Azure

Cisco Cloud ACI ensures that the Microsoft Azure cloud and on-premises ACI are configured appropriately to allow communication between the App EPG and the Web EPG residing on Microsoft Azure, as shown in Figure 35.

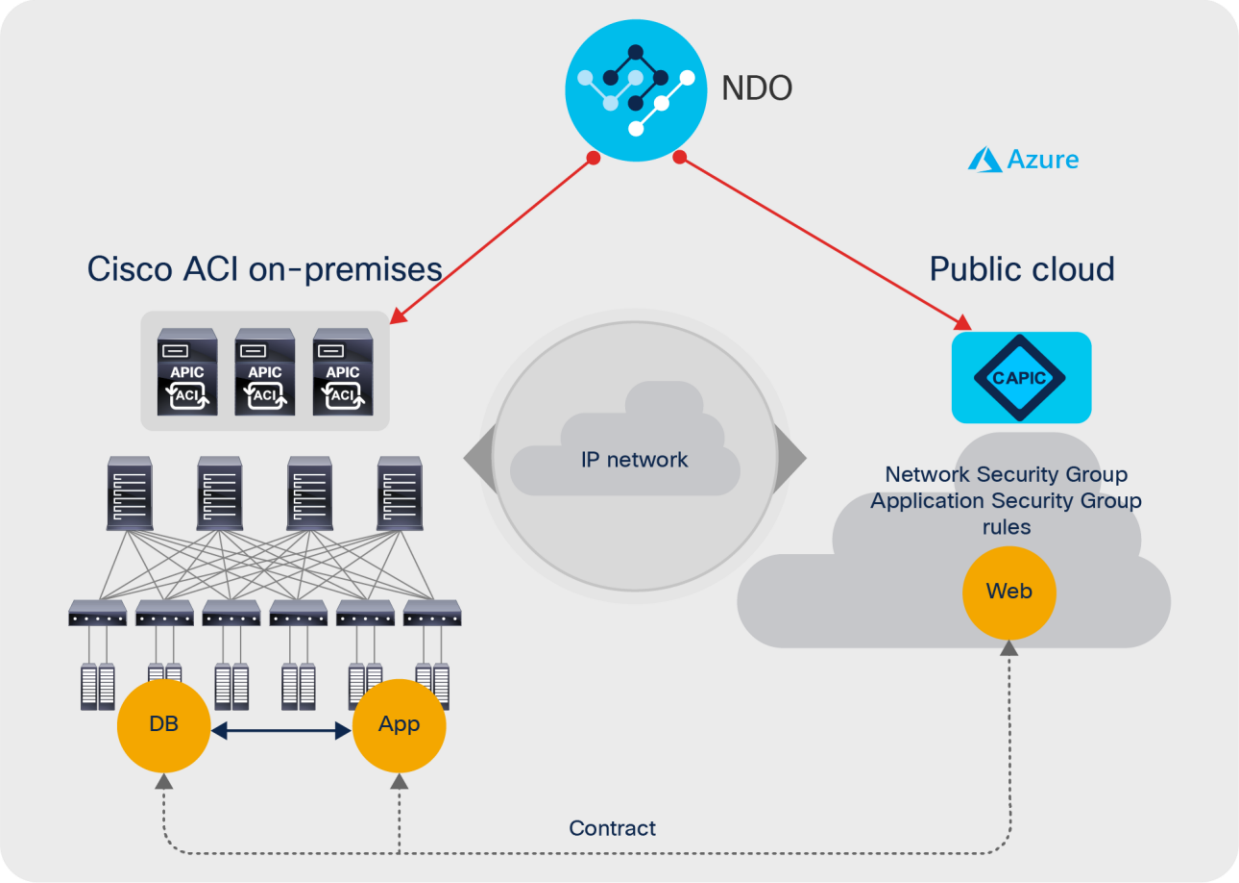

Figure 35.
Three-tier application deployed across on-premises and cloud sites in Microsoft Azure

You can now deploy new Web instances on Microsoft Azure to accommodate your needs.

Summary

The new Cisco Cloud ACI capabilities delivered in Cisco ACI Release 4.1 with AWS and Release 4.2 with Microsoft Azure make it easy for network administrators to quickly tailor the infrastructure to adapt to constantly evolving business requirements. The solution provides ultimate IT agility by greatly facilitating the configuration and day-2 operation of hybrid cloud environments. Cisco Cloud ACI lets you architect complex network topologies and security policies that encompass on-premises locations and public cloud sites. Cross-site orchestration of network connectivity and workload segmentation policies is achieved by Cisco Nexus Dashboard Orchestrator (NDO) working in tandem with Cisco Cloud APIC and on-premises APIC.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA C11-742844-04 01/22