Cisco IT OpenStack ACI Data Center Automation

This is the seventh in a series of white papers that explains how Cisco ACI delivers improved business performance by providing in-depth case studies that cover deployment design, migration to ACI, how contracts enforce network security, the ACI NetApp storage area network deployment, virtualization with AVS, UCS, and VMware, and OpenStack & KVM with AIM multi-cloud management platform. These white papers will enable you to assess the product, plan deployments, and leverage ACI application centric properties to flexibly deploy and manage robust highly scalable data center and network resources.

The audience for this series of white paper includes IT architects, planners, network engineers, developers who will access the system via the APIC API, and APIC fabric administrators.

Version: 1.1, June 2020 – updated with copy edits for clarity.

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

http://www.cisco.com

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883

Table of Contents

Cisco IT OpenStack ACI Orchestration

The Cisco® IT deployment of Cisco Application Centric Infrastructure (ACI) enables its global data center network to deliver the enhanced business value they must have – compelling total cost of ownership, near 100% availability, and agility that includes letting business applications developers directly provision the infrastructure resources they need in a self-service fashion.

OpenStack ACI integration enables Cisco IT to execute on their data center vision.

Cisco IT Data Center Vision

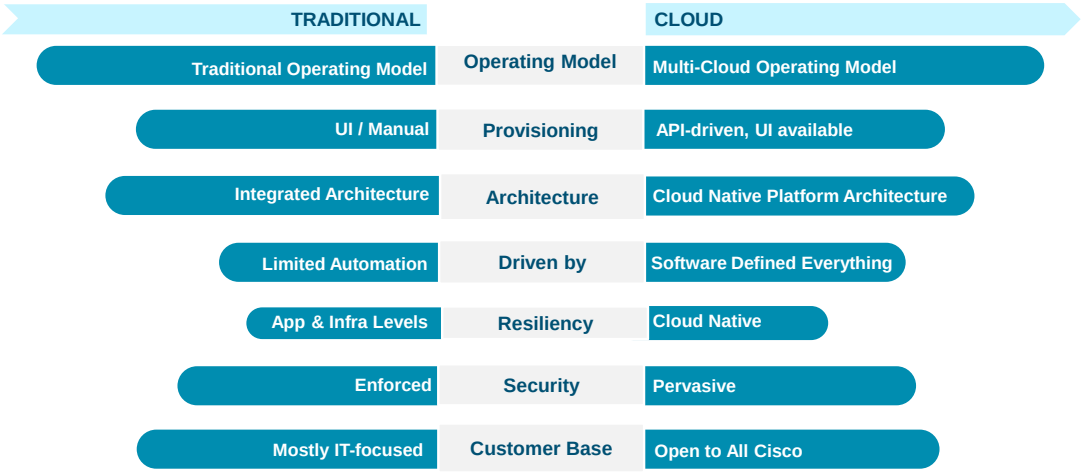

The Cisco IT ACI OpenStack initiative moves their data center operations toward a large scale secure manageable multi-cloud operating model. The OpenStack ACI Integration Module (AIM) enables Cisco IT to add a very large programmable infrastructure as a service (IaaS) private cloud service capability to their ACI fabric. This sets the stage for the enterprise to consume their ACI private cloud and public cloud services in the same manner as those from Amazon AWS, Microsoft Azure, or Google GCP but with greater automation, security, ease of use, and cost-savings.

Cisco IT supports 141,000 employees (71,000 regular employees and 70,000 contractors) in 583 offices across more than 100 countries. ACI enables Cisco IT to achieve its design goals of supporting any workload anywhere managing DC resource pools within virtual boundaries, while maintaining near zero application downtime with enhanced security.

For example, Cisco IT can take an entire UCS compute domain out of service for hardware swaps, upgrades, or maintenance, during normal business hours with zero application down time and put it back in service within 2 hours. The standard data center (DC) has an ACI fabric with an APIC cluster controller, spine switches, and multiple pairs of leaf switches for endpoint connectivity.

Cisco IT Standard ACI Data Center

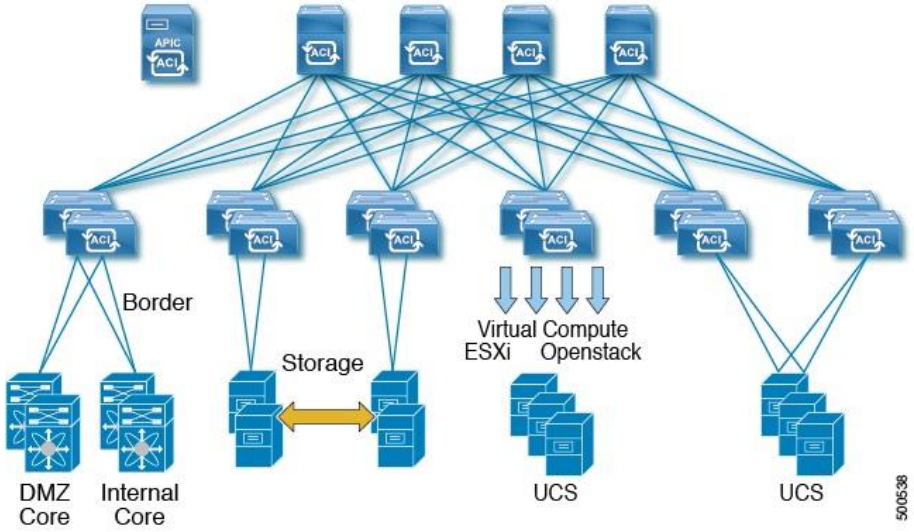

As its data centers and multi-cloud deployments grow, quick and agile application deployment becomes increasingly challenging.

Cisco IT Workload Distribution

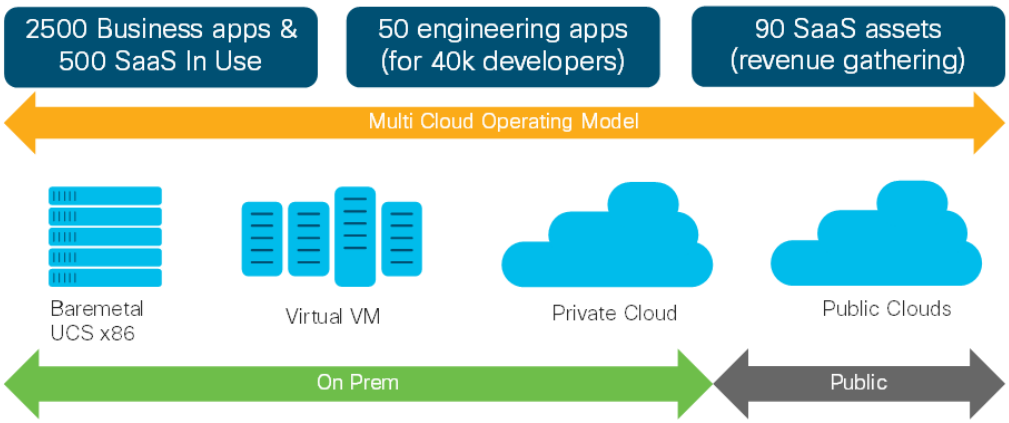

ACI enables Cisco IT to use a common software defined application-aware policy-based operating model across their entire physical and virtual environments.

Deploying a common software defined operating model across their entire DC environments enables Cisco IT to deliver enhanced business value to the enterprise.

Cisco IT New Cloud Platform Benefits

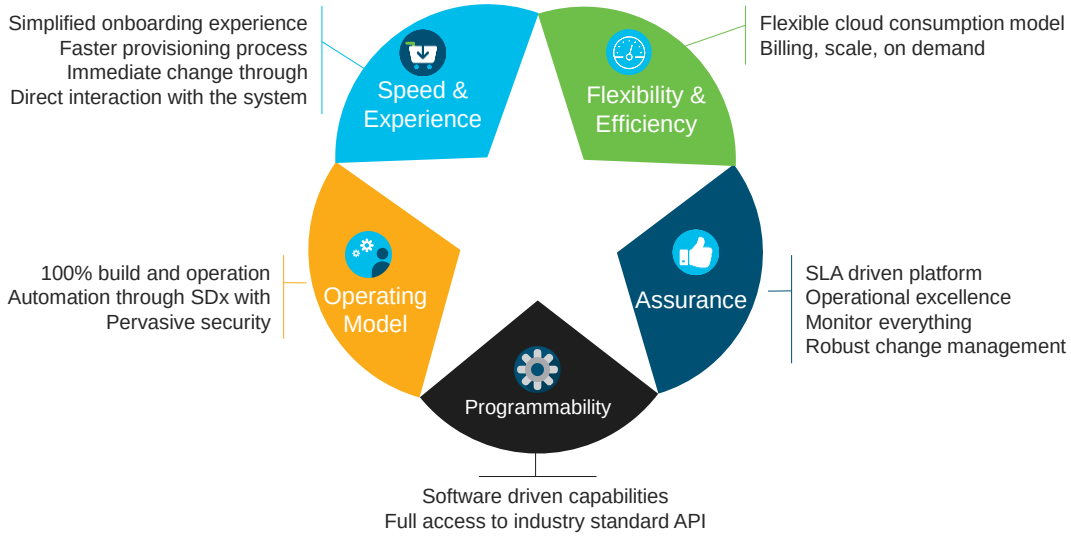

Cisco IT expects the OpenStack ACI project will enable the following improvements in business performance:

- Up to 30% cost reduction to clients for their VMs
- 99.99% SLO Availability
- 4:1 consolidation of cloud offerings with pay as you go and scale on demand
- Fully supported programmability through APIs
- Improved customer experience - for example, 75% fewer clicks for onboarding

As Ken Schroeder, Rob Douglas, and Desh Shukla from the Cisco IT Next Generation Cloud Services team explain, "One of the unique design opportunities for us is to use ACI as the foundation for an OpenStack cloud. This lets applications developers directly consume data center resources wherever they may reside according to their application requirements." This white paper details how Cisco IT designed its OpenStack ACI deployment to do just that.

Cisco IT Cloud

OpenStack ACI orchestration is central to the strategic new Cisco IT Cloud. The new cloud embraces the public cloud provider approach. All of Cisco can self-service consume offerings that are differentiated from traditional enterprise IT.

OpenStack is an open source infrastructure as a service (IaaS) initiative for creating and managing large groups of virtual private servers in a data center. Cisco OpenStack plugins enable OpenStack instances to leverage the ACI fabric as a software defined networking (SDN) platform. This enables dynamic creation of networking constructs that are driven directly from OpenStack, while providing additional visibility and control through the ACI APIC controller. OpenStack supports interoperability between cloud services and allows businesses to build AWS-like cloud services in their own data centers.

For client onboarding and account management, Cisco IT developed the Multicloud Management Platform (MCMP). Cisco business users can choose to purchase and deploy services based on capabilities and locations that best meet their needs. A thin top-level portal provides an onboarding experience for the new Cisco IT private cloud as well as for 3rd party cloud providers. Billing and account usage are tracked from this portal.

Cisco IT Multi-Cloud Management Platform

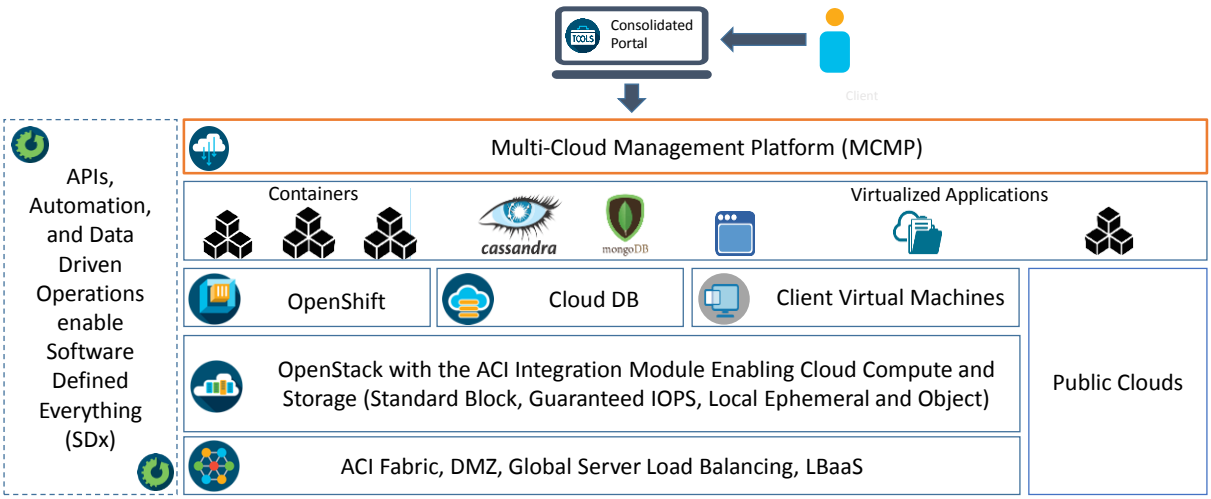

Cisco IT cloud services are built with a software defined everything approach with ACI as the foundational layer and OpenStack providing core abstractions of network resources, Virtual Machine, and storage capabilities.

The industry standardized APIs provided by OpenStack enable offering clients additional consumable platform offerings, in addition to many of the same benefits offered by external cloud providers, including:

- Budget charges based on actual resource use
- Strong security - infrastructure hosted on Cisco solutions in Cisco data centers
- On-demand scaling of infrastructure resources
- Different IT managed service offerings
- Service-level agreements (SLAs) and robust change management processes for high levels of cloud service delivery

Cisco ACI Anywhere

The Cisco ACI platform provides a policy-based software defined solution for managing the network and security of on-site, remote, and public cloud workloads.

Cisco ACI Anywhere – Any Workload, Any Location, Any Cloud

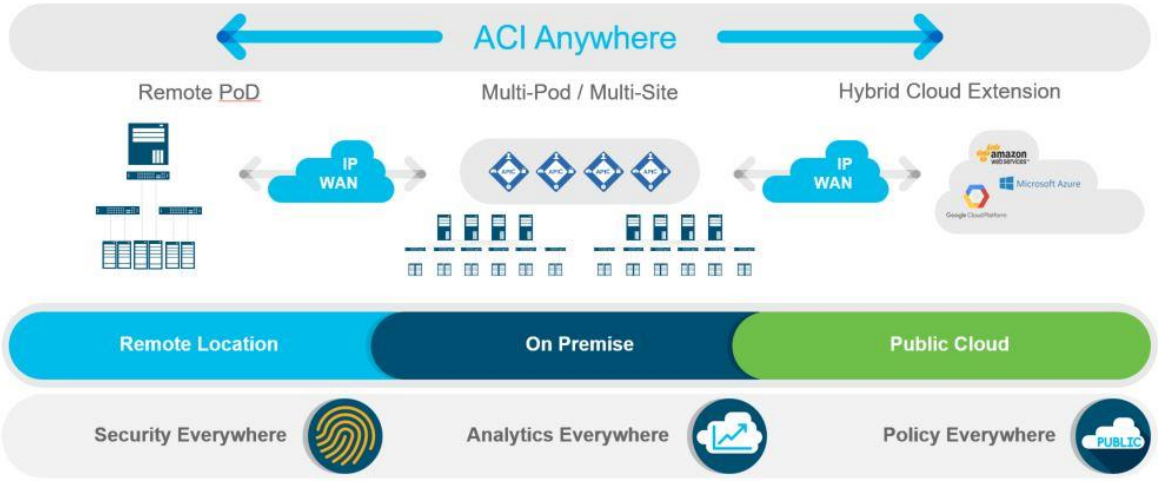

Cisco makes key attributes of the ACI solution, such as unified security policy, concurrent multiple virtual machine management domains with multiple hypervisors, single-pane-of-glass management, and visibility, available in public cloud environments. Customers have the flexibility to run applications across their own private clouds, as well as public clouds of their choice, while maintaining consistent network policies across their entire multicloud domain.

Relevant Cisco ACI Constructs

The Cisco ACI policy model is the software defined basis for managing the entire ACI fabric, including the infrastructure, authentication, security, services, applications, and diagnostics. Logical constructs in the policy model define how the fabric meets the needs of any of the functions of the fabric.

Tenants, Application Profiles, Bridge Domains, EPGs and Contracts

Tenant policies are the core ACI construct that enable application deployment agility. Tenants can map to logical segmentation isolation constructs of both private and public cloud services. Tenants can be isolated from one another or can share resources.

The endpoint group (EPG) is the most important object in the ACI policy model. The ACI fabric can contain the following types of EPGs:

- Application endpoint group
- Layer 2 external outside network instance endpoint group
- Layer 3 external outside network instance endpoint group
- Management access (out-of-band or in-band) endpoint group

Endpoints are devices connected directly or indirectly to the ACI fabric. EPGs are fully decoupled from the physical and logical topology. Endpoint examples include servers, virtual machines, network-attached storage, external Layer 2 or Layer 3 network connections, or clients on the Internet. In general, policies apply to EPGs, not to individual endpoints. An administrator can statically configure an EPG, or automated systems such as VMware vCenter, OpenStack, Cisco Cloud Center or Microsoft Azure Pack can dynamically configure EPGs.

The following figure provides an overview of the Cisco IT ACI tenant implementation.

Cisco IT ACI Tenant Design

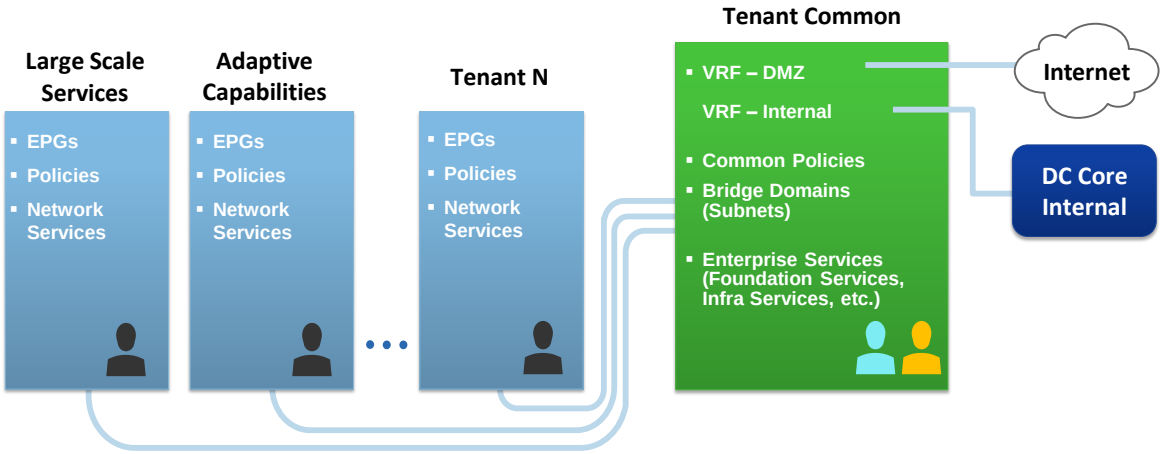

ACI contract security policies regulate network flows within the isolation of tenant application profiles. Cisco IT OpenStack ACI AIM plugin enables application developers to deploy VMs with the memory and CPU resources they need that are isolated within a tenant. If additional network access is needed (Web server, network protocols/ports, etc.,) then the network services team handles such requests.

Within a tenant, bridge domains define a unique Layer 2 MAC address space. A bridge domain must be linked to a context (VRF) and have at least one subnet that is associated with it. A VRF can consist of multiple subnets. Subnets in bridge domains can be *public* (exported to routed connections), *private* (used only within the tenant) or *shared* across VRFs and across tenants.

In the ACI allowed list security model, all communication is blocked by default between EPGs. Contracts explicitly specify what types of communication is allowed between EPGs, and the EPGs specify the source and destination of the communications. Endpoints in EPG 1 can communicate with endpoints in EPG 2 and vice versa, only if a contract allows it. A contract consists of filters and actions. An "allow all" filter would allow any kind of traffic to flow between the EPGs using such a contract. The filter allows traffic based on layer 3 and layer 4 information. For example, a web server might provide the contract filter that only allows http and https traffic.

Virtual Machine Manager Domains

ACI VM networking enables consistent enforcement of policies across both virtual and physical workloads. ACI virtual machine manager (VMM) domains enable an administrator to configure connectivity policies for virtual machine controllers. Multiple VM hypervisors from different vendors can run concurrently on the ACI fabric and can interoperate, regardless of which switches are associated with the ACI VMM domains, and where the compute pods connect to the ACI fabric.

A single ACI leaf switch can connect to multiple types of hypervisor domains. For example, both VMware VMs and OpenStack/KVM VMs can all run on a single UCS compute pod.

ACI Multiple VM Controller Integration

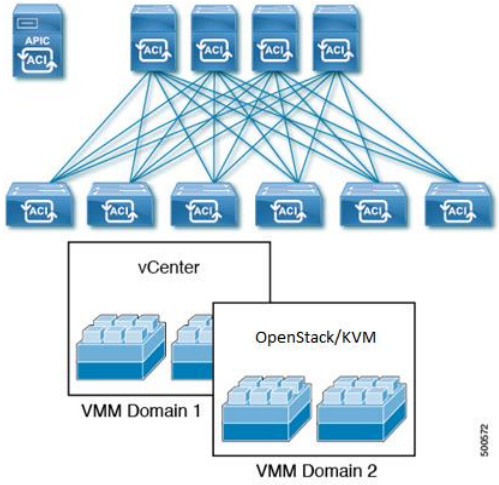

VMM domains provide support for the following:

- Multiple tenants within the ACI fabric.
- Automated static or dynamic VLAN or VXLAN allocations from specified VLAN/VXLAN pools.
- The APIC automates the configuration of a hypervisor virtual switch.
- Micro segmentation that, when combined with intra-EPG isolation or intra-EPG contract for bare metal and VM endpoints, provides policy driven automated complete endpoint isolation within application tiers.

Virtual machine management connectivity to a hypervisor is an example of a

configuration that uses a dynamic EPG. Once the virtual machine management domain is configured in the fabric, the hypervisor automatically triggers the dynamic configuration of EPGs in the ACI fabric leaf switches. This enables virtual machine endpoints to automatically start up, move, and shut down as needed.

The ACI fabric associates tenant application profile EPGs to VMM domains, either automatically by orchestration components such as OpenStack or Microsoft Azure, or by an APIC administrator creating such configurations.

OpenStack ACI Integration Overview

Also known and the Unified Plugin, the ACI Integration Module (AIM) enhances OpenStack deployments by delivering unique critical capabilities that are not available or not as robust in OpenStack, such as distributed L3 routing, distributed NAT, optimized DHCP and metadata services, and clear visibility into the fabric underlay by automatically correlating the OpenStack information with the ACI network topology.

Why Cisco ACI and OpenStack?

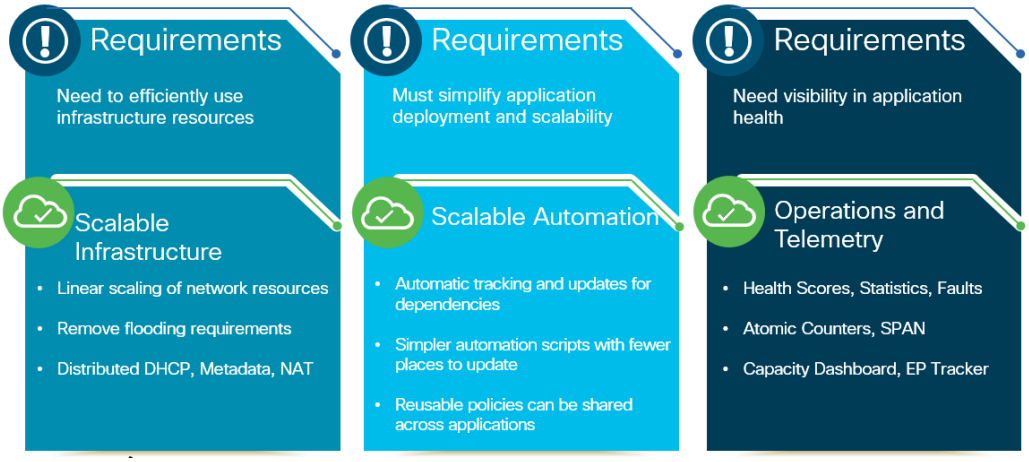

ACI gracefully integrates all physical and virtual networking requirements, easily handling whatever OpenStack requires. Neutron enables providing self-serviced network functions, including network/subnet/router creation and floating IPs on top of the ACI fabric. The ACI distributed Layer 3 services are more robust that what OpenStack provides. Integrating OpenStack with ACI enables you to take advantage of ACI high availability distributed L3 routing, distributed NAT to implement floating IP and SNAT, along with optimized distributed DHCP and metadata services.

ACI offers native service-chaining capabilities to transparently insert or remove services between two endpoints. ACI policy-based routing can automatically perform service insertion of appliances such as firewalls, load balancers, and other L4-L7 devices. Tenant and cloud administrators can automate provisioning complex security policy requirements in the cloud self-service application deployment model.

OpenStack has no visibility into the network underlay but ACI does. ACI offers a fully managed underlay network for the Cisco ACI fabric through the APIC. It provides the capability to connect physical servers and multiple hypervisors to overlay networks. The APIC is the central point for management and provisioning that offers real-time visibility and telemetry information with detailed data about the performance of individual endpoint groups (EPGs) and tenants in the network. This information includes real-time health metrics for the physical and virtual networks. This telemetry information helps reduce troubleshooting time significantly. Faults can be traced from virtual to physical connections down to the virtual machine. Cisco ACI networks can be debugged efficiently, and actions can be orchestrated based on the events triggered by the faults.

The Cisco ACI Integration Module Plugin

The AIM plugin relies on the Neutron OpenStack project to provide "networking as a service" between interface devices (e.g., vNICs) managed by other OpenStack services. The plugin leverages the APIC Rest API to orchestrate the ACI network infrastructure. ACI objects created by the plugin map to Neutron constructs used in OpenStack tenant network activities.

The figure below illustrates the interaction of the OpenStack controller node running the ACI Integration Module, the Cisco ACI fabric, the OpFlex proxy, and the Open vSwitch (OVS) running on the compute host.

The ACI Fabric, and the OpenStack Controller Node with the ACI Integration Module

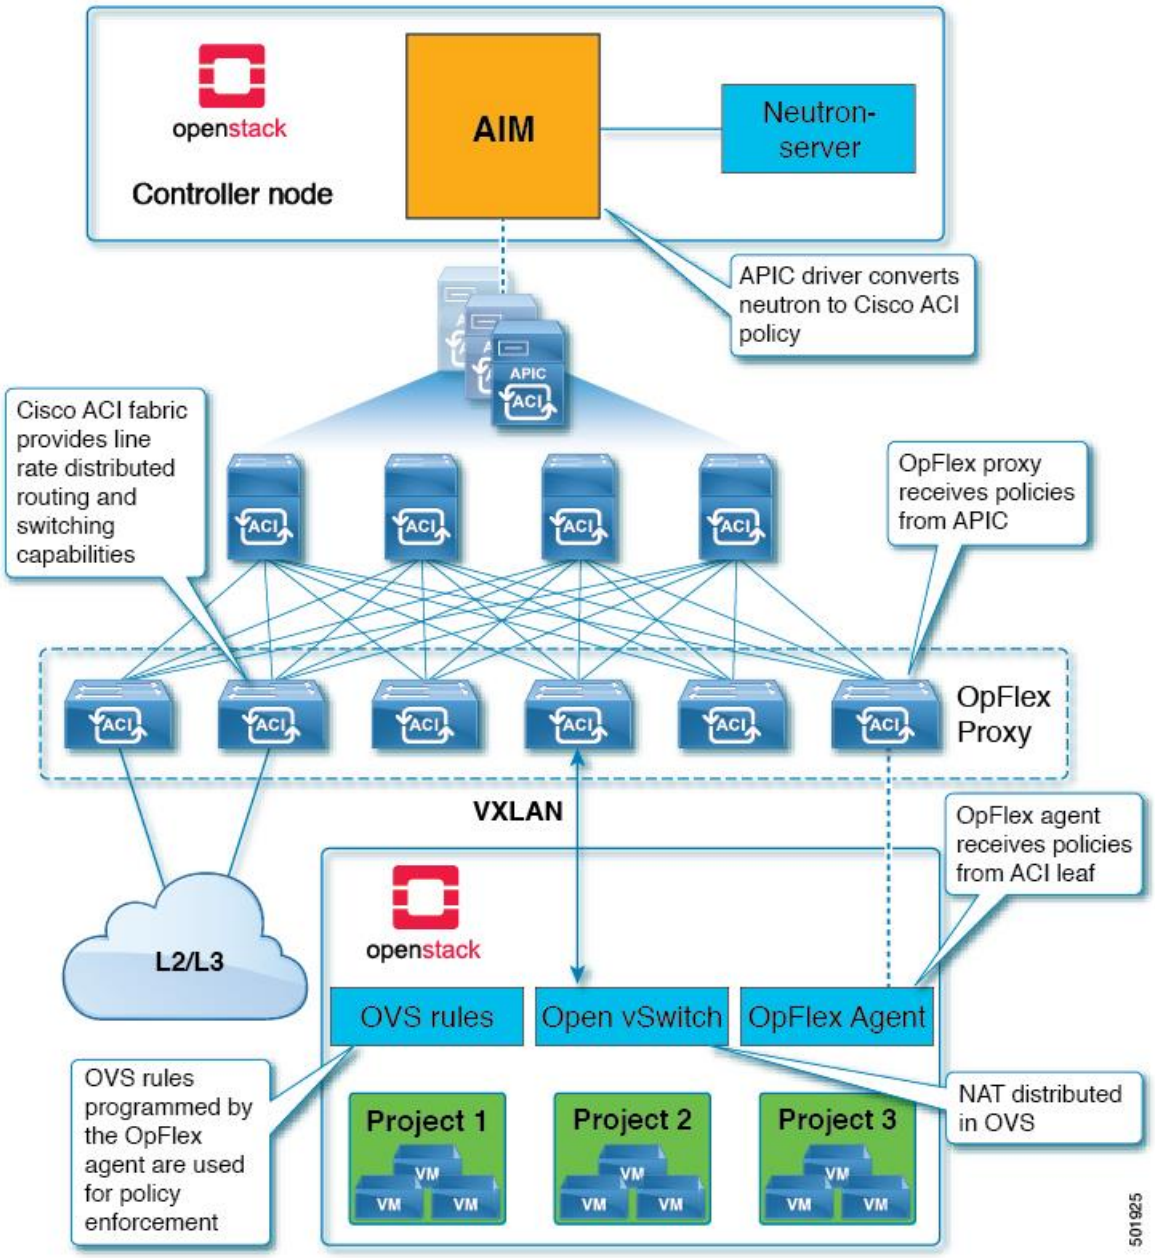

The AIM plugin provides the following:

- **Distributed Network Address Translation (NAT) for external networks**: The AIM package includes the OpFlex ML2 driver, which supports external

networks by distributing OpenStack Source NAT (SNAT) and floating IP functions across the Open vSwitch instances of the compute hosts. Packets destined for IP addresses not defined in the private OpenStack tenant space are automatically translated before they egress a computing host. The translated packets are routed to the external routed network defined in OpenStack and by the APIC. The AIM plugin automates the provisioning of ACI constructs that enable these distributed NAT services.

- **Distributed Layer 3 forwarding**: A combination of the distributed Layer 3 forwarding in the Cisco ACI fabric and local forwarding in the computing node replaces the Neutron Layer 3 agent. If two virtual machines on the same compute node connect to the same OpenStack tenant router, Layer 3 traffic between them is forwarded locally.
- **Distributed Dynamic Host Configuration Protocol (DHCP)**: The Neutron approach centralizes DHCP services on Neutron servers. The OpFlex ML2 driver approach enables using the agent-ovs service to distribute DHCP services. This distributed approach confines DHCP discovery, offer, request, and acknowledge (DORA) traffic local to the host and helps ensure reliable allocation of IP addresses to VM instances.
- **Distributed Metadata proxy**: OpenStack VMs can receive instance-specific information such as instance IDs, host names, and Secure Shell (SSH) keys from the Nova metadata service. In the Neutron approach, a Neutron service acts as a proxy on behalf of OpenStack VMs to receive this metadata. The OpFlex ML2 driver included with the AIM package distributes this proxy function to each compute host.

Depending on the options you use, OpenStack could run on multiple hosts that provide an internal API network, a storage network, a provisioning network, or an external network. AIM continuously synchronizes with the APIC. OpenStack controllers require write privilege access to the APIC.

The Group-based Policy (GBP) abstractions for OpenStack provide an intent-driven declarative policy model that presents simplified application-oriented interfaces to the user. The AIM ML2 drivers coexist with the AIM Group Based Policy (GBP) policy driver, providing full simultaneous support for both Neutron and GBP APIs within the same OpenStack deployment, including sharing of resources between Neutron and GBP workflows where appropriate. OpenStack Group-Based Policies map into the Neutron API and then into the AIM plugin. Neutron APIs map to AIM directly.

OpenStack controller nodes running the Neutron server pass OpenStack configurations through the Cisco AIM plugin, which uses using REST API calls to configure the ACI Application Policy Infrastructure Controller (APIC).

The OpFlex™ Modular Layer 2 (ML2) framework in OpenStack enables integration of networking services based on type drivers and mechanism drivers. Common networking type drivers include local, flat, VLAN, and Virtual Extensible LAN (VXLAN). The OpFlex network type through ML2 enables packet encapsulation of either VXLAN or VLAN on the host as defined in the OpFlex configuration. A mechanism driver communicates networking requirements from Neutron servers to the APIC. The APIC mechanism driver translates Neutron networking elements such as a network (segment), subnet, router, or external network into APIC constructs within the Cisco ACI policy model.

OpFlex ML2 uses the hypervisor native Open vSwitch (OVS) package and leverages local software agents (opflex-agent and agent-ovs) on each OpenStack compute host to communicate with Neutron servers and to program OVS. The OpFlex proxy running on the Cisco ACI leaf switches exchange policy information with the opflex-agent instance in each compute host, effectively extending the Cisco ACI switch fabric and policy model into the OVS virtual switch. This extension results in a cohesive system that can apply networking policy anywhere in the combined virtual and physical switching fabrics, starting from the virtual port at which a virtual machine instance attaches to the network.

OpenStack Nova Compute Hosts: OpFlex Agent Architecture

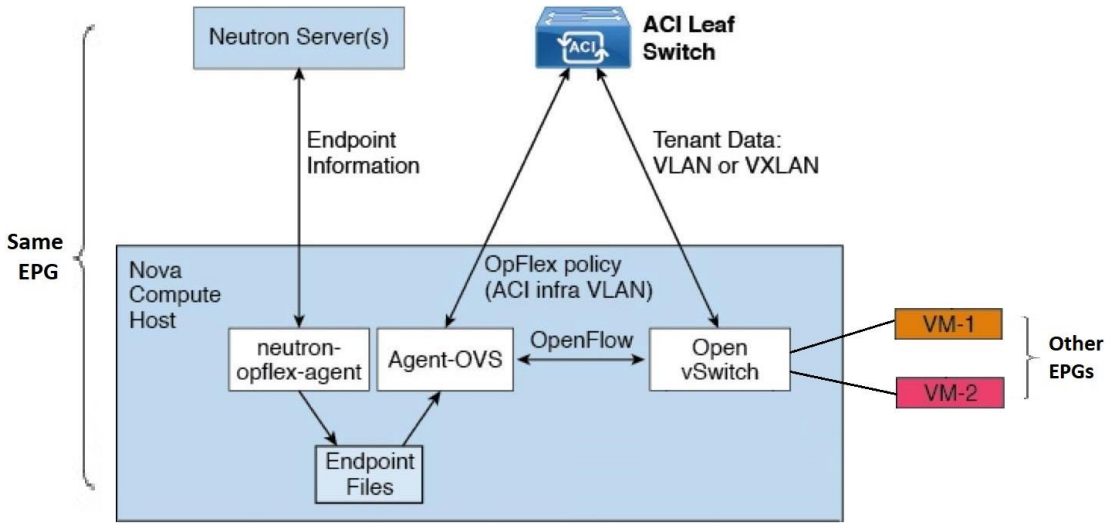

Compute nodes use the OpFlex agent and OVS to connect the VM interface with the network and perform policy enforcement. The neutron-opflex-agent service on the Nova

compute node receives information about OpenStack endpoints from the ML2 driver software on the Neutron server. This information is stored in files on the Nova compute node. The agent-ovs service uses the endpoint information to resolve policy for the endpoints through the OpFlex proxy on the Cisco ACI leaf switch. The agent-ovs service uses OpenFlow to program policy on the OVS for policies that can be enforced locally. Nonlocal policies are enforced on the upstream ACI leaf switch.

The OpenStack AIM plugin enables fully automated orchestration of virtual machines.

OpenStack AIM Plugin VM Orchestration Flow

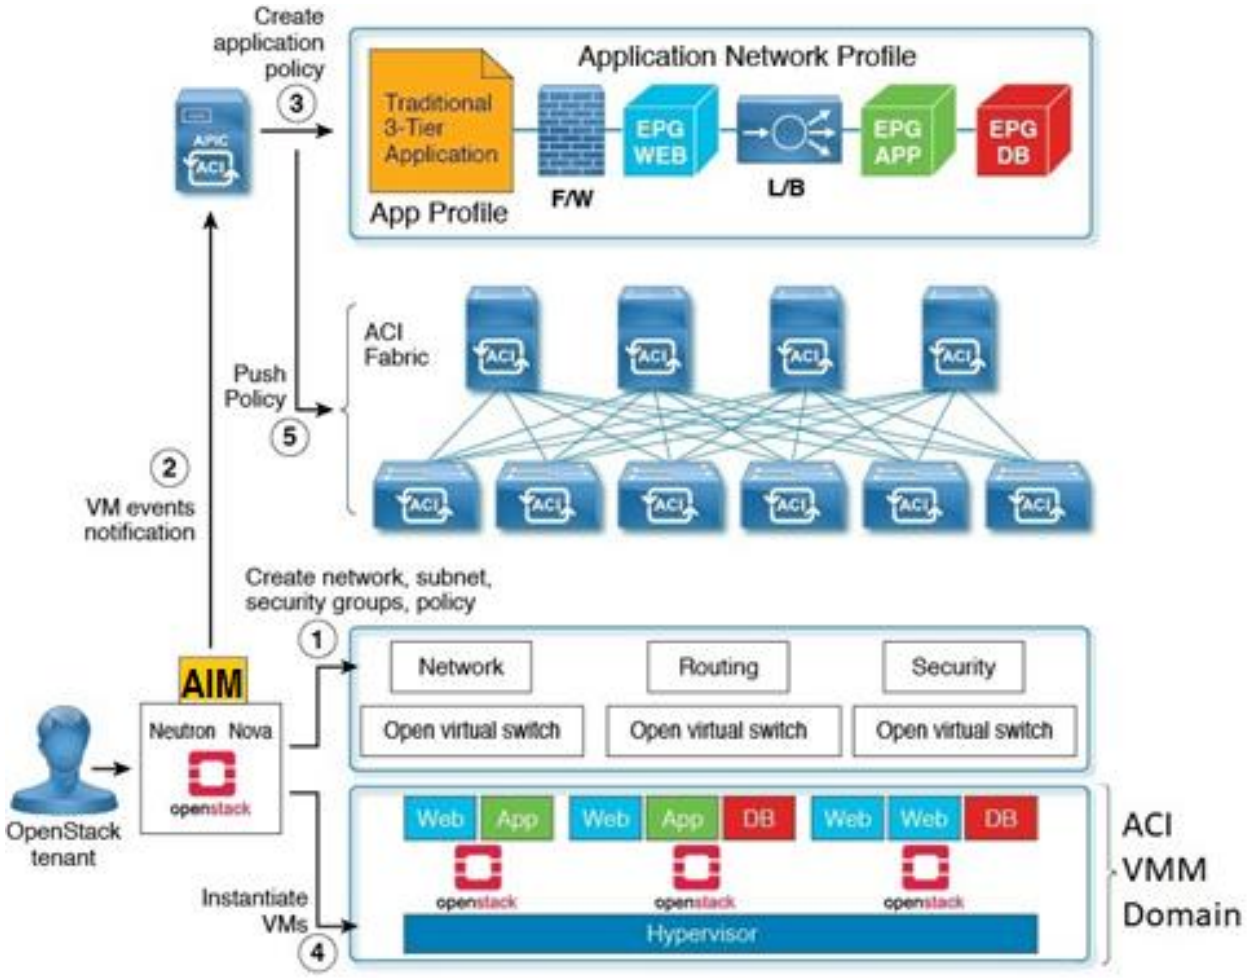

AIM places REST API calls to the APIC for creation of required objects such as the VMM domain, application profiles, EPGs, contracts, and multicast range. The OpFlex agent on a Linux hypervisor performs local traffic switching, routing, and security through Open vSwitch according to policies the AIM plugin pushes into the APIC.

The AIM plugin creates necessary ACI objects according to requests from Neutron on behalf of OpenStack tenants. The Neutron construct mapping to ACI objects are listed below.

OpenStack Neutron ACI Mapping

Neutron	ACI	Cisco IT Naming Convention
Project	Tenant	prj_<uuid of Openstack Project>
Network	EPG + BD	net_<uuid of Neutron Network>
Subnet	Subnet	Subnet Gateway IP/Subnet Prefix
Security Group + Rule	N/A (IP table rules maintained per host)	N/A
Router	Contract + EPG + BD	rtr-<uuid of Neutron router Object>
Network:external	L3Out/Outside EPG	EXT-<name of external EPG>
Address Scope	VRF	DefaultRoutedVRF
Neutron Port	EPG endpoint	N/A

The plugin creates ACI tenants on-demand that map to OpenStack projects.

Cisco IT OpenStack ACI AIM Plugin Case Study

Cisco IT chose to deploy the ACI Integration Module (AIM) for their next-gen OpenStack cloud. The AIM plugin ML2 driver enables OpenStack to maintain the generic Neutron user experience (network/subnet/router/floating IP/security group functions). This is a critical requirement from the larger Cisco user community and business units. This network model allows easy migration of workloads between their private cloud and public cloud where the OpenStack security group is a standard. The engineering IT team requires the generic Neutron interface for its workloads.

ACI Tenant Design Overview

Cisco IT mapped these two consumption models into their OpenStack ACI orchestration design:

- **IT-Managed ACI Tenant**: Deployed as Neutron NO-NAT Networking using single EPG/BD the Cisco IT pre-configured in the ACI Tenant Common. VMs directly attach to the Cisco internal network. IT manages the IP addresses.
- **Self-Managed ACI Tenant**: Deployed as Neutron NAT Networking. Cisco IT chose to have OpenStack orchestrate ACI tenants with 3 EPGs/BDs. The client manages the IP addresses.

Cisco IT set the customizable OpenStack cluster apic_system_id to: `apic_system_id = opens-<cloud_vip_name>`. The AIM plugin uses the `<cloud_vip_name>` names in all ACI objects it creates that are aliased with easier to read names.

Cisco IT configured the ACI Tenant Common functions that it needs, according to what Cisco IT chose to deploy from among the various OpenStack options. Cisco IT chose not to use the OpenStack AIM L4-L7 service graph automation. Instead, they use the OpenStack L4-L7 services that they provision on OpenStack VMs.

Cisco IT OpenStack ACI Tenant Design

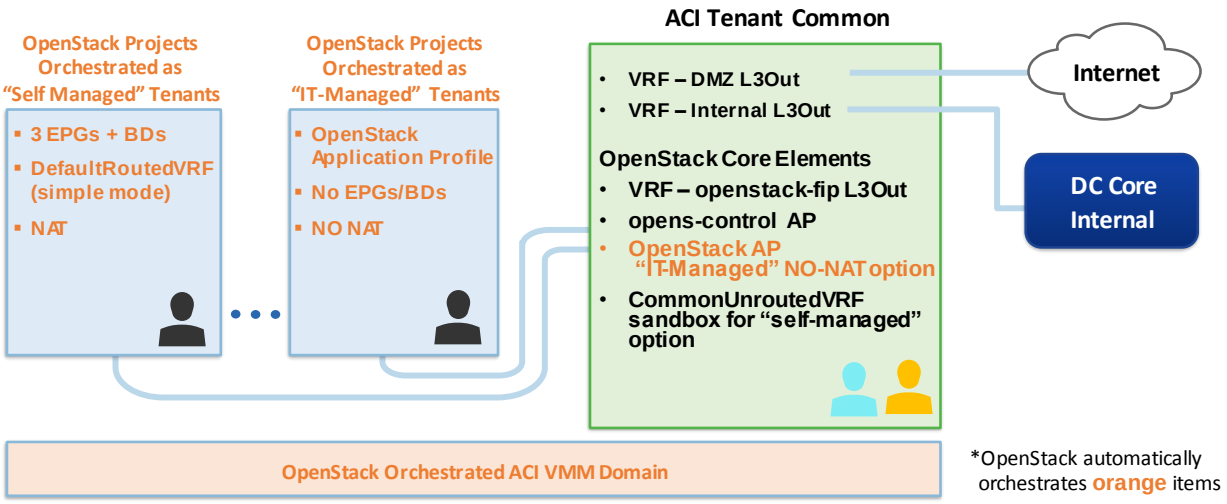

Cisco IT created the `opens-control` application profile in Tenant Common that contains control plane functions. The following EPGs and their corresponding bridge domains enable routable connectivity for OpenStack.

Cisco IT configured `opens-control` application profile EPGs in Tenant Common

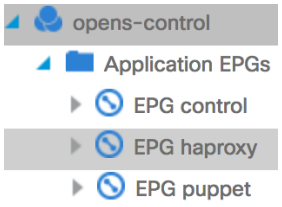

- `control` for OpenStack Neutron controllers & Nova nodes
- `haproxy` for OpenStack endpoints
- `puppet` for Puppet orchestration

Cisco IT also provisions these ACI L3out EPGs for OpenStack external network connectivity:

- `EXT-epg-openstack-fip-l3Out` in Tenant Common that is used for communications between the Neutron network and the external network segments outside the fabric, and for provisioning VMs in OpenStack that require NAT done within OpenStack. Cisco IT provisions a separate VRF (`openstack-fip-vrf`) in Tenant Common for this L3 out
- `EXT-internal-l3Out` in Tenant Common that is used for provisioning VMs in OpenStack provider networks that do not require NAT.

The prefix `EXT-` indicates these BDs and EPGs are used for external connectivity.

VM Orchestration

The figure below illustrates the interaction between the Neutron controller, the AIM plugin, the OpenStack hypervisor, the OpFlex agent, the ACI leaf switches, Nova compute nodes, and VMs.

Cisco IT OpenStack AIM Plugin VM Orchestration

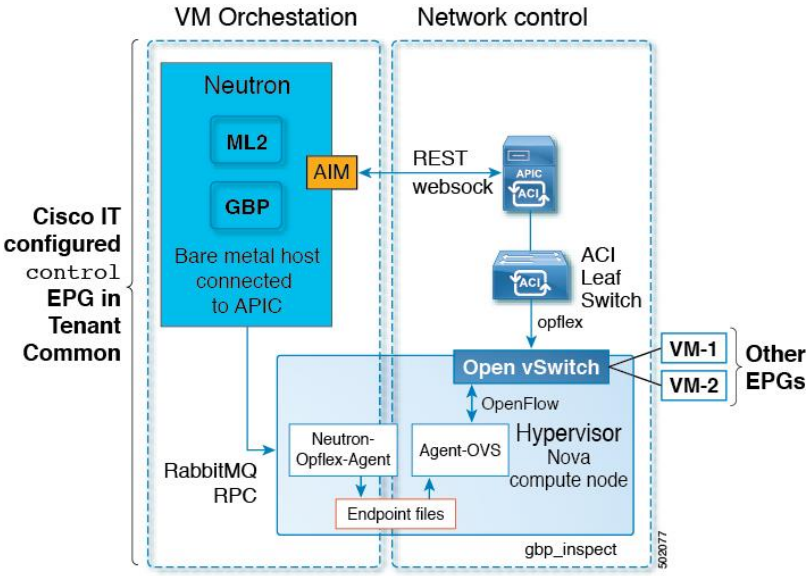

- Neutron-OpFlex-agent - receives updates about new endpoints from the OpenStack Neutron API.
- OpFlex agent - runs the OpFlex protocol with the ACI leaf proxy and programs Open vSwitch via OpenFlow, enforcing ACI policies using OVS OpenFlow rules.
- Endpoint Files - store host local endpoint information in this OpenStack VM located at /var/lib/opflex-agent-ovs.

For sandbox purposes, Cisco IT creates an additional VRF in Tenant Common for VM's without any Neutron router attached. Once OpenStack creates the Neutron router object and attaches it to an external network, the VMs move to the `DefaultVRF` that OpenStack creates under each of the orchestrated "self-managed" ACI tenants.

As illustrated in the figure below, the AIM plugin applies Cisco naming conventions when it orchestrates the ACI VMM domain by prepending `opens-` and appending `-VMM` to the VMM domain name. The AIM plugin associates this VMM domain with the `ucsStnd-AttEntPrfl` AEP Cisco IT created and creates a multicast range for the VMM domain.

OpenStack orchestrated ACI VMM Domain

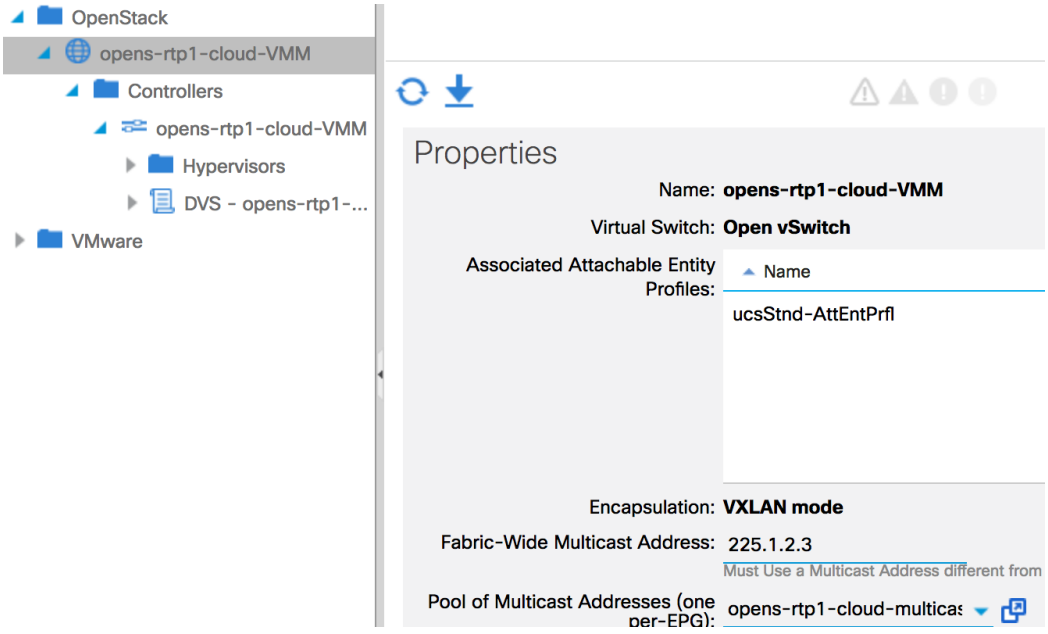

The OpenStack Neutron configuration specifies the multicast range in the `mcast_ranges` variable under the `[ml2_cisco_apic]` section of the neutron.conf file. The plugin prepends the ACI multicast range object name with **opens**.

OpenStack orchestrated multicast range in ACI

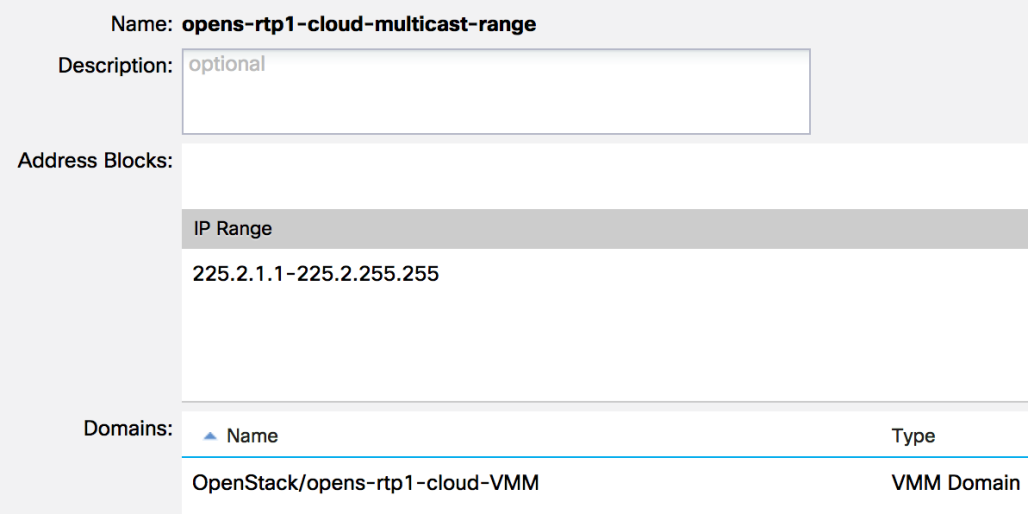

The AIM plugin constantly synchronizes with ACI, identifying which VMs (endpoints) are running on each OpenStack hypervisor and providing information on those EPGs.

OpenStack orchestrated ACI VMM Domain VM endpoint visibility

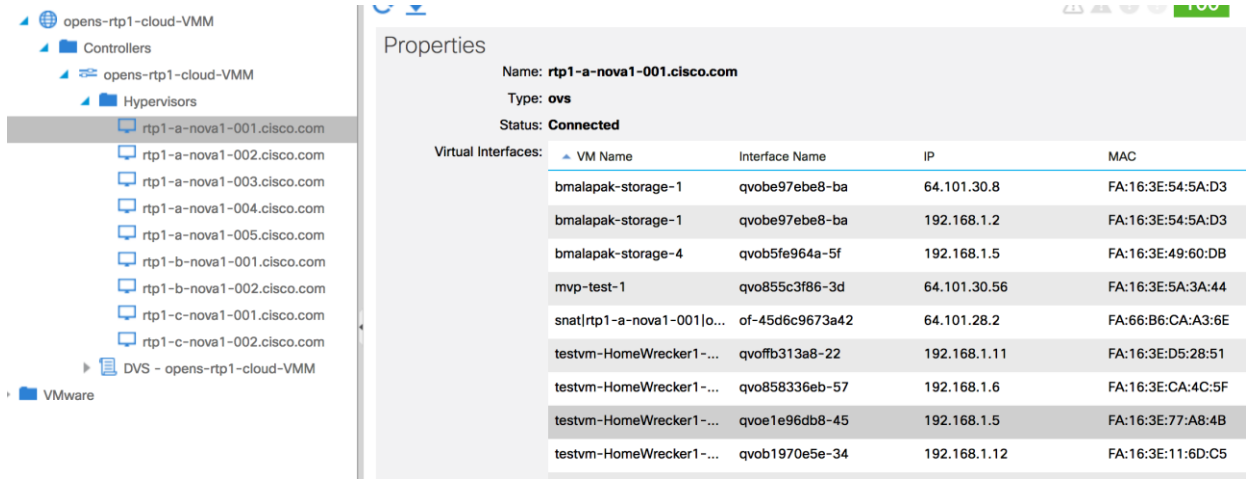

Neutron ports correspond to VMs that are endpoints in an ACI EPG. A port created using Neutron APIs belongs to an ACI EPG. A port created as a GBP PolicyTarget does not use the EPG of the PolicyTargetGroup L2Policy network. Instead, that port belongs to an ACI EPG that OpenStack maps from the port's PolicyTargetGroup.

Cisco "IT- Managed" No-NAT Tenant Option

The Cisco "IT-managed" tenant uses the NO-NAT Neutron Networking option, with VMs attached to the Cisco internal network. Prior to deploying the OpenStack cluster with ACI, Cisco IT pre-allocated three appropriately sized externally routable IP address blocks. Cisco IT manages the assignment of these IP addresses.

OpenStack uses the AIM plugin to orchestrate the creation of the OpenStack Application profile in Tenant Common, and the Tenant IT-Managed tenant.

Cisco "IT-Managed" OpenStack NO-NAT ACI Tenant Design

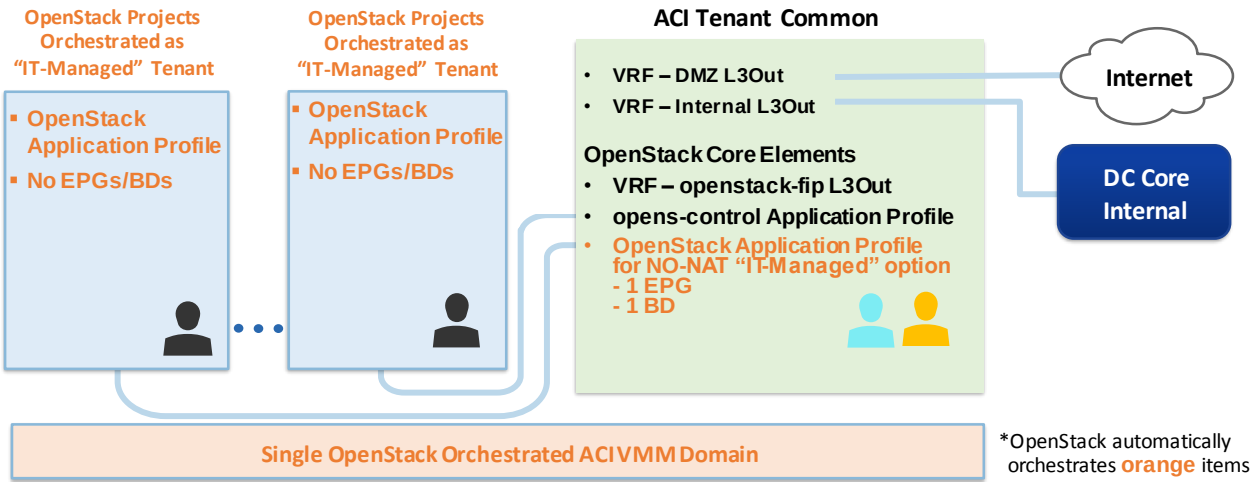

No EPGs are created in the OpenStack application profile under this Tenant IT-Managed tenant. All OpenStack orchestrated "IT-Managed" tenants use the single EPG + BD (corresponds to OpenStack External Network) in the Tenant Common OpenStack application profile and the existing L3 Out EPGs in Tenant Common.

OpenStack orchestrated "IT-Managed" Tenant

The external neutron network shows up on the fabric under Tenant Common of the named with the <cloud-vip> and is assigned the alias "Tenant IT-Managed". Security is

handled in OpenStack security groups. Open vSwitch does the Layer 2 switching. These Cisco IT standard ACI forwarding objects do the Layer 3 forwarding:

- `internal-vrf`
- `internal-l3Out`
- `cisco-internal-extNet`
- `EXT-internal-l3Out`

For traffic engineering purposes, the AIM plugin adds an `allow all` contract from the `EXT-internal-l3Out` EPG to the external EPG (`cisco-internal-extNet`). Cisco IT does not expose the AIM plugin to tenant or end-users.

Cisco IT "Self-Managed" NAT Tenant Option

The Cisco IT "self-managed" option is designed for cloud native workloads with floating IP addressing and VMs using an OpenStack router (ACI BD + EPG + contract) that connects to the Cisco internal network. Users manage their own IP addresses. Cisco IT does not expose the AIM plugin to tenant or end-users directly. Security is handled in OpenStack security groups.

Cisco IT provisions the following ACI forwarding objects in Tenant Common with their respective specified names for NAT Neutron networking configurations:

- VRF: `openstack-fip-vrf`
- EPGs:
    - `cisco-internal-extNet`
    - `openstack-fip-l3Out-DefaultVRF`

Cisco IT OpenStack "Self-Managed" Tenant Design

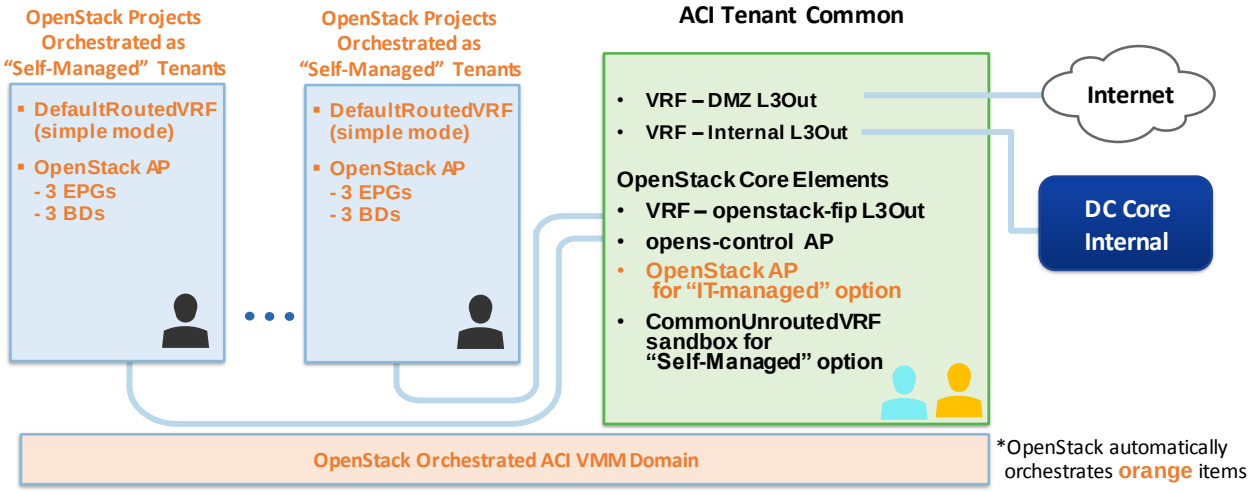

For the "self-managed" NAT mode option, OpenStack uses the AIM plugin to orchestrate ACI tenants with an Application Profile, VRF (`DefaultRoutedVRF`), L3out, and BDs + EPGs that mimic the traditional VLAN model.

OpenStack creates special Neutron networks (`router:external flag set`) that are shared across all the tenants. These Neutron networks provide the following functions:

- SVI for the Neutron network
- Floating IP (DNAT) for external connections to VMs
- A subordinate Neutron network that allows for SNAT configurations and helps VMs reach outside networks

The plugin uses the OpFlex extension agent to implement floating IP (DNAT) in the VM. VM egress traffic (SNAT) is handled inside OVS on the hypervisor.

Through, the Neutron controller uses the VXLAN TypeDrivers and OpFlex Neutron network segments, subnets, routers, or external networks that the AIM plugin translates into ACI constructs. The plugin restricts all tenant specific traffic to the `DefaultRoutedVRF`, which is a VRF the AIM plugin creates.

The external neutron network shows up on the fabric aliased as "Tenant Self-Managed" with the <cloud-vip> name in parentheses.

OpenStack orchestrated "Self-Managed" Tenant

For the "self-managed" option, Cisco IT decided to use 3 Neutron networks, so each self-managed tenant has 3 BDs + 3 EPGs. OVS does the Layer 2 switching and NAT. The ACI fabric does the Layer 3 forwarding.

OpenStack orchestrated "Self-Managed" Tenant ACI constructs for NAT

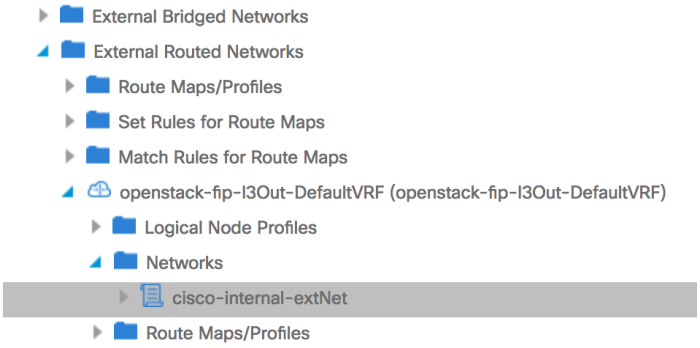

For the "Self-Managed" Tenant NAT functions, the AIM plugin automatically orchestrates all the necessary ACI objects. Open vSwitch instances distributed across the compute nodes automatically do all the necessary NAT transformations that enable traffic flowing from an external network through the ACI fabric to endpoints that have private NAT addresses. In this scenario, a VM has both a private IP interface that is configured in an ACI "self-managed" tenant VRF, and a floating IP interface that is configured in an ACI Tenant Common VRF. For communications between an endpoint from the OpenStack external network on the floating IP interface and an endpoint in the "self-managed" private network, OVS automatically performs the necessary NAT transformations that enable the packet to traverse the Tenant Common VRF, and the "self-managed" VRF.

Cisco IT OpenStack ACI CMDB Integration

ServiceNow can provide real-time visibility across scale-out ACI infrastructure, so that Cisco IT can quickly resolve incidents and drastically reduce mean time to recovery (MTTR).

Cisco IT OpenStack orchestrated ACI Service Now Integration

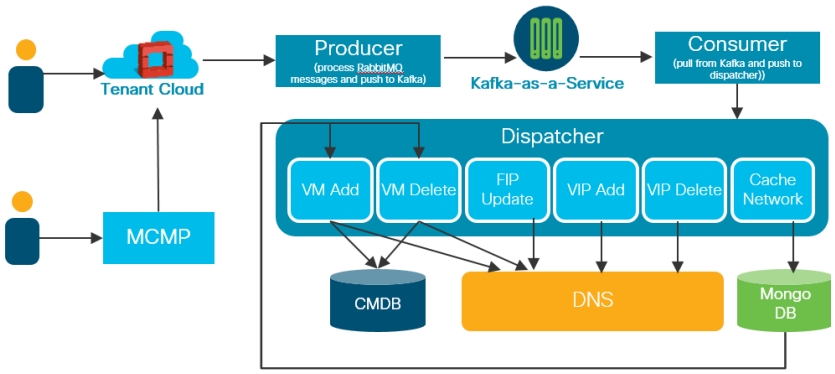

With OpenStack ACI integration, Cisco IT can automate the registration of infrastructure items instantiated through its Multi-Cloud Management Platform (MCMP) in the OpenStack cloud with Cisco CMDBs:

- CMDB maintains a record of infrastructure metadata in the DC cloud
- Maintains DNS entries and ownership

The "data pump" automatically registers new VMs, FIP, and VIPs (LBaaS) with CMDB systems based on:

- Network type: managed, or self-managed
- DMZ or internal
- IT or Engineering

ACI enriches the ServiceNow CMDB with physical inventory as well as configuration tracking, configuration drift analysis, and configuration rollback and roll forward. By correlating infrastructure with existing services, ServiceNow with Cisco ACI discovers and models components that are part of a business service. This unique approach eliminates irrelevant data points and creates an accurate service-aware view that is easily kept up to date with full automation and compliance.

Conclusion

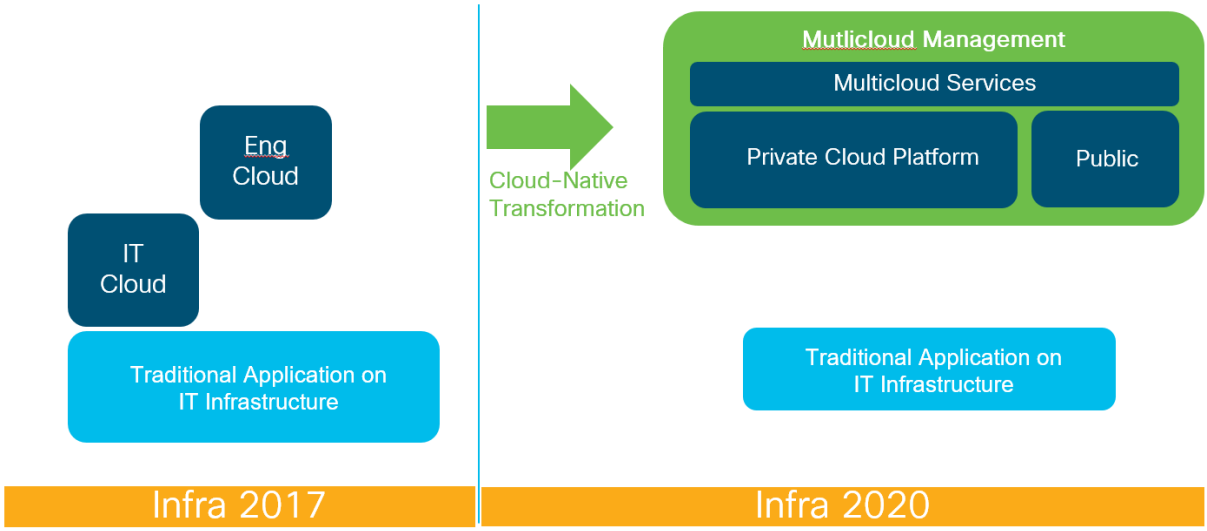

The Cisco® IT deployment of OpenStack ACI AIM plugin enables them to execute their strategic multi-cloud data center vision. This vision for the global data center network delivers the enhanced business value they must have – compelling total cost of ownership, near 100% availability, enhanced network access security, and agility that includes letting business applications developers directly provision the infrastructure resources they need in a self-service fashion.