White Paper

Cisco ACI Fabric and WAN Integration with Cisco Nexus 7000 Series Switches and Cisco ASR Routers

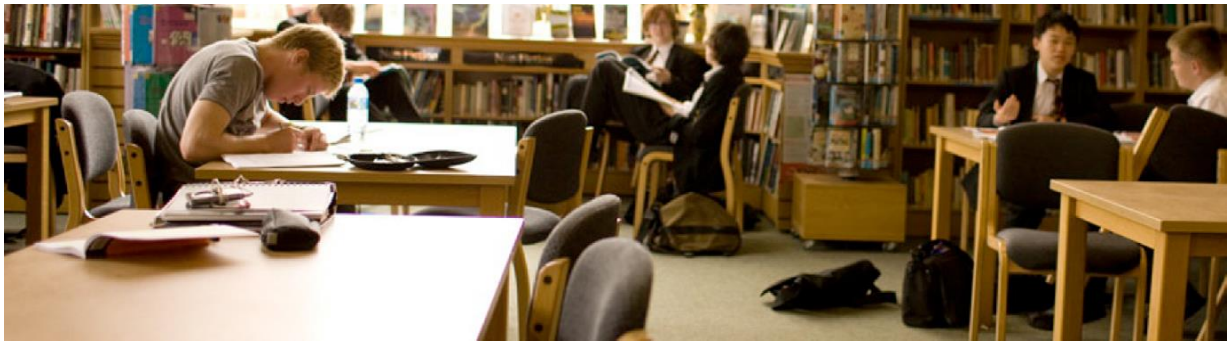

Cisco® Application Centric Infrastructure (Cisco ACI™) is a software-defined networking (SDN) architecture based on a declarative policy model. Cisco ACI combines both software and hardware innovations and provides a robust programmable network for today's dynamic workloads. It is built on a network fabric that combines time-tested protocols with these innovations to create a highly flexible, scalable, and resilient architecture of low-latency, high-bandwidth links. This fabric delivers a network that can support the most demanding data center environments

This document discusses the integration of Cisco ACI with Cisco Nexus® 7000 Series Switches and Cisco ASR 9000 and 1000 Series Aggregation Services Routers. In this architecture, the Cisco Nexus 7000 Series Switches and ASR routers perform the role of WAN edge routers and optionally provide connectivity between multiple Cisco ACI fabrics deployed at the same site or across geographically dispersed sites.

The integration of Cisco ACI into an existing Cisco Nexus environment does not require replacement of existing Cisco Nexus 7000 Series Switches. As demonstrated in the use cases described in this document, customers can transform their businesses by adopting Cisco ACI and integrating this platform with existing Cisco Nexus infrastructure.

Note: Be sure to check the release notes to verify the hardware and software requirements for implementing this solution on Cisco Nexus 9000 Series Switches running in ACI mode and on Cisco Nexus 7000 Series Switches and ASR routers.

The Cisco Nexus 7000 Series Switches and ASR routers function as a peripheral WAN and data center interconnect (DCI) to the Cisco ACI fabric. They use the Cisco OpFlex™ protocol to interface with the Cisco ACI fabric to automate provisioning and to exchange tenant-specific information.

Figure 1 shows the existing integration model using a traditional Layer 2 and Layer 3 hand-off mechanism on a pair of Cisco ACI devices (border leaf nodes) to extend Layer 2 and Layer 3 connectivity outside the Cisco ACI fabric. This is the current model of operation not only for Cisco ACI, but also for other existing fabric technologies and traditional SDN solutions.

Figure 1. Traditional Layer 2 and Layer 3 Hand-Off: Cisco ACI to Cisco Nexus 7000 Series DCI and WAN Edge Routers

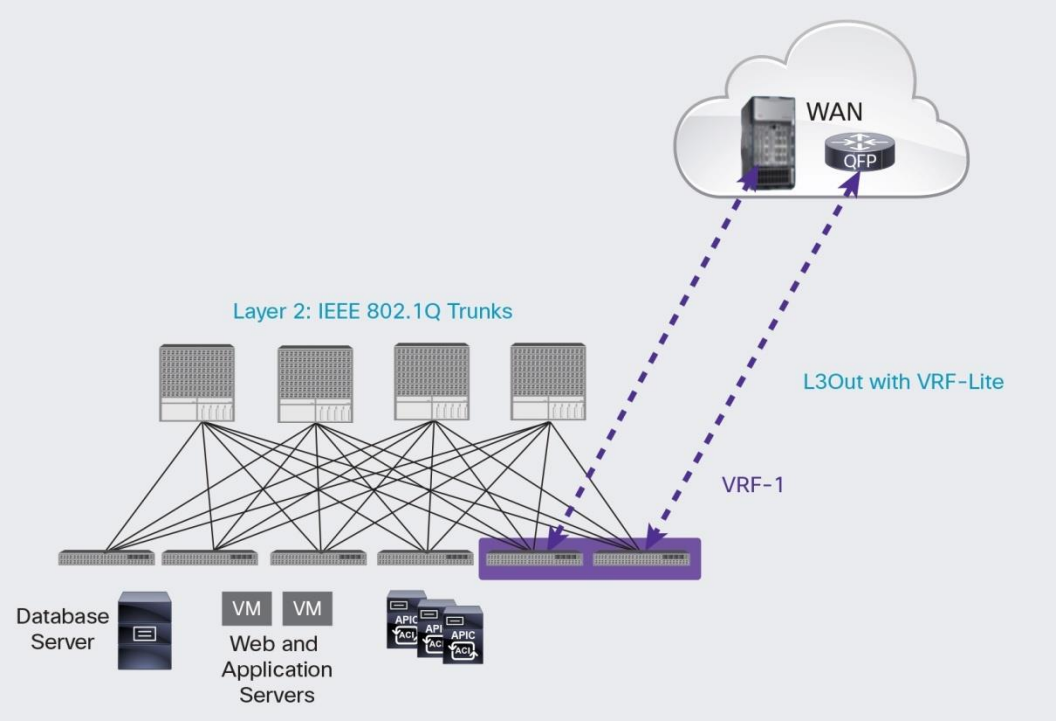

The challenge customers face with the current option is increasing complexity as they increase the number of tenants (Virtual Routing and Forwarding [VRF] instances) defined inside the Cisco ACI fabric, requiring external connectivity services and manual configuration of the WAN edge devices. This problem isn't specific to application-centric infrastructure; the lack of automation of the WAN hand-off from the data center has been a long-time challenge.

Solving the Challenge of WAN Hand-Off Automation with Standard Protocols

With the introduction of Cisco ACI WAN and DCI router integration using Multiprotocol Border Gateway Protocol (MP-BGP) Ethernet VPN (EVPN), customers now can implement complete data center automation, including automation of the hand-off to the WAN and to the DCI routers.

Figure 2 shows the Layer 3 hand-off mechanism offered by the OpFlex and MP-BGP EVPN–based integration model.

Figure 2. Automation of Layer 3 MP-BGP EVPN Hand-Off with OpFlex Policy

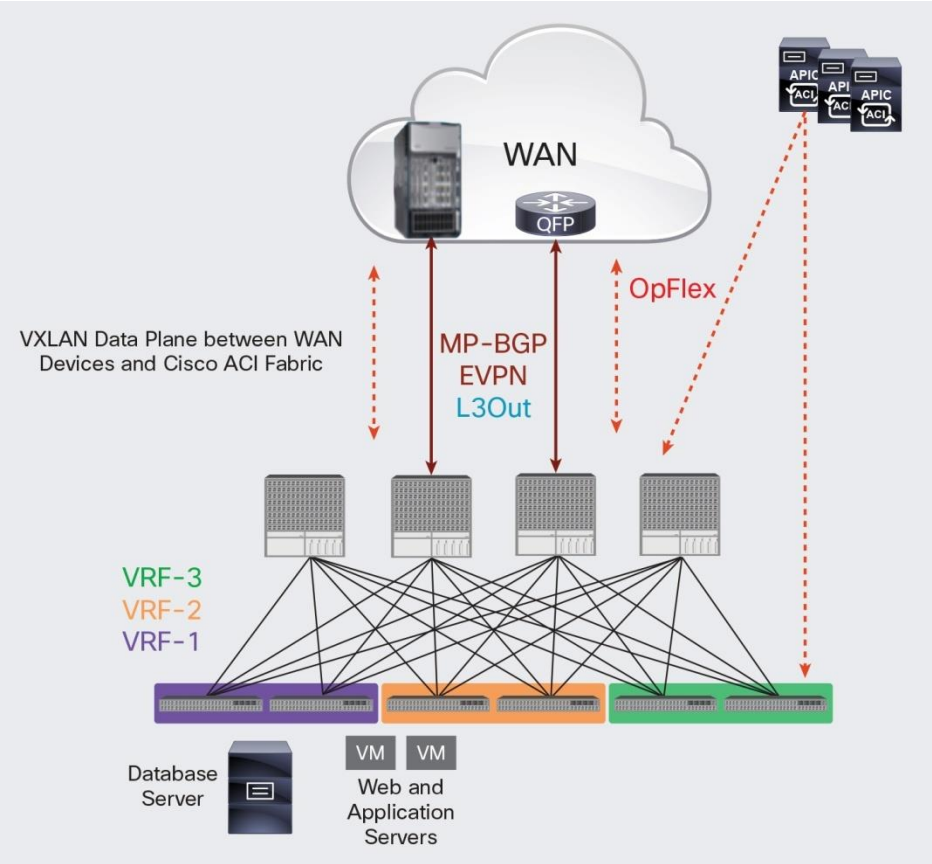

The new model represented in Figure 2 uses the following three main features:

- MP-BGP EVPN standard control plane between the WAN edge devices and the Cisco ACI fabric spine switches: A single BGP session is required to exchange reachability information for multiple tenants (VRF-1, VRF-2, and VRF-3 in the example in Figure 2), removing the per-VRF session requirements of the traditional integration model. This approach is similar to the use of MP-BGP (VPN Version 4 [VPNv4] address family) in Multiprotocol Label Switching (MPLS) VPN deployments.
- OpFlex control plane between the WAN edge devices and the Cisco ACI spine switches to automate fabric-facing tenant provisioning on the DCI and WAN edge devices: The network administrator simply configures a new external Layer 3 outside (L3Out) policy for a tenant on the Cisco Application Policy Infrastructure Controller (APIC). The controller then programs all related information associated with that tenant, such as VRF instance name and BGP extended community route-target attributes, for the Cisco ACI spine switches. The OpFlex proxy server running on the spine switches reads the L3Out managed object and converts it to the OpFlex model. This information is then pushed to the WAN edge device that receives the OpFlex model events, automating the fabric-facing tenant configuration through configuration profile templates.

Note: The WAN-facing tenant configuration is not automated.

- VXLAN data plane between WAN edge devices and the Cisco ACI fabric to establish communication with the external WAN routed domain

Note that although the WAN edge devices perform the functional role of Cisco ACI border leaf nodes (that is, they encapsulate and decapsulate communication with the external network domain), those switches are not fully managed by the APIC controller, and the only control plane interactions performed are the ones described in the preceding paragraphs. The initial phase of the integration between Cisco ACI and Cisco Nexus 7000 Series Switches and ASR routers is limited to the establishment of Layer 3 communication with the WAN (that is, Layer 2 hand-off is not supported).

Extending the VRF connectivity to the WAN edge devices also allows transparent integration with an MPLS VPN service. In this scenario, the WAN edge device can also perform the provider edge (PE) role, and the automation of configuration with the Cisco ACI fabric helps ensure end-to-end connectivity between the MPLS core and the fabric resources.

Connectivity Options

As shown in Figure 3, two connectivity options are supported between the WAN edge devices and the Cisco ACI fabric:

- Direct connection between Cisco ACI spine switches and WAN and DCI switches: This option requires the use of 40- and 100-Gbps interfaces between those devices.
- Use of an IP transport network to provide control plane and data plane connectivity between the Cisco ACI fabric and the WAN devices. This option essentially extends application-centric infrastructure (that is, the underlay) to the WAN edge devices.

Figure 3. Options for Connectivity between WAN Edge Devices and Cisco ACI Fabric

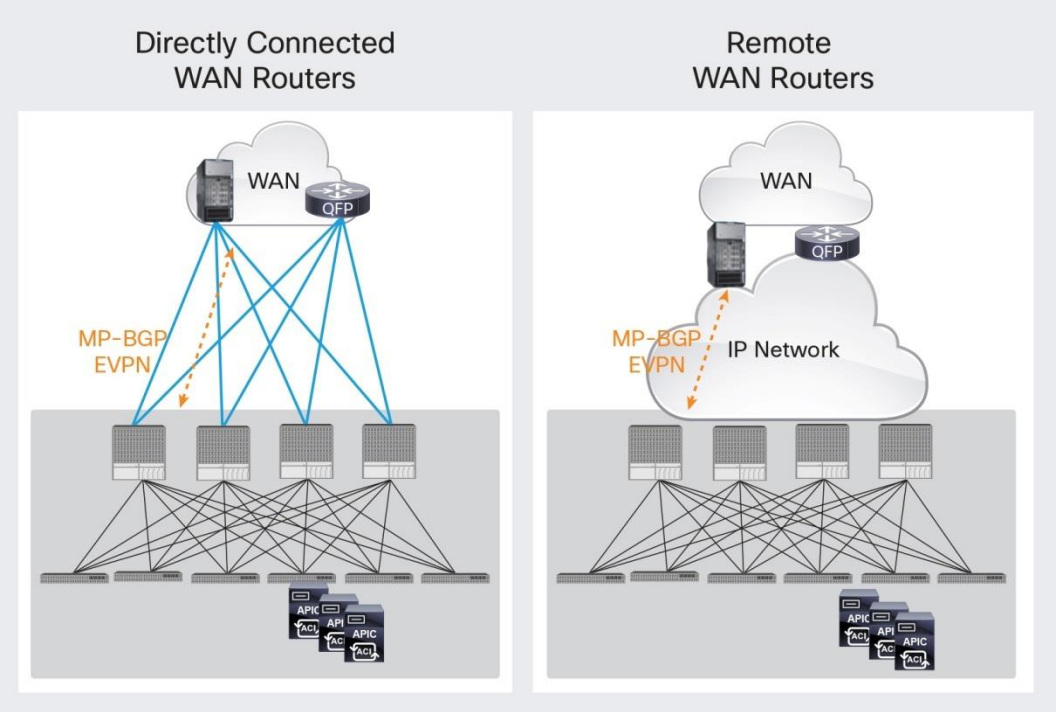

WAN Edge Routers and Security Policy Enforcement

Communication between endpoints deployed within the Cisco ACI fabric does not depend only on reachability information. It is also subject by default to security policy enforcement (the Cisco ACI fabric uses a zero-trust whitelist security model). The same requirements apply to communication with the external routed network domain.

The security policies are still defined on the APIC controller. The external network resources are modeled as one (or more) security groups (external endpoint groups [EPGs]) so that specific policies (contracts) can be applied to provide communication with EPGs internal to the Cisco ACI fabric.

In the MP-BGP EVPN–based integration model, security policies are always enforced on the Cisco ACI leaf nodes (that is, the WAN devices cannot apply policies in hardware). Thus, for communication with the WAN, the policies are enforced at the leaf node to which the endpoint sourcing the traffic is connected (the WAN edge devices ignore policy information contained in the VXLAN header). For traffic originating on the WAN and destined for workloads within the Cisco ACI fabric, security policies are instead enforced at the Cisco ACI leaf nodes to which those workloads are connected. This behavior simplifies the integration with the WAN edge devices while still maintaining the end-to-end Cisco ACI security policy.

Conclusion

The new model for integrating Cisco ACI fabric with Cisco Nexus 7000 Series Switches and ASR routers is based on open protocols such as MP-BGP EVPN and OpFlex, with Virtual Extensible LAN (VXLAN) used as the data plane. This model provides a scalable and operationally simpler mechanism for extending Layer 3 communication between the data center fabric and the external WAN domain, solving a pervasive challenge in today's data centers. This approach eliminates the need to establish a per-VRF session for each tenant to communicate with the outside world.

Direct and indirect connections (through a transport IP network) are supported to establish control and data plane communication between the Cisco ACI fabric and the WAN devices.

In addition, consistent security policies can be defined for communication with the external Layer 3 domain. These policies are always enforced at the Cisco ACI leaf nodes for both inbound and outbound traffic flows.

For More Information

For more detailed information about Cisco ACI functions and deployment considerations, please refer to the resources at http://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html.

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 www.cisco.com/go/offices.

Printed in USA C11-736899-00 03/16