

CSG

Cisco Validated Profile Series

Enterprise Routing

Cisco ACI Fabric and WAN Integration with Cisco ASR 1000 Router

Contents

1. Profile Introduction	3
2. Network Profile	3
a. Topology Diagram & Hardware Specifications	3
2.1. ACI EVPN L3 DCI vrf-lite over MPLS Secure WAN	4
2.2. ACI EVPN L3 DCI as PE	5
2.3. ACI EVPN L3 DCI vrf-lite over MPLS with IPN support	6
2.4. ASR1K SDA Border with ACI and Policy Plane integration	7
i. Key Vertical Features	8
ii. Hardware Profile	10
b. Test Environment	10
3. Use-Case Scenarios	11
3.1. Test Methodology	11
3.2. Use-Cases	11
3.2.1. ACI EVPN L3 DCI vrf-lite over MPLS Secure WAN	12
3.2.1.1 Routing	12
3.2.1.2 Security	12
3.2.1.3 Network Services	12
3.2.1.4 Simplified Management	12
3.2.1.5 System Health Monitoring	12
3.2.1.6 System & Network Resiliency, Robustness	12
3.2.2. ACI EVPN L3 DCI as MPLS PE	12
3.2.2.1 Routing	12
3.2.2.2 Network Services	12
3.2.2.3 Simplified Management	12
3.2.2.4 System Health Monitoring	12
3.2.2.5 System & Network Resiliency, Robustness	13
3.2.3. ACI EVPN L3 DCI vrf-lite over MPLS with IPN support	13
3.2.3.1 Routing	13
3.2.3.2 Network Services	13
3.2.3.4 Simplified Management	13
3.2.3.5 System Health Monitoring	13
3.2.3.6 System & Network Resiliency, Robustness	13
3.2.4. ASR1K SDA Border with ACI and Policy Plane Integration	14
3.2.4.1 Routing	14
3.2.4.2 Security	14
3.2.4.3 Network Services	14
3.2.4.4 Monitoring & Troubleshooting	14
3.2.4.5 Simplified Management	14
3.2.4.6 System Health Monitoring	14
3.2.4.7 System & Network Resiliency, Robustness	14
4. Appendix A	15
5. Acronyms	15
6. Configuration	15
6.1 Usecase: ACI EVPN L3 DCI vrf-lite over MPLS Secure WAN	15
6.2 Usecase: ACI EVPN L3DCI as PE	27
6.3 Usecase: ACI EVPN L3DCI vrf-lite over MPLS with IPN support	33
6.4 Usecase: ASR1K SDA Border with ACI and Policy Plane Integration	36

1. Profile Introduction

TrustSec and ACI (Application Centric Infrastructure) are two architectures that assign hosts to group and administer policies in terms of these groups. They are managed independently by two different domain controllers. (ISE (Identity Services Engine) for TrustSec and APIC (Cisco Application Policy Infrastructure Controller) for ACI). They both have a group-based policy framework. Users and resources are categorized into groups and access control actions are drawn across groups. TrustSec assigns hosts to Security Groups identified by a 16-bit number called Security Group Tag (SGT). ACI assigns hosts to End-Point Groups (EPG) identified by a 16-bit number called ClassId. Different domain controller between TrustSec and ACI requires manual reconciliation of group identifiers across these domains. Due to discrete group namespaces, there is no way to normalize policies for traffic going from TrustSec Enterprise Branch to ACI data center and vice versa.

TrustSec ACI solution integrates with policy plane and allows interconnection of administrative domains of TrustSec border gateway and ACI border gateway to provide a consistent end-to-end policy enforcement experience for our customers.

This Profile is designed in such a way to integrate key requirements in any wan aggregation router and to validate the feature interoperability in a typical deployment.

Table 1. ACI Fabric and WAN Integration Profile Feature Summary

Deployment Areas	Features
Security	DMVPN, SGACL, SGFW, ISE, TrustSec
Network Planning & Troubleshooting	NBAR, FNF
Management & Monitoring	SNMP, SysLog Server
System Resiliency	Interface flaps, ESP/SIP/SPA failovers
Network Services	OSFP, MP BGP EVPN, SGT QOS

2. Network Profile

Based on the research and customer feedback and configuration samples, the Cisco ACI Fabric and WAN Integration with Cisco ASR1000 Router Profile is designed with a deployment topology that is generic and can easily be modified to fit any specific deployment scenario.

a. Topology Diagram & Hardware Specifications

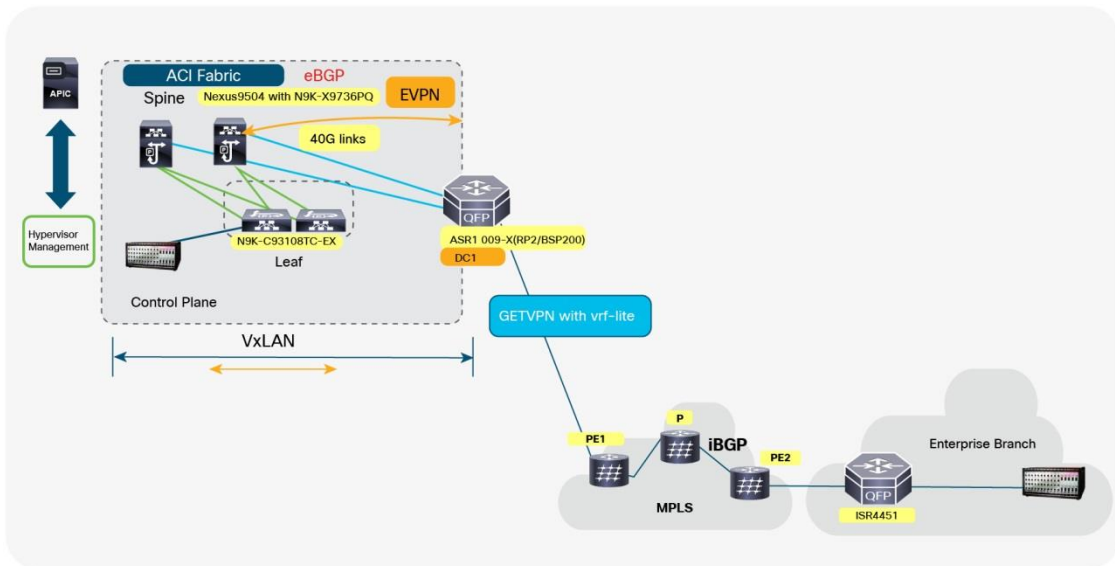
1. **Disclaimer:** The links between the different network layers in the topology are mainly to facilitate this profile validation across different platform combinations and the actual deployment could vary based on specific requirement.

We have the following usecases:

- ACI EVPN L3 DCI vrf-lite over MPLS Secure WAN
- ACI EVPN L3 DCI as PE
- ACI EVPN L3 DCI vrf-lite over MPLS with IPN support
- ASR1K SDA Border with ACI and Policy Plane integration

2.1. ACI EVPN L3 DCI vrf-lite over MPLS Secure WAN

Figure 1. ACI EVPN L3 DCI vrf-lite over MPLS Secure WAN



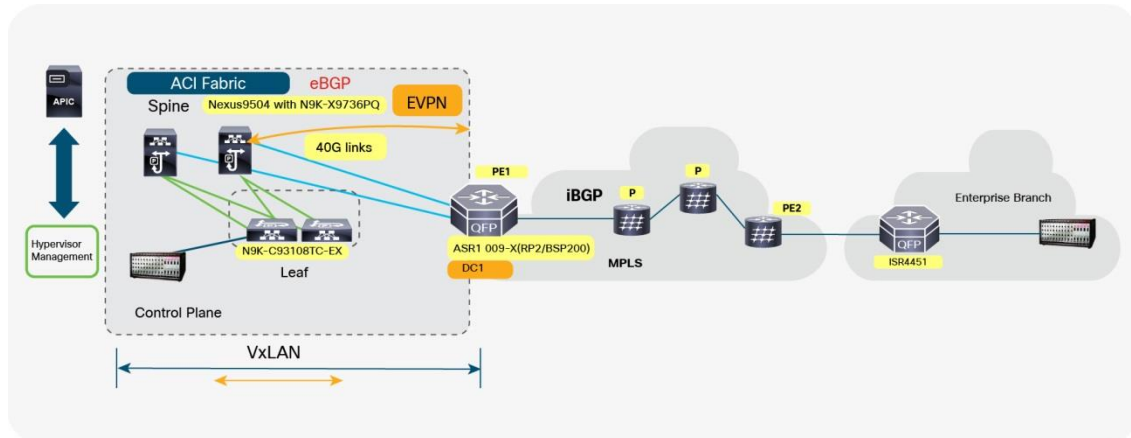
The left-portion of the topology represents ACI fabric consisting of the APIC and the N9k switches. The right portion is the MPLS network that uses the crypto functionality of GETVPN to connect to the Enterprise branch.

Network Device	Platform
DC1	ASR1009-X(RP2/ESP200)
Spine	Nexus9504 with N9K-X9736PQ
Leaf	N9K-C93108TC-EX
PE1	ASR1002-X
PE2	ASR1006(RP2/ESP40)
P	ASR1006(RP2/ESP40)
BRANCH	ISR 4451

Features/Functionalities Tested
<ul style="list-style-type: none"> • DCI connects the ACI fabric with BGP EVPN + Vxlan • Opflex is used to learn the VRF/BD/BDI/VNI on the ACI fabric side • Vrf-lite interfaces on the WAN side connecting the MPLS network • GETVPN is enabled on the vrf-lite interfaces on the WAN side of DCI • Branch is also a GETVPN GM • Policing QoS policy is applied on the BDI and the GETVPN interface on DCI

2.2. ACI EVPN L3 DCI as PE

Figure 2. ACI EVPN L3 DCI as PE



The left-portion of the topology represents ACI fabric consisting of the APIC and the N9k switches. The right portion is the MPLS network to connect to the Enterprise branch. In this use case the DCI also acts as a PE and is part of the MPLS network.

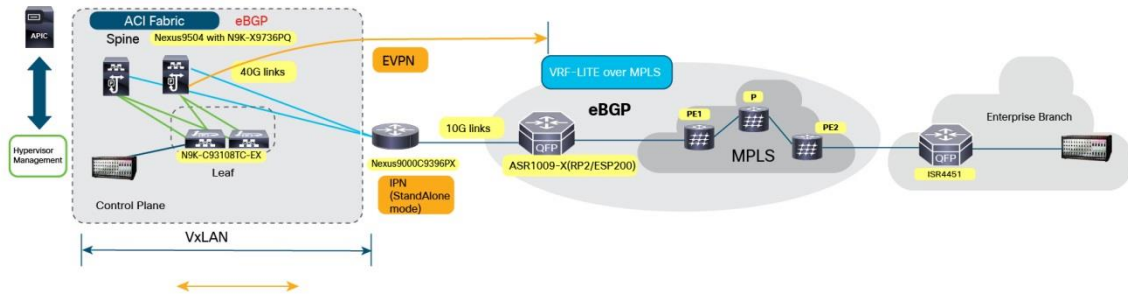
Network Device	Platform
DCI/PE	ASR1009-X(RP2/ESP200)
Spine	Nexus9504 with N9K-X9736PQ
Leaf	N9K-C93108TC-EX
PE1	ASR1002-X
PE2	ASR1006(RP2/ESP40)
P	ASR1006(RP2/ESP40)
BRANCH	ISR 4451

Features/Functionalities Tested
<ul style="list-style-type: none"> DCI connects the ACI fabric with BGP EVPN + Vxlan Opflex is used to learn the VRF/BD/BDI/VNI on the ACI fabric side DCI acts as PE on the WAN side EVPN BGP prefixes are converted to vpnv4 routes and vice-versa on the DCI

2.3. ACI EVPN L3 DCI vrf-lite over MPLS with IPN support

In a Multi-Pod deployments, very often customers ask to connect the GOLF devices to the already existing IPN network. Customers use this deployment to use 10G interfaces on the GOLF devices and 40G on the Spines.

Figure 3. ACI EVPN L3 DCI vrf-lite over MPLS with IPN support



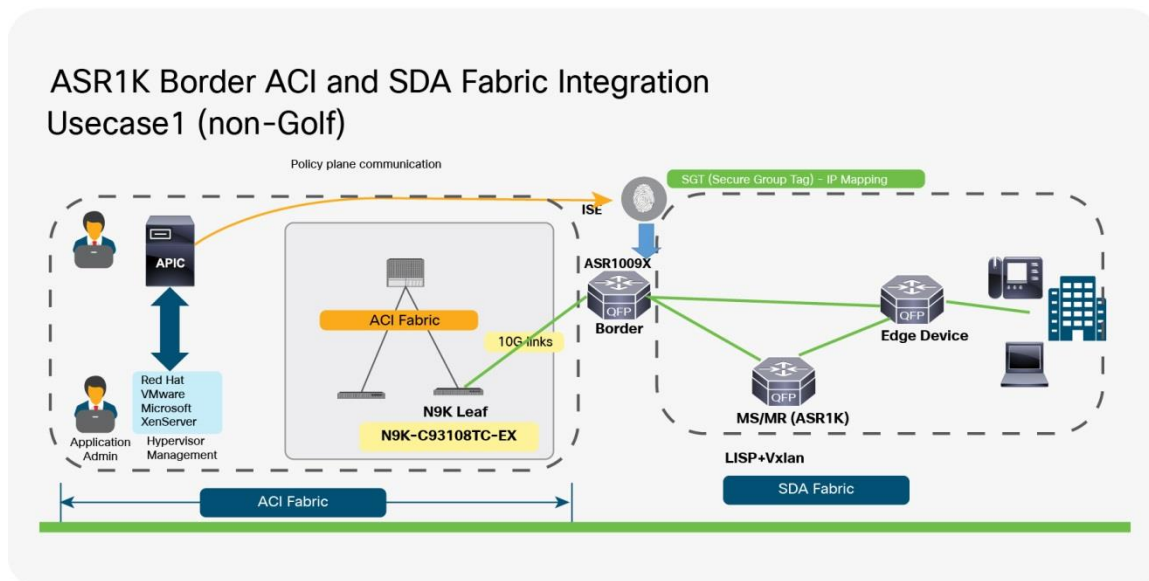
The left-portion of the topology represents ACI fabric consisting of the APIC and the N9k switches. The connectivity to the DCI is not direct but via a IPN (standalone).The connectivity to the DCI from the IPN is via a 10 G link. The right portion of the topology is the vrf-lite MPLS network.

Network Device	Platform
PE	ASR1009-X(RP2/ESP200)
IPN	Nexus9000 C9396PX
Spine	Nexus9504 with N9K-X9736PQ
Leaf	N9K-C93108TC-EX
PE1	ASR1002-X
PE2	ASR1006(RP2/ESP40)
P	ASR1006(RP2/ESP40)
BRANCH	ISR 4451

Features/Functionalities Tested
<ul style="list-style-type: none"> • 40G link between Spine and IPN. 10G link between IPN and DCI • OSPF is underlay routing protocol. Its configured on Spine, IPN and DCI • L3OUT interface is on the Spine connecting the IPN • BGP EVPN is between the Spine and DCI • Opflex agent is on DCI with peer as Spine L3 out interface

2.4. ASR1K SDA Border with ACI and Policy Plane integration

Figure 4. ASR1K SDA Border with ACI and Policy Plane integration



The left-portion of the topology represents ACI fabric consisting of the APIC and the N9k switches. DCI is connected to N9K Border Leaf. It also connects to MS/MR and SDA fabric. We have BGP between ASR1K Border and ACI Border Leaf.

ASR1K acts as DCI Border and it integrates both ACI Fabric Data Center and SDA Fabric – Policy Plane integration.

- ASR1K is directly connected to N9K Leaf (ACI Fabric).
- 10 VRFs on the ACI side are mapped to one common VRF on the Leaf. Interface between N9K Leaf and ASR1K DCI Border are on the ACI VRF.
- eBGP between ASR1K and N9K Leaf.
- All prefixes from 10 VRFs are learnt on ACI VRF on ASR1K DCI Border.
- LISP+VxLAN in the SDA Fabric.
- ASR1K has LISP on the SDA Fabric Side and acts as an xTR.
- LISP on ASR1K DCI Border imports the prefixes learnt from ACI fabric via bgp.
- 30 SDA VRFs.
- 30 VRFs on SDA side talk to one ACI VRF on ASR1K DCI Border using LISP extranet.
- MS/MR has LISP extranet or the MS/MR based route-leaking between 30 SDA VRFs and ACI VRF.
- ACI VRF is the provider, SDA VRFs are the subscribers.
- MS/MR has all the VRFs (ACI and SDA VRFs).
- LISP Scale – 200 EIDs per subscriber VRF (Total 6000 EIDs).
- iBGP between MS/MR and ASR1K DCI. Redistributing LISP to BGP is done on MS/MR.

- Trustsec integration.
- Policy-plane integration on ISE with APIC using SXP service. ISE and APIC exchange SGT — IP bindings.
- Policy enforcement on APIC.

Network Device	Platform
ASR1K SDA Border	ASR1009-X(RP2/ESP200)
Leaf1	Nexus9000 C9396PX
Spine	Nexus9504 with N9K-X9736PQ
Border Leaf	N9K-C93108TC-EX
MS/MR	ASR1001-X
Edge	ASR1001-X

Features/Functionalities Tested
<ul style="list-style-type: none"> • 10 VRFs on the ACI side are mapped to one common VRF on the Leaf. Interface between N0K Leaf and ASR1K DCI Border are on the ACI VRF. • eBGP between ASR1K and N9K Leaf. • All prefixes from 10 VRFs are learnt on ACI VRF on ASR1K DCI Border. • LISP+VxLAN in the SDA Fabric. • ASR1K has LISP on the SDA Fabric Side and acts as an xTR. • ASR1K has LISP on the SDA Fabric Side and acts as an xTR. • LISP on ASR1K DCI Border imports the prefixes learnt from ACI fabric via bgp. • 30 SDA VRFs. • 30 VRFs on SDA side talk to one ACI VRF on ASR1K DCI Border using LISP extranet. • MS/MR has LISP extranet or the MS/MR based route-leaking between 30 SDA VRFs and ACI VRF. • ACI VRF is the provider, SDA VRFs are the subscribers. • MS/MR has all the VRFs (ACI and SDA VRFs). • LISP Scale – 200 EIDs per subscriber VRF (Total 6000 EIDs). • iBGP between MS/MR and ASR1K DCI. Redistributing LISP to BGP is done on MS/MR. • Trustsec integration. • Policy-plane integration on ISE with APIC using SXP service. ISE and APIC exchange. • SGT — IP bindings. • Policy enforcement on APIC.

i. Key Vertical Features

Table-2 defines the 3-D hardware, Place-In-Network (PIN) and the features deployed. The scale of these configured features, the test environment, list of end-points and hardware software versions of the network topology will be defined the subsequent sections of this guide.

2.4.1.1. ACI EVPN L3 DCI vrf-lite over MPLS Secure WAN

Deployment Layer	Platforms	Critical Vertical Features
DCI	ASR1009x(ESP200/ASR1000-MIP100/40G EPA)	MP-BGP EVPN + Vxlan Opflex ACI OSPF as underlay eBGP with Vrf-lite interface towards the Border GETVPN Group Member GDOI Group QoS (Policing) on BDI and GETVPN wan interface
PE1, P, PE2	ASR1006(RP2/ESP40)	MPLS VPN
Branch	ISR4451	GETVPN Group Member eBGP towards PE2
KeyServer	ISR3925E	GETVPN KeyServer GDOI group
ACI fabric	Spine: Nexus9504 with N9K-X9736PQ module Leaf: N9K-C93108TC-EX APIC-M2	MP-BGP EVPN Opflex Tenant with multiple VRFs ACI Mode

Disclaimer: Refer to appropriate CCO documentation for release/feature support across different platforms.

2.4.1.2. ACI EVPN L3 DCI as PE

Deployment Layer	Platforms	Critical Vertical Features
DCI	ASR1009x(ESP200/ASR1000-MIP100/40G EPA)	MP-BGP EVPN + Vxlan Opflex ACI OSPF as underlay MPLS PE on the WAN side QoS (Policing) on BDI interface Imports vpv4 routes to evpn and vice versa
P, PE2	ASR1006(RP2/ESP40)	MPLS VPN
ACI fabric	Spine: Nexus9504 with N9K-X9736PQ module Leaf: N9K-C93108TC-EX APIC-M2	MP-BGP EVPN Opflex Tenant with multiple VRFs ACI mode

Disclaimer: Refer to appropriate CCO documentation for release/feature support across different platforms

2.4.1.3. ACI EVPN L3 DCI vrf-lite over MPLS with IPN support

Deployment Layer	Platforms	Critical Vertical Features
DCI	ASR1009x(ESP200/ASR1000-MIP100/40G EPA)	MP-BGP EVPN + Vxlan Opflex ACI OSPF for underlay eBGP with Vrf-lite interface towards the Border QoS (Policing) on BDI and GETVPN wan interface
PE1, P, PE2	ASR1006(RP2/ESP40)	MPLS VPN
IPN	N9K-C9396PX	Standalone mode OSPF for underlay

Deployment Layer	Platforms	Critical Vertical Features
ACI fabric	Spine: Nexus9504 with N9K-X9736PQ module Leaf: N9K-C93108TC-EX APIC-M2	MP-BGP EVPN Opflex Tenant with multiple VRFs ACI Mode

Disclaimer: Refer to appropriate CCO documentation for release/feature support across different platforms.

2.4.1.4. ASR1K SDA Border with ACI and Policy Plane integration

Deployment Layer	Platforms	Critical Vertical Features
DCI /Border	ASR1009x(ESP200/10G Links), ASR1006x(RP2/ESP40)	eBGP between N9K Border Leaf and ASR1K border ACI OSPF for underlay ASR1K border has BGP and LISP towards SDA fabric VRF route leaking between user VRFs and common VRF
MS/MR	ASR1001X, ASR1006(RP2/ESP40)	MS/MR has LISP extranet to leak routes between VRFs
Edge	ASR1001X, ASR1013 (RP2/ESP200)	LISP XTR 30 VRFs with 200 EIDs in each VRF
ACI fabric	Spine: Nexus9504 with N9K-X9736PQ module Border Leaf: N9K-C93108TC-EX Leaf: N9K-C9396PX APIC-M2	MP-BGP between the ACI Leafs and Spine Non-Golf Scenario Common tenant has multiple user VRFs and default VRF

Disclaimer: Refer to appropriate CCO documentation for release/feature support across different platforms.

ii. Hardware Profile

Table-3 defines the set of relevant hardware, servers, test equipment and end-points that are used to complete the end-to-end Retail Vertical Profile deployment.

List of hardware, along with the relevant software versions and the role of these devices complement the actual physical topology that is defined in Figure-1 of the previous section.

Table 2. Hardware Profile of Servers and End-Points

VM and HW	Software Versions	Description
APIC	apic-2.2(1.104a)	Controller for the ACI fabric
ISE	2.2.470 2.3 (SDA Border usecase)	For authentication and authorization for network devices, SGT creation and translation of SGT to EPG and vice-versa
UCS Server	ESXi 5.5.0	To manage and host the Windows Virtual Machines, IXIA traffic tool etc.
Ixia	IxLoad	Test tool to generate HTTP, FTP, DNS & telnet traffic
Windows VM Clients	Windows 7	End-points to test end to end traffic
Cisco Prime	3.1.5	For Opflex agent configuration using PNP

b. Test Environment

This section contains the relevant scales at which the features are deployed across the physical topology. Table-4 lists out the scale for each respective feature.

Disclaimer:

Table below captures a sample set of scale values used in one of the use cases.

Please refer to appropriate CCO documentation / Datasheets for comprehensive scale data.

Feature	Scale
LISP Extranet	30 VRFs on SDA side are mapped to one VRF on the ASR1K Border
VRF-Leaking in ACI	10 VRFs on the Leaf are mapped to one VRF on Border Leaf. Border Leaf connects to ASR1K on Common/default VRF
Vrfs in tenant on APIC	11
VRFs in SDA	30
Number of EIDs in each VRF in SDA	200 (Total 6000 EIDs from SDA side)
VRF on ASR1K Border	1
VRFs on MS/MR	31 (30 SDA VRFs + 1 ACI VRF)

3. Use-Case Scenarios

3.1. Test Methodology

The Use-cases listed in Table-5 below, will be executed using the Topology defined in Figure-1 along with the Test environment, Table-4, already explained in this document.

Images are loaded on the devices under test via the tftp server using the Management interface.

To validate a new release, the network topology is upgraded with the new software image with existing configuration that comprises of the use-cases and relevant traffic profiles. Addition of new use-cases acquired from the field or customer deployments are added on top of the existing configuration.

During each use case execution, syslog would be monitored closely across the devices for any relevant system events, errors or alarms. With respect to longevity for this profile setup, CPU and memory usage/ leaks would be monitored during the validation phase. Furthermore, to test the robustness of the software release and platform under test, typical networks events would be triggered during the use-case execution process.

3.2. Use-Cases

Table-5 describes the Use-Cases that were executed on the Cisco ACI Fabric and WAN Integration with Cisco ASR 1000 Router Profile. These Use-cases are divided into buckets of technology areas to see the complete coverage of the deployment scenarios. Use-cases continuously evolve based on the feedback from the field.

These technology buckets comprises of Security, Network Services, Monitoring & Troubleshooting, simplified management, system health monitoring along with System resiliency.

3.2.1. ACI EVPN L3 DCI vrf-lite over MPLS Secure WAN

No.	Focus Area	Use Cases
3.2.1.1 Routing		
5	BGP EVPN	<ul style="list-style-type: none"> BGP EVPN is configured on the ACI side OSPF is used for the underlay Opflex agent is enabled on DCI to download the VRF/BD/BDI/VNI/BGP configuration from the spine
3.2.1.2 Security		
1	GETVPN GM	<ul style="list-style-type: none"> GETVPN gdoi crypto map is applied on the WAN vrf-lite interface
3.2.1.3 Network Services		
3	QoS	<ul style="list-style-type: none"> QoS policy with policing for different dscp values is applied on the BDI interface for both input and output QoS is also applied on the vrf-lite WAN interface
4	NAT	<ul style="list-style-type: none"> Static NAT is applied on the BDI and vrf-lite crypto interface
3.2.1.4 Simplified Management		
6	Monitoring	<ul style="list-style-type: none"> Exporting and monitoring logs from the syslog server
3.2.1.5 System Health Monitoring		
7	System Health	Monitor system health for CPU usage, memory consumption and memory leaks during longevity
3.2.1.6 System & Network Resiliency, Robustness		
8	System Resiliency	Verify system level resiliency during the following events, <ul style="list-style-type: none"> Active/Standby ESP failure WAN/DCI/BDI/NVE Interface flaps SIP/SPA reload/OIR Active/Standby RP failure One-shot ISSU
9	Negative Events, Triggers	Verify that the system holds good and recovers to working condition after the following negative events are triggered. <ul style="list-style-type: none"> Config Changes - Add/Remove config snippets, config replace Routing protocol Interface Flaps QoS events like adding/removing QoS policy, modifying the ACL, modifying the class map

3.2.2. ACI EVPN L3 DCI as MPLS PE

No.	Focus Area	Use Cases
3.2.2.1 Routing		
5	BGP EVPN	<ul style="list-style-type: none"> BGP EVPN is configured on the ACI side OSPF is used for the underlay Opflex agent is enabled on DCI to download the VRF/BD/BDI/VNI/BGP configuration from the spine
6	PE	<ul style="list-style-type: none"> DCI is configured as MPLS PE on the WAN side BGP EVPN prefixes are imported into the vpnv4 and vice versa on the DCI
3.2.2.2 Network Services		
3	QoS	<ul style="list-style-type: none"> QoS policy with policing for different dscp values is applied on the BDI interface for both input and output
3.2.2.3 Simplified Management		
6	Monitoring	<ul style="list-style-type: none"> Exporting and monitoring logs from the syslog server
3.2.2.4 System Health Monitoring		
7	System Health	Monitor system health for CPU usage, memory consumption and memory leaks during longevity

No.	Focus Area	Use Cases
3.2.2.5 System & Network Resiliency, Robustness		
8	System Resiliency	Verify system level resiliency during the following events, <ul style="list-style-type: none"> • Active/Standby ESP failure • WAN/DCI/BDI/NVE Interface flaps • SIP/SPA reload/OIR • Active/Standby RP failure • One-shot ISSU
9	Negative Events, Triggers	Verify that the system holds good and recovers to working condition after the following negative events are triggered. <ul style="list-style-type: none"> • Config Changes - Add/Remove config snippets, config replace • Routing protocol Interface Flaps • QoS events like adding/removing QoS policy, modifying the ACL, modifying the class map

3.2.3. ACI EVPN L3 DCI vrf-lite over MPLS with IPN support

No.	Focus Area	Use Cases
3.2.3.1 Routing		
5	BGP EVPN	<ul style="list-style-type: none"> • BGP EVPN is configured on the ACI side • OSPF is used for the underlay • Opflex agent is enabled on DCI to download the VRF/BD/BDI/VNI/BGP configuration from the spine
6	OSPF	<ul style="list-style-type: none"> • IPN is configured in standalone mode • OSPF is enabled on the interface connecting the Spine and DCI
3.2.3.2 Network Services		
3	QoS	<ul style="list-style-type: none"> • QoS policy with policing for different dscp values is applied on the BDI interface for both input and output
3.2.3.4 Simplified Management		
6	Monitoring	<ul style="list-style-type: none"> • Exporting and monitoring logs from the syslog server
3.2.3.5 System Health Monitoring		
7	System Health	Monitor system health for CPU usage, memory consumption and memory leaks during longevity
3.2.3.6 System & Network Resiliency, Robustness		
8	System Resiliency	Verify system level resiliency during the following events, <ul style="list-style-type: none"> • Active/Standby ESP failure • WAN/DCI/BDI/NVE Interface flaps • SIP/SPA reload/OIR • Active/Standby RP failure • One-shot ISSU
9	Negative Events, Triggers	Verify that the system holds good and recovers to working condition after the following negative events are triggered. <ul style="list-style-type: none"> • Config Changes - Add/Remove config snippets, config replace • Routing protocol Interface Flaps • QoS events like adding/removing QoS policy, modifying the ACL, modifying the class map

3.2.4. ASR1K SDA Border with ACI and Policy Plane Integration

No.	Focus Area	Use Cases
3.2.4.1 Routing		
1	BGP	<ul style="list-style-type: none"> eBGP is used on ASR1K towards ACI fabric BGP is used to learn the prefixes from ACI side. Route-leaking between VRFs and common VRF happens on Border Leaf. Border Leaf sends all the ACI prefixes to ASR1K on a single VRF BGP is configured again between ASR1K Border and MS/MR BGP prefixes learnt from ACI side are imported to LISP on the Border
2	LISP	<ul style="list-style-type: none"> LISP+Vxlan is configured on the ASR1K Border and Edge BGP prefixes are imported into LISP on ASR1K Border LISP extranet is used on MS/MR to provide the connectivity between Border and Edge ACI prefixes are configured as Providers and Edge side prefixes are configured as Subscribers Verify connectivity is established and traffic between the hosts connected to Edge and a host connected to the N9K ACI Leaf goes through fine
3.2.4.2 Security		
2	TrustSec	<ul style="list-style-type: none"> Policy-Plane integration between ISE and APIC TrustSec static mappings ie ip to SGT, ip subnet to SGT and L3 if is enabled on Edge Dynamic SGT mappings are learnt from ISE and propagated to the ACI fabric SXP between ISE and APIC. ISE learns the EPGs and assigns SGTs SGT-IP bindings on ISE These mappings are propagated to APIC as networks under the I3out of common tenant. Verify the routes are propagated to the Leaf from the border leaf SGT inline is enabled for LISP Verify traffic is working end-to-end Apply contracts on these networks on APIC for policy enforcement
3.2.4.3 Network Services		
3	QoS	<ul style="list-style-type: none"> QoS policy with policing for different dscp values is applied on the BDI interface for both input and output
3.2.4.4 Monitoring & Troubleshooting		
5	FNF	<ul style="list-style-type: none"> FNF is configured on the interface to match the source and destination tag. Verify traffic is going through with the right tag
3.2.4.5 Simplified Management		
6	Monitoring	<ul style="list-style-type: none"> Exporting and monitoring logs from the syslog server
3.2.4.6 System Health Monitoring		
7	System Health	Monitor system health for CPU usage, memory consumption and memory leaks during longevity
3.2.4.7 System & Network Resiliency, Robustness		
8	System Resiliency	Verify system level resiliency during the following events, <ul style="list-style-type: none"> Active/Standby ESP failure WAN/DCI/BDI/NVE/LISP Interface flaps SIP/SPA reload/OIR
9	Negative Events, Triggers	Verify that the sytem holds good and recovers to working condition after the following negative events are triggered. <ul style="list-style-type: none"> Config Changes - Add/Remove config snippets, config replace Routing protocol Interface Flaps LISP events like flapping extranet policy, reloading MS/MR, modifying extranet prefix-list. TrustSec events like changing the static sgt values, adding and deleting EPG on APIC, adding and deleting SGT on ISE

4. Appendix A

Disclaimer

Below are some sample configuration snippets to give a general idea about the configuration used in some of the use-cases and would require further customization for actual deployments. For detailed configuration options/ best practices, please refer to the CCO documentation.

Cisco Reference Links:

VXLAN –GPO IETF draft-<https://tools.ietf.org/html/draft-smith-vxlan-group-policy-03>

<https://www.cisco.com/c/en/us/solutions/enterprise-networks/trustsec/solution-overview-listing.html>.

5. Acronyms

Here is the list of Acronyms used in the CVP

- ACI = Application Centric Infrastructure
- ACI Domain = A data center operating in ACI mode
- ISE=Identity Services Engine
- EPG = End Point Group, a group assigned to a host in ACI
- SG = Scalable Group, a group assigned to a host in TrustSec domain
- ClassId = Refers to a 16-bit number that identifies an EPG
- SGT = Refers to a 16-bit number that identifies an SG
- Tenant = A specific instance of policy enforcement (e.g. customer in a service provider setup)
- Vxlan= Virtual Extensible LAN (**VXLAN**) is a network virtualization technology that attempts to improve the scalability problems associated with large cloud computing deployments
- VNI= VXLAN Network Identifier .Each VXLAN segment is identified through a 24-bit segment ID

For any feedback/questions, please send an email to: eas-cvp-feedback@cisco.com

6. Configuration

6.1 Usecase: ACI EVPN L3 DCI vrf-lite over MPLS Secure WAN

The following section describes the configuration needed for the deployment of the solution

Topology:

ACI Fabric –(Fo1/1/1.4)DCI(Te2/1/0.2) – (Te1/3/0.2)PE1(Te1/1/0) – P – PE2 –
Branch – Traffic

DCI Configuration:

Configure interface going to ACI fabric and OSPF underlay:

```
Interface FortyGigabitEthernet1/1/1
  description "Connected to Spine"
  no ip address
!
interface FortyGigabitEthernet1/1/1.4
  encapsulation dot1Q 4
```

```
ip address 88.0.0.4 255.255.255.0
ip ospf mtu-ignore
!
interface Loopback0
 ip address 31.1.1.1 255.255.255.255
!
router ospf 100
 nsf ietf
 area 1 nssa
 area 100 nssa
 network 31.1.1.1 0.0.0.0 area 0
 network 88.0.0.0 0.0.0.255 area 100
!
```

Configure BGP EVPN neighbor:

Configure BGP on ASR1K DCI with evpn neighbor as N9K Spine loopback IP.

```
router bgp 101
 bgp router-id 31.1.1.1
 bgp log-neighbor-changes
 bgp listen limit 5000
 bgp graceful-restart
 timers bgp 120 360
 neighbor 102.102.102.102 remote-as 100 neighbor
 102.102.102.102 ebgp-multihop 255 neighbor
 102.102.102.102 update-source Loopback0 neighbor
 102.102.102.102 ha-mode graceful-restart
!
address-family ipv4
 neighbor 102.102.102.102 activate
exit-address-family
!
address-family l2vpn evpn
 import vpv4 unicast re-originate
 neighbor 102.102.102.102 activate
 neighbor 102.102.102.102 send-community both
exit-address-family
!
```

Configure NVE interface:

Configure the NVE interface with source as loopback, host-reachability protocol as bgp.

Set the vxlan udp port to 48879 for ACI mode.

```
vxlan udp port 48879
!
interface nve1
  no ip address
  source-interface Loopback0
  host-reachability protocol bgp
  vxlan udp port 48879
!
```

Configure opflex:

Configure ASR1k as opflex agent so that the VRF/BDI/BD/VNI data is pushed from N9K Spine which is the opflex server.

Peer IP address is the connected sub-interface ip on spine. Source IP address is the vlan-4 sub-interface IP address on ASR1K DCI.

Identity should be configured as dci-[Loopback IP address of ASR1K DCI]

```
opflex agent
  service vxlan-evpn
  nve-id 1
  bdi-ip 100.1.1.1 255.255.255.0
  domain DC1
  identity dci-[31.1.1.1]
  peer 1 ip-address 88.0.0.3 tcp-port 8009 src-ip-address 88.0.0.4
```

Configure P/PE:

Configure the vrf-lite interface for cust2 on DCI:

```
interface TenGigabitEthernet2/1/0.2
  description "Connection to MPLS PE1"
  encapsulation dot1Q 3
  vrf forwarding cust2
  ip address 30.1.2.1 255.255.255.0
!
```

```
router bgp 101
  address-family ipv4 vrf cust2
  advertise l2vpn evpn
  redistribute connected
  neighbor 30.1.2.2 remote-as 102
  neighbor 30.1.2.2 ebgp-multihop 255
  neighbor 30.1.2.2 activate
  exit-address-family
```

!

Configure MPLS PE1:

```
vrf definition cust2
  rd 101:2
  !
  address-family ipv4
    route-target export 101:2
    route-target export 200:2
    route-target import 101:2
    route-target import 200:1
  exit-address-family
  !

interface TenGigabitEthernet1/3/0.2
  description "Connection to DCI"
  encapsulation dot1Q 3
  vrf forwarding cust2
  ip address 30.1.2.2 255.255.255.0
  !
interface Loopback0
  ip address 10.1.1.1 255.255.255.255
  !
interface TenGigabitEthernet1/1/0
  ip address 10.2.1.1 255.255.255.0
  mpls ip
  mpls label protocol ldp
  !
mpls label protocol ldp
mpls ldp router-id Loopback0 force
  !

router ospf 100
  router-id 10.1.1.1
  network 10.1.1.1 0.0.0.0 area 0
  network 10.2.1.0 0.0.0.255 area 0
  network 10.3.1.0 0.0.0.255 area 0
  network 10.4.1.0 0.0.0.255 area 0
  !

router bgp 102
  no bgp log-neighbor-changes
  neighbor INTERNAL peer-group
  neighbor INTERNAL remote-as 102
  neighbor INTERNAL update-source Loopback0
```

```

neighbor 10.1.1.2 peer-group INTERNAL
neighbor 10.1.1.2 update-source Loopback0
!
address-family ipv4
  neighbor 10.1.1.2 activate
exit-address-family
!
address-family vpnv4
  neighbor INTERNAL send-community both
  neighbor INTERNAL next-hop-self
  neighbor 10.1.1.2 activate
exit-address-family
!
address-family ipv4 vrf cust2
  redistribute connected
  neighbor 30.1.2.1 remote-as 101
  neighbor 30.1.2.1 ebgp-multihop 255
  neighbor 30.1.2.1 activate
exit-address-family
!

```

Configure Core Device :

```

interface TenGigabitEthernet2/1/0
  description "Connecting PE2"
  ip address 10.2.2.1 255.255.255.0
  mpls ip
  mpls label protocol ldp
!
interface TenGigabitEthernet2/3/0
  descriptin "Connecting PE1"
  ip address 10.2.1.2 255.255.255.0
  mpls ip
  mpls label protocol ldp
!
interface Loopback0
  ip address 10.1.1.2 255.255.255.255
!
vrf definition KS-VRF1
  description ## V4 VRF for Cust : KS-VRF1 ##
  rd 100:1
!
address-family ipv4
  import map route-filter
  route-target export 200:1
  route-target import 200:2

```

```

    exit-address-family
  !
interface GigabitEthernet2/0/5
  description "Connecting KeyServer"
  vrf forwarding KS-VRF1
  ip address 15.1.1.1 255.255.255.0
  negotiation auto
  !
router ospf 200 vrf KS-VRF1
  network 13.1.1.0 0.0.0.255 area 0
  network 15.1.1.0 0.0.0.255 area 0
  network 16.1.1.0 0.0.0.255 area 0
router ospf 100
  router-id 10.1.1.2
  network 10.1.1.2 0.0.0.0 area 0
  network 10.2.1.0 0.0.0.255 area 0
  network 10.2.2.0 0.0.0.255 area 0
  network 10.2.4.0 0.0.0.255 area 0
  !
router bgp 102
  bgp log-neighbor-changes
  neighbor INTERNAL peer-group
  neighbor INTERNAL remote-as 102
  neighbor INTERNAL update-source Loopback0
  neighbor 10.1.1.1 peer-group INTERNAL
  neighbor 10.1.1.3 peer-group INTERNAL
  neighbor 10.2.4.2 remote-as 102
  neighbor 10.2.4.2 peer-group INTERNAL
  !
address-family ipv4
  neighbor 10.1.1.1 activate
  neighbor 10.1.1.3 activate
  neighbor 10.2.4.2 activate
exit-address-family
  !
address-family vpv4
  neighbor INTERNAL send-community both
  neighbor INTERNAL route-reflector-client
  neighbor 10.1.1.1 activate
  neighbor 10.1.1.3 activate
  neighbor 10.2.4.2 activate
exit-address-family
  !
address-family ipv4 vrf KS-VRF1
  redistribute ospf 200

```

```
exit-address-family
!
mpls label protocol ldp
mpls ldp router-id Loopback0 force
!
```

Configure MPLS PE2:

```
interface Loopback0
 ip address 10.1.1.3 255.255.255.255
!
interface TenGigabitEthernet0/2/0
 description "##Connected to P "
 ip address 10.2.2.2 255.255.255.0
 mpls ip
 mpls label protocol ldp
!
router ospf 100
 router-id 10.1.1.3
 network 10.1.1.3 0.0.0.0 area 0
 network 10.2.2.0 0.0.0.255 area 0
!
vrf definition cust2
 description ## Dual-Stack V4V6 VRF for Cust : cust2 ##
 rd 100:2
!
 address-family ipv4
  route-target export 100:2
  route-target export 200:2
  route-target export 101:2
  route-target import 100:2
  route-target import 200:1
  route-target import 101:2
 exit-address-family
!
mpls label protocol ldp
mpls ldp router-id Loopback0 force
!
interface TenGigabitEthernet0/3/0.2
 description "Connecting Branch"
 encapsulation dot1Q 3
 vrf forwarding cust2
 ip address 60.1.2.2 255.255.255.0

!
router bgp 102
```

```

no bgp log-neighbor-changes
neighbor INTERNAL peer-group
neighbor INTERNAL remote-as 102
neighbor INTERNAL update-source Loopback0
neighbor 10.1.1.2 peer-group INTERNAL
neighbor 10.1.1.2 update-source Loopback0
!
address-family ipv4
neighbor 10.1.1.2 activate
exit-address-family
!
address-family vpnv4
neighbor INTERNAL send-community both
neighbor INTERNAL next-hop-self
neighbor 10.1.1.2 activate
exit-address-family
!
address-family ipv4 vrf cust2
redistribute connected
neighbor 60.1.2.1 remote-as 103
neighbor 60.1.2.1 ebgp-multihop 255
neighbor 60.1.2.1 activate
exit-address-family
!

```

Configure Branch

Branch Config:

```

vrf definition cust2
description ## V4 VRF for Cust : cust2 ##
rd 103:2
!
address-family ipv4
route-target export 103:2
route-target import 103:2
exit-address-family
!
interface TenGigabitEthernet0/1/2.2
description "Connecting PE2"
encapsulation dot1Q 3
vrf forwarding cust2
ip address 60.1.2.1 255.255.255.0
!
router bgp 103
no bgp log-neighbor-changes
!

```

```

address-family ipv4 vrf cust2
  redistribute connected
  neighbor 60.1.2.2 remote-as 102
  neighbor 60.1.2.2 ebgp-multihop 255
  neighbor 60.1.2.2 activate
exit-address-family
!
interface TenGigabitEthernet0/1/5.2
  description "Traffic interface"
  encapsulation dot1q 3
  vrf forwarding cust2
  ip address 64.1.2.1 255.255.255.0
!

```

GETVPN Config:

KeyServer is connected to the Core P router.

KeyServer Config:

```

interface GigabitEthernet0/1
  description "Connected to Core router"
  ip address 15.1.1.2 255.255.255.0
  no ip redirects
  no ip proxy-arp
  media-type sfp
!
interface Loopback0
  ip address 15.0.0.1 255.255.255.255
!
ip access-list extended bw600-crypto-policy
deny ip host 15.0.0.1 40.0.0.0 0.255.255.255
deny ip 40.0.0.0 0.255.255.255 host 15.0.0.1
deny ip host 16.0.0.1 40.0.0.0 0.255.255.255
deny ip 40.0.0.0 0.255.255.255 host 16.0.0.1
deny ip 50.0.0.0 0.255.255.255 host 16.0.0.1
deny ip host 16.0.0.1 50.0.0.0 0.255.255.255
deny ip 50.0.0.0 0.255.255.255 host 15.0.0.1
deny ip host 15.0.0.1 50.0.0.0 0.255.255.255
deny ip host 15.0.0.1 60.0.0.0 0.255.255.255
deny ip 60.0.0.0 0.255.255.255 host 15.0.0.1
deny ip host 16.0.0.1 60.0.0.0 0.255.255.255
deny ip 60.0.0.0 0.255.255.255 host 16.0.0.1
deny ip host 15.0.0.1 70.0.0.0 0.255.255.255
deny ip 70.0.0.0 0.255.255.255 host 15.0.0.1
deny ip host 16.0.0.1 70.0.0.0 0.255.255.255
deny ip host 14.1.1.1 50.0.0.0 0.255.255.255

```

```

deny ip host 14.1.1.1 40.0.0.0 0.255.255.255
deny ip host 14.1.1.1 60.0.0.0 0.255.255.255
deny ip host 14.1.1.1 70.0.0.0 0.255.255.255
deny ip 40.0.0.0 0.255.255.255 host 14.1.1.1
deny ip 50.0.0.0 0.255.255.255 host 14.1.1.1
deny ip 60.0.0.0 0.255.255.255 host 14.1.1.1
deny tcp any any eq bgp
deny udp any any eq 848
deny udp any any eq isakmp
deny tcp any eq bgp any
deny udp any eq 848 any
deny udp any eq isakmp any
deny udp any any eq ntp
deny udp any eq ntp any
permit ip any any
!
crypto gdoi group bw600
identity number 600
server local

rekey algorithm aes 256
rekey retransmit 10 number 3
rekey authentication mypubkey rsa KeyServer.cisco.com
rekey transport unicast
sa ipsec 10
profile bw600
match address ipv4 bw600-crypto-policy
replay time window-size 50
no tag
address ipv4 15.0.0.1
!
crypto isakmp policy 10
encr aes 256
authentication pre-share
group 2
!
crypto isakmp policy 20
encr aes 256
group 2
lifetime 60
crypto isakmp key KS-cisco address 0.0.0.0
crypto isakmp identity dn
crypto isakmp keepalive 10 periodic
!

```



```
crypto ipsec transform-set bw600-cryptoset esp-aes 256 esp-sha-hmac
 mode tunnel
!
crypto ipsec profile bw600
 set security-association lifetime seconds 7200
 set transform-set bw600-cryptoset
!
ip route 0.0.0.0 0.0.0.0 15.1.1.1
!
```

DCI Config for Group Member:

```
crypto isakmp policy 10
 encr aes 256
 authentication pre-share
 group 2
!
crypto isakmp keepalive 10 3
crypto isakmp profile ikev1
 keyring key
 match identity address 0.0.0.0
!
crypto gdoi group bw600
 identity number 600
 server address ipv4 15.0.0.1
!
crypto map getvpn-bw600 10 gdoi
 set group bw600
 crypto map getvpn-bw600
!
interface TenGigabitEthernet2/1/0.2
 description "Connected to PE1"
 encapsulation dot1Q 3
 vrf forwarding cust2
 ip address 30.1.2.1 255.255.255.0
 crypto map getvpn-bw600
!
```

Branch Configuration for Group Member:

```
crypto isakmp policy 10
 encr aes 256
 authentication pre-share
 group 2
!
crypto gdoi group bw600
 identity number 600
```

```
server address ipv4 15.0.0.1
!
crypto map getvpn-bw600 10 gdoi
set group bw600
crypto map getvpn-bw600
!
interface TenGigabitEthernet0/1/2.2
encapsulation dot1Q 3
vrf forwarding cust2
ip address 60.1.2.1 255.255.255.0
crypto map getvpn-bw600
!
```

Apply Qos on the GETVPN interface on DCI:

```
class-map match-any gold-cust
match ip dscp cs4
class-map match-any premium-cust
match ip dscp cs6
class-map match-any silver-cust
match ip dscp cs2
!
policy-map grandchild-cust
class premium-cust
police cir percent 50
class gold-cust
police cir percent 10
class silver-cust
police cir percent 35
!
interface TenGigabitEthernet2/1/0.2
encapsulation dot1Q 3
vrf forwarding cust2
ip address 30.1.2.1 255.255.255.0
crypto map getvpn-bw600
service-policy input grandchild-cust
service-policy output grandchild-cust
!
```

Configuration

6.2 Usecase: ACI EVPN L3DCI as PE

The following section describes the configuration needed for the deployment of the solution.

BGP EVPN/Opflex Bringup on ASR1K

ACI Fabric --- DCI ---- P1 --- P --- PE2 --- Traffic

DCI configuration

```
interface FortyGigabitEthernet1/1/1.4
  description "Connection to the spine"
  encapsulation dot1Q 4
  ip address 88.0.0.4 255.255.255.0
  ip ospf mtu-ignore
!
interface Loopback0
  ip address 31.1.1.1 255.255.255.255
!
interface TenGigabitEthernet2/1/0
  description "Connection to Core P router"
  ip address 10.3.1.1 255.255.255.0
  mpls ip
  mpls label protocol ldp
!
router ospf 100
  nsf ietf
  area 1 nssa
  area 100 nssa
  network 10.3.1.0 0.0.0.255 area 0
  network 31.1.1.1 0.0.0.0 area 0
  network 88.0.0.0 0.0.0.255 area 100
  network 89.0.0.0 0.0.0.255 area 100
  network 90.1.1.0 0.0.0.255 area 100
!
```

Configure BGP EVPN neighbor:

```
mpls ldp router-id Loopback0 force
!
vrf definition cust1
  rd 1:1001
!
address-family ipv4
  route-target export 100:1
  route-target import 100:1
  route-target import 101:1
  route-target export 100:1 stitching
```

```

    route-target import 100:1 stitching
exit-address-family
!
router bgp 101
  bgp router-id 31.1.1.1
  bgp log-neighbor-changes
  bgp listen limit 5000
  bgp graceful-restart
  timers bgp 120 360
  neighbor INTERNAL peer-group
  neighbor INTERNAL remote-as 101
  neighbor INTERNAL update-source Loopback0
  neighbor 10.1.1.2 peer-group INTERNAL → MPLS P neighbor
  neighbor 10.1.1.2 update-source Loopback0
  neighbor 102.102.102.102 remote-as 100 → EVPN Neighbor
  neighbor 102.102.102.102 ebgp-multihop 255 neighbor
  neighbor 102.102.102.102 update-source Loopback0 neighbor
  neighbor 102.102.102.102 ha-mode graceful-restart
!
  address-family ipv4
    neighbor 10.1.1.2 activate
    neighbor 102.102.102.102 activate
  exit-address-family
!
  address-family vpnv4
    import l2vpn evpn re-originate
    neighbor INTERNAL send-community both
    neighbor INTERNAL next-hop-self
    neighbor 10.1.1.2 activate
  exit-address-family
!
  address-family l2vpn evpn
    import vpnv4 unicast re-originate
    neighbor 102.102.102.102 activate
    neighbor 102.102.102.102 send-community both
  exit-address-family
!

```

Configure opflex:

Configure ASR1k as opflex agent so that the VRF/BDI/BD/VNI data is pushed from N9K Spine which is the opflex server.

Peer IP address is the connected sub-interface ip on spine. Source IP address is the vlan-4 sub-interface IP address on ASR1K DCI.

Identity should be configured as dci-[Loopback IP address of ASR1K DCI]

```
opflex agent
service vxlan-evpn
  nve-id 1
  bdi-ip 100.1.1.1 255.255.255.0
domain DC1
  identity dci-[31.1.1.1]
  peer 1 ip-address 88.0.0.3 tcp-port 8009 src-ip-address 88.0.0.4
```

Configure NVE interface:

Configure the NVE interface with source as loopback, host-reachability protocol as bgp.

Set the vxlan udp port to 48879 for ACI mode.

```
vxlan udp port 48879
!
interface nve1
  no ip address
  source-interface Loopback0
  host-reachability protocol bgp
  vxlan udp port 48879
!
```

P1 Configuration

```
interface TenGigabitEthernet1/3/0
  description "Connecting DCI PE"
  ip address 10.3.1.2 255.255.255.0
  mpls ip
  mpls label protocol ldp
!
interface TenGigabitEthernet1/1/0
  description "Connecting P"
  ip address 10.2.1.1 255.255.255.0
  mpls ip
  mpls label protocol ldp
!
mpls ldp router-id Loopback0 force
!
interface Loopback0
  ip address 10.1.1.1 255.255.255.255
!
router ospf 100
  router-id 10.1.1.1
  network 10.1.1.1 0.0.0.0 area 0
  network 10.2.1.0 0.0.0.255 area 0
  network 10.3.1.0 0.0.0.255 area 0
```

```
network 10.4.1.0 0.0.0.255 area 0
!
```

P Configuration

```
interface Loopback0
 ip address 10.1.1.2 255.255.255.255
!
interface TenGigabitEthernet2/1/0
 description "Connecting PE2"
 ip address 10.2.2.1 255.255.255.0
 mpls ip
 mpls label protocol ldp
!
interface TenGigabitEthernet2/3/0
 description "Connecting P1"
 ip address 10.2.1.2 255.255.255.0
 mpls ip
 mpls label protocol ldp
!
mpls ldp router-id Loopback0 force
!
router ospf 100
 router-id 10.1.1.2
 network 10.1.1.2 0.0.0.0 area 0
 network 10.2.1.0 0.0.0.255 area 0
 network 10.2.2.0 0.0.0.255 area 0
 network 10.2.4.0 0.0.0.255 area 0
!
router bgp 101
 bgp log-neighbor-changes
 neighbor INTERNAL peer-group
 neighbor INTERNAL remote-as 101
 neighbor INTERNAL update-source Loopback0
neighbor 10.1.1.3 peer-group INTERNAL → PE2 neighbor
neighbor 31.1.1.1 peer-group INTERNAL → DCI PE
!
address-family ipv4
 neighbor 10.1.1.3 activate
 neighbor 31.1.1.1 activate
 exit-address-family
!
address-family vpnv4
 neighbor INTERNAL send-community both
 neighbor INTERNAL route-reflector-client
 neighbor 10.1.1.3 activate
```

```
neighbor 31.1.1.1 activate
exit-address-family
!
```

PE2 Configuration

```
vrf definition cust1
  rd 101:1
  !
  address-family ipv4
    route-target export 101:1
    route-target import 101:1
  exit-address-family
  !
interface TenGigabitEthernet0/3/0.1
  description "Connecting the Branch"
  encapsulation dot1Q 2
  vrf forwarding cust1
  ip address 60.1.1.2 255.255.255.0
  !
interface TenGigabitEthernet0/2/0
  description "Connected to P "
  ip address 10.2.2.2 255.255.255.0
  mpls ip
  mpls label protocol ldp
  !
interface Loopback0
  ip address 10.1.1.3 255.255.255.255
  !
router ospf 100
  router-id 10.1.1.3
  network 10.1.1.3 0.0.0.0 area 0
  network 10.2.2.0 0.0.0.255 area 0
  !
router bgp 101
  no bgp log-neighbor-changes
  neighbor INTERNAL peer-group
  neighbor INTERNAL remote-as 101
  neighbor INTERNAL update-source Loopback0
  neighbor 10.1.1.2 peer-group INTERNAL
  neighbor 10.1.1.2 update-source Loopback0
  !
address-family ipv4
  neighbor 10.1.1.2 activate
exit-address-family
!
```

```
address-family vpnv4
  neighbor INTERNAL send-community both
  neighbor INTERNAL next-hop-self
  neighbor 10.1.1.2 activate
exit-address-family
!
address-family ipv4 vrf cust1
  redistribute connected
  neighbor 60.1.1.1 remote-as 103
  neighbor 60.1.1.1 ebgp-multihop 255
  neighbor 60.1.1.1 activate
exit-address-family
!
```

Branch Configuration

```
vrf definition cust1
  description ## V4 VRF for Cust : cust1 ##
  rd 103:1
!
address-family ipv4
  route-target export 103:1
  route-target import 103:1
!
interface TenGigabitEthernet0/1/2.1
  description "Connected to PE2"
  encapsulation dot1Q 2
  vrf forwarding cust1
  ip address 60.1.1.1 255.255.255.0
!
interface TenGigabitEthernet0/1/7.1
  description "Connected to Traffic"
  encapsulation dot1Q 2
  vrf forwarding cust1
  ip address 65.1.1.2 255.255.255.0
!
interface Loopback0
  ip address 61.1.1.1 255.255.255.255
!
router bgp 103
  no bgp log-neighbor-changes
!
address-family ipv4 vrf cust1
  redistribute connected
  neighbor 60.1.1.2 remote-as 101
  neighbor 60.1.1.2 ebgp-multihop 255
```



```
neighbor 60.1.1.2 activate
exit-address-family
!
```

Configuration

6.3 Usecase: ACI EVPN L3DCI vrf-lite over MPLS with IPN support

The following section describes the configuration needed for the deployment of the solution

BGP EVPN/Opflex Bringup on ASR1K

ACI Leaf --- ACI Spine --- (eth2/6.4) IPN (eth1/48) --- (Te2/2/0) DCI

DCI Configuration:

```
interface TenGigabitEthernet2/2/0
description "connected to IPN"
ip address 90.1.1.1 255.255.255.0
ip ospf mtu-ignore
!
interface Loopback0
ip address 31.1.1.1 255.255.255.255
!
router ospf 100
nsf ietf
area 1 nssa
area 100 nssa
network 31.1.1.1 0.0.0.0 area 0
network 88.0.0.0 0.0.0.255 area 100

network 89.0.0.0 0.0.0.255 area 100
network 90.1.1.0 0.0.0.255 area 100
!
```

Configure BGP EVPN neighbor:

Configure BGP on ASR1K DCI with evpn neighbor as N9K Spine loopback IP. "import vpv4

unicast re-originate" is needed only if the DCI acts as a PE.

The BGP EVPN control plane in the VXLAN BGP EVPN fabric ensures distribution of routes between ToR/leaf switch VTEPs within the fabric. ToRs will forward the attached end host IP and Layer-3 VXLAN VNIs using the EVPN Route Type 5 option. Based on the import route target (RT) configured on the border leaf switch, the switch will import the /32 routes into appropriate VRF tables.

The GOLF router uses the “import vpv4 unicast re-originate” cli to re-originate the vpv4 WAN routes as EVPN routes towards the ACI . “advertise l2vpn evpn” is used to re-originate the EVPN routes from ACI fabric to l3vpn address-family on the WAN side.

```
router bgp 101
  bgp router-id 31.1.1.1
  bgp log-neighbor-changes
  bgp listen limit 5000
  bgp graceful-restart
  timers bgp 120 360
  neighbor 102.102.102.102 remote-as 100
  neighbor 102.102.102.102 ebgp-multihop 255
  neighbor 102.102.102.102 update-source Loopback0
  neighbor 102.102.102.102 ha-mode graceful-restart
!
  address-family ipv4
    neighbor 102.102.102.102 activate
  exit-address-family
!
  address-family l2vpn evpn
    import vpv4 unicast re-originate
    neighbor 102.102.102.102 activate
    neighbor 102.102.102.102 send-community both
  exit-address-family
!
```

Configure NVE interface:

Configure the NVE interface with source as loopback, host-reachability protocol as bgp.

Set the vxlan udp port to 48879 for ACI mode.

```
vxlan udp port 48879
!
interface nve1
  no ip address
  source-interface Loopback0
  host-reachability protocol bgp
  vxlan udp port 48879
!
```

Configure opflex:

Configure ASR1k as opflex agent so that the VRF/BDI/BD/VNI data is pushed from N9K Spine which is the opflex server.

Peer IP address is the connected sub-interface ip on spine. Source IP address is the vlan-4 sub-interface IP address on ASR1K DCI.

Identity should be configured as dci-[Loopback IP address of ASR1K DCI]

```
opflex agent
service vxlan-evpn
  nve-id 1
  bdi-ip 100.1.1.1 255.255.255.0
domain DC1
  identity dci-[31.1.1.1]
  peer 2 ip-address 86.0.0.3 tcp-port 8009 src-ip-address 90.1.1.1
!
```

N9K IPN Configuration:

```
feature ospf
feature bgp
feature pim
feature interface-vlan
feature vn-segment-vlan-based
feature lacp
feature vpc
feature lldp
feature bfd
feature nv overlay
interface loopback0
  ip address 3.3.3.3/32
!
interface Ethernet1/48 → Connecting ASR1K
  no switchport
  ip address 90.1.1.2/24
  ip router ospf 100 area 0.0.0.100
  no shutdown
!
interface Ethernet2/6.4 → Connection to Spine
  encapsulation dot1q 4
  ip address 86.0.0.4/24
  ip ospf mtu-ignore
  ip router ospf 100 area 0.0.0.100
  no shutdown
!
router ospf 100
  router-id 3.3.3.3
  network 3.3.3.3/32 area 0.0.0.100
  network 86.0.0.0/24 area 0.0.0.100
  network 89.0.0.0/24 area 0.0.0.100
  network 90.1.1.0/24 area 0.0.0.100
  area 0.0.0.100 nssa
!
```

Configuration

6.4 Usecase: ASR1K SDA Border with ACI and Policy Plane Integration

The following section describes the configuration needed for the deployment of the solution

(80.1.100.2) ACI Leaf (30.1.1.2) ---- (30.1.1.1)DCI Border (149.1.1.1) --- MS/MR ---- Edge(72.1.1.2)

Configure BGP on DCI Border with neighbor as N9K Border Leaf

```
vrf definition aci1
  rd 101:1
  route-target export 101:1
  route-target import 101:1
  !
  address-family ipv4
  exit-address-family
  !
interface TenGigabitEthernet2/0/0
  description "Interface connected to N9K Border Leaf"
  vrf forwarding aci1
  ip address 30.1.1.1 255.255.255.0
  !
interface Loopback0
  ip address 50.1.20.1 255.255.255.255
  !
router bgp 101
  bgp router-id 50.1.20.1
  bgp log-neighbor-changes
  neighbor 30.1.1.2 remote-as 108
  neighbor 30.1.1.2 ebgp-multihop 255
  neighbor 30.1.1.2 update-source TenGigabitEthernet2/0/0
  !
  address-family ipv4
    redistribute connected
    neighbor 30.1.1.2 activate
  exit-address-family
  !
  address-family ipv4 vrf aci1
    redistribute connected
    neighbor 30.1.1.2 remote-as 108
    neighbor 30.1.1.2 ebgp-multihop 255
    neighbor 30.1.1.2 update-source TenGigabitEthernet2/0/0
    neighbor 30.1.1.2 activate
  exit-address-family
```

Configure OSPF DCI Border, MS/MR and Edge XTR

DCI1

```
interface GigabitEthernet0/0/1
  description "Connected to MS/MR"
  ip address 149.1.1.1 255.255.255.0
  ip lisp source-locator Loopback0
  negotiation auto
!
interface Loopback0
  ip address 50.1.20.1 255.255.255.255
!
router ospf 101
  network 50.1.20.1 0.0.0.0 area 0
  network 149.1.1.0 0.0.0.255 area 0
!
```

MS/MR

```
interface Loopback0
  ip address 124.1.1.1 255.255.255.0
!
interface GigabitEthernet0/0/0
  description "Connected to DCI Border"
  ip address 149.1.1.2 255.255.255.0
  negotiation auto
!
interface GigabitEthernet0/0/2
  description "connected to Edge/XTR"
  ip address 60.1.1.2 255.255.255.0
  negotiation auto
!
router ospf 101
  network 60.1.1.0 0.0.0.255 area 0
  network 124.1.1.1 0.0.0.0 area 0
  network 149.1.1.0 0.0.0.255 area 0
!
```

Edge/XTR

```
interface Loopback0
  ip address 30.3.1.1 255.255.255.0
!
interface GigabitEthernet0/0/0
  ip address 60.1.1.1 255.255.255.0
  negotiation auto
!
```

```
router ospf 101
 network 30.3.1.1 0.0.0.0 area 0
 network 60.1.1.0 0.0.0.255 area 0
!
```

Configure LISP and LISP extranet to leak routes between ACI and SDA VRFs

Configure DCI as Border:

```
ip community-list 1 permit 655370
!
route-map database deny 10
 match community 1
route-map database permit 20
!
router lisp
 locator-set border
  50.1.20.1 priority 1 weight 100
  IPv4-interface Loopback0 priority 0 weight 10
  auto-discover-rlocs
  exit-locator-set
!
service ipv4
 encapsulation vxlan
 map-cache-limit 10000
 itr map-resolver 124.1.1.1
 itr
 etr map-server 124.1.1.1 key lisp
 etr
 sgt
 exit-service-ipv4
!
instance-id 100
 service ipv4
  eid-table vrf acil
  database-mapping 30.1.1.0/24 50.1.20.1 priority 1 weight 100
  route-import map-cache maximum-prefix 10000
  route-import map-cache bgp 101
  route-import database bgp 101 route-map database locator-set border
  map-cache-limit 10000
  itr map-resolver 124.1.1.1
  itr
  etr map-server 124.1.1.1 key lisp
  etr
  sgt
  exit-service-ipv4
!
```

```
    exit-instance-id
  !
  exit-router-lisp

vrf definition acil
  rd 101:1
  route-target export 101:1
  route-target import 101:1
  route-target import 201:1 → Import sda1 RT
  route-target import 202:1 → Import sda2 RT
  !
  address-family ipv4
  exit-address-family
!
```

Configure MS/MR

```
vrf definition acil
  rd 101:1
  route-target export 101:1
  route-target import 101:1
  route-target import 201:1 → Import sda1 RT
  route-target export 202:1 → Import sda2 RT
  !
  address-family ipv4
  exit-address-family
!

vrf definition sda1
  rd 201:1
  route-target export 201:1
  route-target import 201:1
  !
  address-family ipv4
  exit-address-family
!

vrf definition sda2
  rd 202:1
  route-target export 202:1
  route-target import 202:1
  !
  address-family ipv4
  exit-address-family
!

router lisp
```

```
locator-set msmr
  IPv4-interface Loopback0 priority 1 weight 100
  exit-locator-set
!
service ipv4
  map-server
  map-resolver
  exit-service-ipv4
!
instance-id 1
  service ipv4
    eid-table vrf sda1
    route-export site-registrations
    distance site-registrations 10
    exit-service-ipv4
  !
  exit-instance-id
!
instance-id 2
  service ipv4
    eid-table vrf sda2
    route-export site-registrations
    distance site-registrations 10
    exit-service-ipv4
  !
  exit-instance-id
!
instance-id 100
  service ipv4
    eid-table vrf aci1
    distance site-registrations 10
    exit-service-ipv4
  !
  exit-instance-id
!
site border
  authentication-key lisp
  exit-site
!
site edge
  authentication-key lisp
  eid-record instance-id 1 72.1.1.0/24 accept-more-specifics
  eid-record instance-id 1 72.1.2.0/24
  eid-record instance-id 2 72.1.3.0/24 accept-more-specifics
  eid-record instance-id 2 72.1.4.0/24
```



```

eid-record instance-id 100 80.1.100.0/24 accept-more-specifics
eid-record instance-id 100 80.1.101.0/24 accept-more-specifics
exit-site
!
extranet ex
eid-record-provider instance-id 100 30.1.1.0/24 bidirectional
eid-record-provider instance-id 100 80.1.100.0/24 bidirectional
eid-record-provider instance-id 100 80.1.101.0/24 bidirectional
eid-record-subscriber instance-id 1 72.1.1.0/24 bidirectional
eid-record-subscriber instance-id 2 72.1.3.0/24 bidirectional
eid-record-subscriber instance-id 11 72.1.11.0/24 bidirectional
exit-extranet
!
exit-router-lisp
!
route-map tag permit 10
set community 655370
!
router bgp 101
bgp router-id 124.1.1.1
bgp log-neighbor-changes
neighbor 149.1.1.1 remote-as 101
!
address-family vpnv4
neighbor 149.1.1.1 activate
neighbor 149.1.1.1 send-community both
neighbor 149.1.1.1 route-map tag out
exit-address-family
!
address-family ipv4 vrf sda1
redistribute lisp
exit-address-family
!
address-family ipv4 vrf sda2
redistribute lisp
exit-address-family
!

```

Configure Edge as an XTR

```

vrf definition sda1
rd 201:1
route-target export 201:1
route-target import 201:1
!

```

```

address-family ipv4
exit-address-family
!
vrf definition sda2
rd 202:1
route-target export 202:1
route-target import 202:1
!
address-family ipv4
exit-address-family
!
!
interface GigabitEthernet0/0/5.2
encapsulation dot1Q 2
vrf forwarding sda1
ip address 72.1.1.1 255.255.255.0
lisp mobility edge
!
interface GigabitEthernet0/0/5.3
encapsulation dot1Q 3
vrf forwarding sda1
ip address 72.1.2.1 255.255.255.0
lisp mobility edge
!
router lisp
locator-set edge
30.3.1.1 priority 1 weight 100
IPv4-interface Loopback0 priority 0 weight 10
auto-discover-rlocs
exit-locator-set
!
service ipv4
encapsulation vxlan
map-cache-limit 10000
itr map-resolver 124.1.1.1
itr
etr map-server 124.1.1.1 key lisp
etr
sgt
exit-service-ipv4
!
instance-id 1
dynamic-eid edge
database-mapping 72.1.1.0/24 locator-set edge
exit-dynamic-eid

```

```
!
service ipv4
  eid-table vrf sda1
  itr map-resolver 124.1.1.1
  itr
  etr map-server 124.1.1.1 key lisp
  etr
  sgt
  exit-service-ipv4
!
exit-instance-id
!
instance-id 2
dynamic-eid edge
  database-mapping 72.1.3.0/24 locator-set edge
  exit-dynamic-eid
!
service ipv4
  eid-table vrf sda2
  itr map-resolver 124.1.1.1
  itr
  etr map-server 124.1.1.1 key lisp
  etr
  exit-service-ipv4
!
exit-instance-id
!
exit-router-lisp
Branch1#
```




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>.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)